



nV System Configuration Guide for Cisco ASR 9000 Series Routers, IOS XR Release 6.7.x

First Published: 2020-01-20

Americas Headquarters

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

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

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

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

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

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

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

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

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

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

© 2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface vii

Communications, Services, and Additional Information vii

CHAPTER 1

Configuring the Satellite Network Virtualization (nV) System 1

Prerequisites for Configuration 2

Overview of Satellite nV System 3

Benefits of Satellite nV System 4

Cisco ASR 9000 Series Router Satellite nV Hardware Compatibility Matrix 5

Overview of Port Extender Model 6

Satellite System Physical Topology 8

Dual Home Mode 9

Delayed Switchback for nV Dual Hosts 12

Features Supported in the Satellite nV System 13

Inter-Chassis Link Redundancy Modes and Load Balancing 13

Satellite Discovery and Control Protocols 14

Satellite Discovery and Control Protocol IP Connectivity 14

ICL Fabric Port Monitoring 14

Dynamic ICL 16

Multiple ICCP Groups for Satellite nV System Network Topologies 17

Satellite nV System Access Port Bundles along with ICL Bundles 17

Use Cases for Satellite nV System Access Port Bundles 18

Port Partitioning 19

BFD over Satellite Interfaces 19

BNG Interoperability 20

Layer-2 and L2VPN Features 20

Layer-3 and L3VPN Features 20

Layer-2 and Layer-3 Multicast Features	20
Multicast IRB	20
Quality of Service	21
Time of Day Synchronization	21
Satellite Chassis Management	21
ARP Redundancy Support for Dual Head Mode	22
Ethernet Link OAM	22
802.3ah Loopback Support on Satellite nV System	22
Native Image and SMU Push Support for Cisco NCS 5000 Series Satellites	23
Support of 10x10G ICLs on Cisco NCS 5000 Series Satellites	24
Support for 40G ICLs on Cisco NCS 5000 Series Satellites	24
Support for MACsec Sessions on nV Satellite Access Interface	24
Restrictions of the Satellite nV System	24
Restrictions of the Cisco NCS 5000 Series Satellite	27
Satellite nV Usability Enhancements	29
nV Satellite Auto Image Upgrade	29
Simplified Access to Satellites	30
MPP Check Skip for Satellite Image Download with Auto IP	31
Implementing a Satellite nV System	31
Defining the Satellite nV System	31
Auto-IP	33
Configuring the Host IP Address	34
Configuring Inter-Chassis Links and IP Connectivity	35
Configuring the Inter-Chassis Links in a Dual Home Mode	37
Configuring the Satellite nV Access Interfaces	38
Configuring the Fabric CFM for L1 Physical ICL	39
Plug and Play Satellite nV Switch Turn up: (Rack, Plug, and Go installation)	41
Synchronous Ethernet (SyncE) Offload in nV System	43
Restrictions of SyncE Offload on the nV System	44
Restrictions of SyncE Offload on the Cisco NCS 5000 Series Satellite	44
Hub-and-Spoke Topology for SyncE Offload	45
Configuring SyncE Offload on Cisco ASR 9000 Series Router Hosts	45
Configuration Examples of SyncE Offload on Cisco ASR 9000 Series Router Hosts	46
Configuration to enable SyncE offload with Cisco NCS 5000 Series Satellite	48

Minimum Active Links on nV System	50
Soft Minimum Active Links for Dual Home Mode	50
Configuring Soft Minimum Active Links on Dual Home Mode	50
Hard Minimum Active Links for Single Host Mode	51
Configuring Hard Minimum Active Links on Single Host Mode	52
Upgrading and Managing Satellite nV Software	53
Image Upgrade for nV Satellite	53
Prerequisites	54
Installing a Satellite	55
Installing native images and SMUs for Cisco NCS 5000 Series Satellite	56
Monitoring the Satellite Software	58
Show Commands for Dual Home Mode	61
Dual Home Network Mode	61
Monitoring the Satellite Protocol Status	62
Monitoring the Satellite Inventory	63
Monitoring the Satellite Environment	65
Reloading the Satellite and Monitoring DOM Parameters	65
Port Level Parameters Configured on a Satellite	66
Loopback Types on Satellite Ports	67
Configuration Examples for Satellite nV System	68
Satellite System Configuration: Example	68
Satellite Global Configuration	68
ICL (satellite-fabric-link) Interface Configuration	68
Satellite Interface Configuration	69
Satellite Management Using Private VRF	70
Configuration of Satellite using Auto-IP	70
Satellite Configuration with Dual-Homed Hosts	71
Dual-Home for Multiple Satellites with Single Physical ICLs on Both Hosts and Satellites	71
Dual Home Configuration for SAT1	71
Dual Home Configuration for SAT2	72
Configuring nV Satellite Auto Image Upgrade: Example	73
Additional References	73
Image Upgrade for Cisco NCS 500x Satellite from Cisco IOS XR Software Release 6.0.0	73
Related Documents	74

Standards	74
MIBs	74
RFCs	75
Technical Assistance	75

CHAPTER 2

Configuring Multicast Offload on the Satellite nV System	77
Need for Multicast Offload	77
Scope and Prerequisites	78
Multicast Offload Terminology	80
Overview of Multicast Offload	80
Behavior of Multicast Offload During Satellite Topology Events	83
Difference between Layer 2 and Layer 3 (over BVI) core support	84
Uplink/Core Isolation Tracker	84
Multicast Offload Scalability	84
Multicast Offload Use Case Scenario	85
Restrictions for Multicast Offload	85
Configuring Satellite nV Multicast Offload	87



Preface



Note This product has reached end-of-life status. For more information, see the [End-of-Life and End-of-Sale Notices](#).

From Release 6.1.2 onwards, Cisco introduces support for the 64-bit Linux-based IOS XR operating system. However, support for nV Satellite System on the Cisco IOS XR 64-bit operating system is available only from the Release 6.2.1. Extensive feature parity is maintained between the 32-bit and 64-bit environments. Unless explicitly marked otherwise, the contents of this document are applicable for both environments. For more details on Cisco IOS XR 64 bit, refer to the [Release Notes](#) for Cisco ASR 9000 Series Routers, Release 6.1.2 document.

The Preface contains the following sections:

- [Communications, Services, and Additional Information, on page vii](#)

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

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



CHAPTER 1

Configuring the Satellite Network Virtualization (nV) System

This module describes Satellite Network Virtualization (Satellite nV) system configurations on Cisco ASR 9000 Series Aggregation Services Routers .

Table 1: Feature History for Configuring Satellite System

Release	Modification
Release 6.2.1	nV system basic feature support in hub and spoke topology on the host running Cisco IOS XR 64-bit operating system for Cisco ASR 9000v Satellite and Cisco NCS 5000 Series Satellites
Release 6.3.1	Support for small-frame-padding feature on satellite nV access interfaces. Support for hard minimum active links feature on bundle ICL of a single host mode. SMU install enhancements for Cisco NCS 5000 Series Satellites such as, installing updates, reference lists and native Cisco NCS 5000 Series Satellite SMUs. Added a keyword <code>on-first-connect</code> to upgrade <code>on-connect</code> command.
Release 6.3.2	Support for Provider Backbone Bridging Ethernet VPN (PBB-EVPN) on nV Satellite access interface bundle over ICL bundle.
Release 6.4.1	Support for MACsec sessions over Cisco ASR 9000v Satellite and Cisco NCS 5000 Series Satellite access interfaces. Support for upto 400 nV satellite access interfaces per NPU of the SE (Services Edge Optimized) variant of Cisco ASR 9000 High Density 100GE Ethernet line cards. Support for 40G ICL between the Cisco NCS 5000 Series Satellite and the Cisco ASR 9000 Series Router Host. Support on Cisco NCS 5000 Series Satellite for 10GE ports to operate at restricted speeds of 10/100/1000 M on qualified optics

Release 6.5.1	Support for SyncE offload on the nV System with the host running Cisco IOS XR 64-bit operating system. Support for the nV system on the Cisco ASR 9901 Router.
Release 6.6.1	Support for auto upgrade to a reference list of software packages on the Cisco NCS 5000 Series Satellite by specifying an optional reference keyword and reference name in the command: <code>upgrade [on-connect on-first-connect] <reference ref_name></code> .
Release 7.0.1	Cisco IOS XR 64-bit operating system supports nV system on the Cisco ASR 9000 4th Generation QSFP28 based dense 100GE Line Cards.

- [Prerequisites for Configuration, on page 2](#)
- [Overview of Satellite nV System, on page 3](#)
- [Benefits of Satellite nV System, on page 4](#)
- [Cisco ASR 9000 Series Router Satellite nV Hardware Compatibility Matrix, on page 5](#)
- [Overview of Port Extender Model, on page 6](#)
- [Satellite System Physical Topology, on page 8](#)
- [Dual Home Mode, on page 9](#)
- [Features Supported in the Satellite nV System, on page 13](#)
- [Restrictions of the Satellite nV System, on page 24](#)
- [Satellite nV Usability Enhancements, on page 29](#)
- [Implementing a Satellite nV System, on page 31](#)
- [Synchronous Ethernet \(SyncE\) Offload in nV System, on page 43](#)
- [Minimum Active Links on nV System, on page 50](#)
- [Upgrading and Managing Satellite nV Software, on page 53](#)
- [Configuration Examples for Satellite nV System, on page 68](#)
- [Additional References, on page 73](#)

Prerequisites for Configuration

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Before configuring the Satellite nV system, you must have these hardware and software installed in your chassis:

- Hardware (Host):
 - Cisco ASR 9000 Series Aggregation Services Routers, the Cisco ASR 9000 High Density 100GE Ethernet line cards, the Cisco ASR 9000 4th Generation QSFP28 based dense 100GE line cards as the location of Inter Chassis Links. Cisco ASR 9000 Ethernet Line Cards can co-exist in the Satellite nV System but cannot be used for Satellite ICLs and also with ISM/VSM.
- Hardware (Satellite) :
 - Cisco ASR 9000v Satellite (Cisco ASR9000v-V2)
 - Cisco NCS 5002 Series Satellite (PID: NCS-5002-ACSR, NCS-5002-FN-BK, NCS-5002-FLT-BK, NCS-5002-FN-FR, NCS-5002-FLT-FR)

- Cisco NCS 5001 Series Satellite (PID: NCS-5001-ACSR, NCS-5001-FN-BK, NCS-5001-FLT-BK, NCS-5001-FN-FR, NCS-5001-FLT-FR)

- Cisco IOS XR Software Release 6.2.1 or later must be installed on the host.



Note The Cisco ASR 9000 Series Routers running Cisco IOS XR 64-bit operating system supports the nV Satellite System in hub and spoke topology from Cisco IOS XR Software Release 6.2.1 onwards.

The Cisco ASR 9000 Series Routers running Cisco IOS XR 64-bit operating system supports the Cisco ASR 9000 High Density 100GE Ethernet line cards, the Cisco ASR 9000 4th Generation QSFP28 based dense 100GE line cards.

For more information on other hardware requirements and list of TMG optics supported, see *Cisco ASR 9000 Series Aggregation Services Router Hardware Installation Guide* and *Cisco ASR 9000 Series Aggregated Services Router Satellite Systems Installation Guide*.

Overview of Satellite nV System

The Cisco ASR 9000 Series Router Satellite Network Virtualization (nV) service or the Satellite Switching System enables you to configure a topology in which one or more satellite switches complement one or more Cisco ASR 9000 Series routers, to collectively realize a single virtual switching system. In this system, the satellite switches act under the management control of the routers. The complete configuration and management of the satellite chassis and features is performed through the control plane and management plane of the Cisco ASR 9000 Series Router, which is referred to as the host.



Note Cisco ASR 9006, Cisco ASR 9010, Cisco ASR 9901, Cisco ASR 9904, Cisco ASR 9906, Cisco ASR 9910, Cisco ASR 9912, Cisco ASR 9922 Series Routers with Modular Services Line Card can also be used as hosts in the Satellite nV System.

Interconnection between the Cisco ASR 9000 Series Router and its satellites is through standard Ethernet interfaces.

In general, the type of interface used on the host is decided on the basis of the satellite device used as shown in the figure.




Note 1-Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces can be used as ICL ports. The Cisco ASR 9000 High Density 100GE Ethernet line cards support 1GE, 10GE, 40GE and 100GE ICL ports. The Cisco ASR 9000 4th Generation QSFP28 based dense 100GE line cards support 10GE and 100GE ICL ports.



Note With the Cisco NCS 5000 Series Satellite, only 100G ICL is supported from Cisco IOS XR Software Release 6.0.0 whereas even 10GigE dynamic ICL is supported from Cisco IOS XR Software Release 6.0.1. The Cisco NCS 5000 Series Satellites support 40G ICL from Cisco IOS XR Software Release 6.4.1.

Figure 1: Cisco ASR 9000 Series Router nV Switching System



Note Though the above figure shows Nx10GigE links, Satellite nV System also supports Nx1GigE and Nx100GigE links.

This type of architecture can be realized in a carrier Ethernet transport network, with the satellite switches used as either access switches, or pre-aggregation and aggregation switches. These switches feed into an edge router, such as the Cisco ASR 9000 Series Router where more advanced Layer 2 and Layer 3 services are provisioned. The network topology depicted in the figure is called the Hub and Spoke network topology.

You can also utilize this model in a Fiber To The Business (FTTB) network application, where business internet and VPN services are offered on a commercial basis. Further, it can also be used in other networks, such as wireless or Radio Access Network (RAN) backhaul aggregation networks.

Benefits of Satellite nV System

The Cisco ASR 9000 Series satellite nV system offers these benefits:

- 1. Extended port scalability and density** - You can create a virtual line card with more than 400 physical Gigabit Ethernet ports. There is a significant increase of Ethernet port density in the resulting logical Cisco ASR 9000 Series Router. For example, a single 24-port Ten Gigabit Ethernet line card on the Cisco ASR 9000 Series Router could integrate up to 24 satellite switches each with 44 GigE ports; this results in an effective port density of 1056 Gigabit Ethernet ports for each Cisco ASR 9000 Series Router line card slot. In other configurations, even higher port density can be achieved. This is beneficial because the Cisco ASR 9000 Series Router has a per-slot non blocking capacity of up to 400 Gbps (with appropriate RSPs) and there is no other way of physically fitting hundreds of gigabit ethernet ports/ SFPs on the face plate of a single Cisco ASR 9000 Series line card. As a result, in order to utilize the full capacity of an Cisco ASR 9000 Series line card, it is necessary to physically separate out the ethernet ports, while

maintaining logical management control. This would appear as if all ports were physically on a single large line card of the Cisco ASR 9000 Series Router.



Note Port scalability can be achieved for 10GigE access ports and 100GigE uplinks using the Cisco NCS 5000 Series satellites .

2. **Reduced cost** - All the edge-routing capabilities and application features of the Cisco IOS XR Software are available on low cost access switches.
3. **Reduced operating expense** - You can upgrade software images, and also manage the chassis and services from a common point. This includes a single logical router view, single point of applying CLI or XML interface for the entire system of switches, a single point of monitoring the entire system of switches and a single point of image management and software upgrades for the entire system.
4. **Enhanced feature consistency** - All the features on the regular GigE ports and 10GigE ports of Cisco ASR 9000 Series Router are also available on the access ports of a satellite access switch in a functionally identical and consistent manner. The typical application of a satellite system would be in the access and aggregation layers of a network. By integrating the access switches along with the aggregation or core switch, you can ensure that there are no feature gaps between the access switch and the aggregation or core switch. All features, such as carrier ethernet features, QoS and OAM, function consistently, from access to core, because of this integrated approach.
5. **Improved feature velocity** - With the satellite solution, every feature that is implemented on the Cisco ASR 9000 Series Router becomes instantly available at the same time in the access switch, resulting in an ideal feature velocity for the edge switch.
6. **Better resiliency** - The nV satellite solution enables better multi-chassis resiliency, as well as better end-to-end QoS. For more information on QoS capabilities, see *Cisco ASR 9000 Series Aggregation Services Router QoS Configuration Guide* .

Cisco ASR 9000 Series Router Satellite nV Hardware Compatibility Matrix

The following table lists the IOS XR 64 bit Satellite nV hardware compatibility matrix.

Table 2: IOS XR 64 Bit Satellite nV Hardware Compatibility Matrix

Line Cards / Chassis / Port Expansion Cards	Cisco ASR 9000v Satellite Supported Version	Cisco NCS 5000 Series Satellite Supported Version
A9K-MPA-4X10GE on MOD400	6.2.1	6.2.1
A9K-MPA-20X10GE on MOD400	6.2.1	6.2.1
A9K-MPA-20X1GE on MOD400	6.2.1	–
A9K-MPA-4X10GE on MOD200	6.2.1	6.2.1
A9K-MPA-20X10GE on MOD200	6.2.1	6.2.1

Line Cards / Chassis / Port Expansion Cards	Cisco ASR 9000v Satellite Supported Version	Cisco NCS 5000 Series Satellite Supported Version
A9K-MPA-20X1GE on MOD200	6.2.1	–
A9K-4X100G-SE	6.2.1	6.2.1
A9K-4X100G-TR	6.2.1	6.2.1
A9K-8X100G-SE	6.2.1	6.2.1
A9K-8X100G-TR	6.2.1	6.2.1
A9K-MPA-1X100GE on MOD200	–	6.3.1
A9K-MPA-1X100GE on MOD400	–	6.3.1
A9K-MPA-2X100GE on MOD200	–	6.3.1
A9K-MPA-2X100GE on MOD400	–	6.3.1
A9K-MPA-8X10GE on MOD200	6.3.2	6.3.2
A9K-MPA-8X10GE on MOD400	6.3.2	6.3.2
A9K-48X10GE-1G-SE	6.5.3	6.5.3
A9K-48X10GE-1G-TR	6.5.3	6.5.3
A9K-24X10GE-1G-SE	6.5.3	6.5.3
A9K-24X10GE-1G-TR	6.5.3	6.5.3
A99-48X10GE-1G-SE	6.5.3	6.5.3
A99-48X10GE-1G-TR	6.5.3	6.5.3
A9K-8X100GE-X-TR	7.0.1	7.0.1
A9K-16X100GE-TR	7.0.1	7.0.1
A9K-32X100GE-TR	7.0.1	7.0.1
A99-16x100-X-SE	7.0.1	7.0.1


Overview of Port Extender Model

In the Port Extender Satellite switching system also called as Hub and Spoke model, a satellite switch is attached to its host through physical ethernet ports.

The parent router, Cisco ASR 9000 Series Router is referred to as the host in this model. From a management or a provisioning point of view, the physical access ports of the satellite switch are equivalent to the physical ethernet ports on the Cisco ASR 9000 Series Router. You do not need a specific console connection for managing the Satellite Switching System, except for debugging purposes. The interface and chassis level

features of the satellite are visible in the control plane of Cisco IOS XR software running on the host. This allows the complete management of the satellites and the host as a single logical router.

Figure 2: Port Extender Satellite Switching System



In this model, a single Cisco ASR 9000 Series Router hosts two satellite switches, SAT1 and SAT2, to form an overall virtual Cisco ASR 9000 switching system; represented by the dotted line surrounding the Cisco ASR 9000 Series Router, SAT1, and SAT2 in the Figure.

This structure effectively appears as a single logical Cisco ASR 9000 Series Router to the external network. External access switches (A1, A2 and A3) connect to this overall virtual switch by physically connecting to SAT1 and SAT2 using normal ethernet links. The links between the satellite switches and the Cisco ASR 9000 Series Router are ethernet links referred to as the Inter-Chassis Links (ICL). The Cisco ASR 9000 Series

Router is referred to as the Host. When there is congestion on the ICLs, an inbuilt QoS protection mechanism is available for the traffic.



Note SAT1, SAT2, and the host Cisco ASR 9000 Series Router need not be located in the same geographic location. This means that the ICLs need not be of nominal length for only intra-location or intra-building use. The ICLs may be tens, hundreds, or even thousands of miles in length, thereby creating a logical satellite switch spanning a large geography.

Satellite System Physical Topology

The satellite nV system supports Hub and Spoke network topology in single and dual home modes.

Hub and Spoke network topology (or the port extender satellite switching system) is the basic topology for Satellite nV system and is covered in the section [Overview of Port Extender Model, on page 6](#). Dual Home, the Advanced Satellite nV network mode is covered in the section [Dual Home Mode, on page 9](#).

Hub and Spoke network topology allows a physical Ethernet MAC layer connection from the satellite to the host. This can be realized using a direct Ethernet over Fiber or Ethernet over Optical transport (such as Ethernet over a SONET/ SDH/ CWDM/ DWDM network).

This topology also allows a satellite switch to geographically be at a different location, other than that of the host. There is no limit set for the distance, and the solution works even when the satellite is placed at a distance of tens, hundreds, or even thousands of miles from the host.



Note For tunable DWDM optics, the same wavelength needs to be configured on both host and the satellite.

For the ASR9000v series satellite, if this is the first ICL, then the following CLI command should be used from a direct console connection on the ASR9000v series satellite to set the wavelength channel number on the satellite side:

```
test dwdm wavelength set port_number_as_on_faceplate wavelength_channel_number
```

For subsequent ICLs, the satellite can be accessed via telnet from the host to configure this command for those ICLs.

For ASR9k host side configuration, please refer to "Configuring Dense Wavelength Division Multiplexing Controllers" chapter in the *Interface and Hardware Component Configuration Guide for Cisco ASR 9000 Series Routers*.




Note Hub and Spoke topology with single and dual home modes use untagged LLC SNAP packet encapsulation for SDAC Discovery protocol and untagged TCP for SDAC control protocol packets. Once the satellite is discovered and connected, the data packets are then encapsulated in 802.1ad between the ASR9000 host and the nV satellite.

Dual Home Mode

In the dual home mode, two hosts are connected to a satellite through the Satellite Discovery And Control (SDAC) Protocol. The SDAC Protocol provides the behavioral, semantic, and syntactic definition of the relationship between a satellite device and its host.

Both these dual-homed hosts act in the active/standby mode for the satellite. The standby host takes control of the satellite only when the active host is down. The two hosts can leverage the ICCP infrastructure to provide redundant Layer 2 and Layer 3 services for Satellite Ethernet interfaces. The network traffic is switched through the active host. In case of connection loss to the active host due to various types of failure such as cut cable and host or client connection interface failure, the standby host becomes the active host and the active host becomes the new standby host. The hosts communicate with each other using ORBIT/ICCP protocols. The Satellite Discovery and Control (SDAC) session is established from both the active and standby hosts, and it is only the traffic flows that are in the active/standby mode.

Figure 3: Dual Home Mode Architecture



Features of the Dual Home Mode

These are some of the enhanced features offered by the dual home mode:

- **Shared control for chassis functionality:** Chassis control functionality which includes software upgrade, chassis reload, and environment monitoring is completely shared by all hosts connected to the Satellite. Both the hosts get equal access to the information, and have full control of any available actions. As a result, a disruptive change initiated by one host, such as an upgrade is noticed by the other host as well. This means that there is no segregation of the chassis functionality and provides multiple views to the same information.
- **Active/Standby determination:** Active/Standby determination is controlled by the hosts. They exchange the pertinent information through ORBIT protocol, which includes electing a priority selection algorithm to use. This algorithm determines the factors that are taken into account when calculating priority information. The hosts then each send a single numerical priority value to the Satellite. The Satellite only picks the lowest-host priority value, and forwards data to that host. Independently, the hosts make the

same determination, and the traffic flows bi-directionally between the Active host and the Satellite. The hosts take a number of parameters into account when calculating the priority value, including the user-configured priority, from the host to the Satellite, and a tie-break of the chassis MAC for each host.

Cisco IOS XR Software uses these parameters to calculate the priority, where each item is more important than any of the subsequent ones. This means that the subsequent parameters are only taken into account, if the higher-priority items are identical across the two hosts.

- **Connectivity** – Indicates whether the Host and Satellite can currently exchange data.
- **PE isolation** – Indicates that if the PE is isolated, then it defers to the other host.
- **Minimum Links** – Indicates that if the number of active links is less than the configured value in bundled ICL, then it defers to the other host. For more information, see [Soft Minimum Active Links for Dual Home Mode, on page 50](#).
- **Configured Priority** – This is as early as possible to allow the greatest flexibility for the operator.



Note

The host priority switchover functionality does not function when one of the hosts is running an IOS XR release older than 6.1.x and the other host is running 6.1.x or later software releases. In such cases, you must disconnect the host running the older IOS XR release from the satellites. After a successful disconnection, the satellites switch over to the remaining hosts. You can also upgrade the disconnected host to the latest software release.

- **Parity of the Satellite ID** – This is used as a late tie breaker to provide some load balancing across two Hosts with numerous hub-and-spoke Satellites, in which the even-numbered Satellites prefer one host, while the odd-numbered Satellites prefer the other host.

On a tie-breaker of all the previous priorities, it falls back to the Primary host, which is the one with the lowest chassis MAC address based on byte size.

- **Support for seamless Split Brain handling:** A Split Brain is a condition in which the two hosts are no longer able to communicate with each other, but each of them can still communicate with the Satellite device. This scenario can be hazardous, because the devices can reach different conclusions, resulting in traffic loss, or loops and broadcast storms.

The Satellite protocol has these features to cope with such a situation:

- When connected to each other, the two hosts publish a shared System MAC. This allows the Satellites to recognize probes from what appear to be different hosts, but in fact come from a paired set of hosts.
- Whenever a host-to-host connection is lost, each peer publishes the Chassis MAC as the System MAC. This operation is seamless and does not require a reset of the state machines, and hence causes no traffic loss. This allows the Satellite to react, most likely by dropping its connection to the standby host.
- Whenever the connection is restored, the hosts again start publishing the System MAC seamlessly and allowing the Satellite to restore functionality to the standby host.
- If the host-to-host connection is lost while the host is PE-isolated, it immediately drops discovery to the satellite. This ensures that the satellite uses the host with an available path to the core, if one exists.

- **Separate records of interface counters of satellite access ports on either host:** - The SDAC protocol that runs between the hosts and satellite reports only the difference in statistics to each host, that is, the number of packets/bytes since the last update. Clearing the counters on either host clears only the local copy of the statistics of the satellite access ports. It doesn't update the interface counters on the satellite itself and the other host won't get notified that the counters have been reset on the peer. So, it's expected that these counters may be different on each host if you have cleared the counters at different times or if the host-to-satellite connections came up at different times.

To clear the satellite interface counters on the host, execute the command **clear counters** on the host router.

To view the interface counters of satellite access ports, execute the command **show interface satellite-access-port** on the host router.

After clearing counters on Host 1, the statistics for satellite access ports on both hosts are different because only the local copy of the satellite interface counters on Host 1 got reset.

```
On Host 1:
Router#clear counters
Clear "show interface" counters on all interfaces [confirm]
Router#show int GigabitEthernet160/0/0/0
GigabitEthernet160/0/0/0 is up, line protocol is up
  Interface state transitions: 3
  Hardware is GigabitEthernet/IEEE 802.3 interface(s), address is 8478.ac07.6155 (bia
8478.ac07.6155)
  Internet address is 10.10.1.1/24
  MTU 1514 bytes, BW 1000000 Kbit (Max: 1000000 Kbit)
    reliability 255/255, txload 190/255, rxload 127/255
  Encapsulation ARPA,
  Full-duplex, 1000Mb/s, SX, link type is force-up
  output flow control is off, input flow control is off
  Carrier delay (up) is 100 msec, Carrier delay (down) is 100 msec
  loopback not set,
  Last link flapped 23:16:57
  ARP type ARPA, ARP timeout 04:00:00
  Last input 00:00:00, output 00:00:00
  Last clearing of "show interface" counters 00:00:04
  5 minute input rate 499313000 bits/sec, 61311 packets/sec
  5 minute output rate 748970000 bits/sec, 91966 packets/sec
    327104 packets input, 332990972 bytes, 0 total input drops
  0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 1 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
490655 packets output, 499485772 bytes, 0 total output drops
  Output 0 broadcast packets, 163552 multicast packets
  0 output errors, 0 underruns, 0 applique, 0 resets
  0 output buffer failures, 0 output buffers swapped out
  0 carrier transitions
```

Satellite access port counters on Host 2 are not reset and so different from the values on Host 1:

```
On Host 2:
Router#show int GigabitEthernet160/0/0/0
GigabitEthernet160/0/0/0 is up, line protocol is up
  Interface state transitions: 3
  Hardware is GigabitEthernet/IEEE 802.3 interface(s), address is 8478.ac07.6155 (bia
8478.ac07.6155)
  Internet address is 10.10.1.1/24
  MTU 1514 bytes, BW 1000000 Kbit (Max: 1000000 Kbit)
    reliability 255/255, txload 191/255, rxload 127/255
  Encapsulation ARPA,
```

```

Full-duplex, 1000Mb/s, SX, link type is force-up
output flow control is off, input flow control is off
Carrier delay (up) is 100 msec, Carrier delay (down) is 100 msec
loopback not set,
Last link flapped 23:17:03
ARP type ARPA, ARP timeout 04:00:00
Last input 00:00:00, output 00:00:00
Last clearing of "show interface" counters 02:32:13
5 minute input rate 499495000 bits/sec, 61333 packets/sec
5 minute output rate 749243000 bits/sec, 92000 packets/sec
 541290577 packets input, 551031306258 bytes, 0 total input drops
 0 drops for unrecognized upper-level protocol
Received 0 broadcast packets, 2628 multicast packets
 0 runts, 0 giants, 0 throttles, 0 parity
 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
807585148 packets output, 822119157464 bytes, 0 total output drops
Output 0 broadcast packets, 269196783 multicast packets
 0 output errors, 0 underruns, 0 applique, 0 resets
 0 output buffer failures, 0 output buffers swapped out
 0 carrier transitions

```

Delayed Switchback for nV Dual Hosts

In a highly scaled dual host nV system, a switchback to the original active host (after it is back online) can sometimes lead to traffic outages if the switchback happens before that host is ready to forward traffic. The delayed switchback feature sets a hold-time, which is the duration in seconds a higher-priority standby device that has just started waits, before preempting the current active host. Before allowing a switchover to the active host, the delayed switchback feature checks whether the current standby host pending switchover has met the following criteria:

- Creation of satellite port instances and application of basic user configuration on the satellite ports in the hardware forwarding path on the host that is to take up the active role (This does not include the verification of access port feature programming in that hardware.)
- Full synchronization of information with the satellites that would failover, to minimize the chances of any outage post failover
- Hold timer of 5 minutes to allow for the programming of related features that set up the overall forwarding path

The delayed switchback feature is primarily useful when switchback needs to be delayed after a host has reloaded and its various modules are still initializing. Also, in a scenario where the satellite is discovered by one host for which another host acts as the active host and successfully forwards traffic, delayed switchback helps ensure that the redundancy mechanism of dual host failover does not end up causing traffic outages by a premature failover.



Note

- Even when the partner device is unable to forward traffic and cannot continue as an active host, a switchover can still occur.
- For ICL partitioned systems, the delayed switchback is guaranteed to consider only one partition. Issues in forwarding path for remaining partitions may not be considered for the hold timer.
- Currently, the feature applies only at a per-satellite level.

You can use the **show nv satellite redundancy** command to check if host is active or standby, figure out priorities for the local and partner device and understand different priority levels.

For details of this command, see the *Cisco ASR 9000 Series Aggregation Services Router nV System Command Reference*.

Features Supported in the Satellite nV System

This section provides details of the features of a satellite system.

Inter-Chassis Link Redundancy Modes and Load Balancing

The Cisco ASR 9000 Series Satellite system supports these redundancy modes:

- **Non-redundant inter-chassis links mode** - In this mode, there is no link level redundancy between inter-chassis links of a satellite.
- **Redundant inter-chassis links mode** - In this mode, the link level redundancy between inter-chassis links are provided using a single link aggregation (LAG) bundle.

In the redundant ICL mode, the host device load-balances the host to satellite traffic between the members of the ICL bundle using a simple modulo of access port number and number of active ICL bundle member links.

The satellite devices use a similar hashing function to load-balance satellite to host traffic between the members of the ICL bundle.

Each device independently implements the hash function to pin the access ports to the ICL member links for a traffic direction. L2 or L3 header contents from the packet are not used for flow based hashing of the access-port to the ICL bundle member links. This ensures that all packets use a single ICL for a given satellite access-port per direction. Although, each device might pick a different member for egress as they independently implement the hash function. As a result, the actions applied for QoS and other per-direction features still consider all the packets as belonging to a single physical port.

In order to find out the ICL bundle member link to which a given satellite access port is pinned to, check the **Base Interface** in the show command, **show sits interface satellite-access-port**

```
Router# show sits interface tenGigE 600/0/0/70

Interface Name           : TenGigE600/0/0/70
Interface IFHandle       : 0x00002ce0
Base Interface         : HundredGigE0/3/0/4
Base Interface IFHandle  : 0x0a0003c0
LAG Hash Index          : 72
Published Base interface : HundredGigE0/3/0/4
Published Base ifh      : 0x0a0003c0
Published LAG Hash      : 72
Sat Interface State in DB: Caps Publish Success
```



Note If a satellite system is operating in redundant ICL mode, then Ethernet OAM features are not supported on the access ports of that satellite.

For more details on QoS application and configuration on ICLs, see *Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide*.

Satellite Discovery and Control Protocols

Cisco's proprietary discovery and control protocols are used between the satellite switches and the host Cisco ASR 9000 Series Router devices, to handle discovery, provisioning, and monitoring of the satellite devices from the host Cisco ASR 9000 Series Satellite System in-band over the ICLs. The Satellite Discovery And Control (SDAC) Protocol provides the behavioral, semantic, and syntactic definition of the relationship between a satellite device and its host.

Satellite Discovery and Control Protocol IP Connectivity

The connectivity for the SDAC protocol is provided through a normal in-band IP routed path over the ICLs using private and public IP addresses appropriate for the carrier's network.

You can configure a management IP address on the host CLI for each satellite switch and corresponding IP addresses on the ICLs. You can select addresses from the private IPv4 address space (for example, 10.0.0.0/8 or 192.1.168.0/24) in order to prevent any conflict with normal service level IPv4 addresses being used in the IPv4 FIB. You can also configure a private VRF that is used for only satellite management traffic, so that the IP addresses assigned to the satellites can be within this private VRF. This reduces the risk of address conflict or IP address management complexity compared to other IP addresses and VRFs that are used on the router.



Note Auto-IP is the recommended mode of configuration. For more information on Auto-IP mode, see [Auto-IP, on page 33](#).

ICL Fabric Port Monitoring

This feature enables the host to create virtual fabric port interfaces locally to represent each of the fabric facing interfaces on the satellite.

Interfaces are created when either:

- Candidate fabric port configuration is applied to the satellite.
- The satellite informs the host that the port exists on the candidate fabric port channel.
- The satellite informs the host that the port exists and is in Active state on the topology channel.

This allows you to monitor the information related to the links within fabric ports and get their Layer 1 and Layer 2 parameters from the host. The satellite sends the ICL interface details to the host through the topology channel messages. You can run SNMP queries on these interfaces. The **show interface nvFabric-TenGigE/GigE interface** provides L1 details and L2 statistics for these interfaces. The host uses the information received on the topology channel to create the virtual fabric port interfaces locally. The topology channel TLVs are modified such that they provide the detail port definition in the format of *slot/subslot/port*. Currently, there is no TLV in Layer 1 and Layer 2 channel for MTU to be exchanged from host to satellite. Hence, the host just hard codes it to 9212.



Note In the case of Cisco ASR 9000v, nVFabric-TenGigE interface is created irrespective of whether 1 GE or 10GE ICL is used.

For Cisco NCS 5000 Series Satellites nVFabric-HundredGigE and nVFabric-TengGigE interfaces are created for 100GE and 10GE ICLs.

For Cisco NCS 5000 Series Satellites nVFabric-HundredGigE interface is created irrespective of whether 40GE or 100GE ICL is used. And nVFabric-TengGigE interface is created for 10GE ICLs.

The virtual interfaces are deleted if they are not configured via the candidate fabric port configuration and:

- If the satellite informs the host that the port is no longer used as a fabric link.
- If the satellite is removed from the network topology.
- If the topology channel moves to “Down” state for other reasons, such as authentication failing due to an incorrect serial ID.

Interfaces created due to candidate fabric port configuration will never be deleted and will persist over the above conditions.



Note The **show nv satellite topology** command is updated to use the name of the Satellite nV Fabric interface.

```
RP/0/0/CPU0:Router#show nv satellite topology

GigabitEthernet0/1/0/11
-----
Redundancy-Group: 10
Ring Whole: False
Discovery status: Running
Satellites:
  Satellite 200 (BVID 2)
  -----
  Received Serial Number: CAT1709U06V
  MAC address: 7cad.7404.6028
  Satellite fabric links:
    GE/0/0/11 (nVFabric-GigE200/0/0/11) (Remote ID: 0xb):
      Host (GigabitEthernet0/1/0/11)
    GE/0/0/10 (nVFabric-GigE200/0/0/10) (Remote ID: 0xa):
      Remote port not yet identified
  -----

RP/0/RSP1/CPU0:vkg1#sh nv sat topology
TenGigE0/3/1/2
-----
Redundancy-Group: 10
Ring Whole: False
Discovery status: Running
Satellites:
  Satellite 100 (BVID 2)
  -----
  Received Serial Number: CAT1721U0ED
  MAC address: 8478.ac03.6404
  Satellite fabric links:
    TenGE/0/0/46 (nVFabric-TenGigE100/0/0/46) (Remote ID: 0x3):
```

```
Host (TenGigE0/3/1/2)
TenGE/0/0/44 (nVFabric-TenGigE100/0/0/44) (Remote ID: 0x1):
Remote port not yet identified
```

Dynamic ICL

The Dynamic ICL feature allows you to configure ICL on the 1 GigE satellite ports 42 and 43 apart from the designated four 10GigE ports on the Cisco ASR 9000v and Cisco ASR 9000v2 satellites. The Cisco ASR 9000v satellites can have, by default, up to 6 potential ICLs. These are the four 10 GigE ports(44-47) and two 1 GigE ports (42-43) and 42 fixed 1 GigE access ports (0-41). The two 1 GigE ports (42-43) that are considered as potential ICLs are special ports that can be used either as ICLs or as access ports. When these two ports (42-43) are configured as access ports, they are not considered as potential ICLs thereafter and therefore behave like any other 1 GigE access ports till the next reload. When the satellite reloads and comes back up, those two ports again become as potential ICLs. When the satellite gets discovered and connects back to the host and if the host cross-link configuration still has the ports 42 and 43 as access ports, then these ports again switch to the access port mode and exhibit the behavior like any other 1 GigE access ports. This feature is supported on all the existing satellite topologies.

In case of dual head satellite mode, if any of these ports (42-43) is configured as access port in at least one of the host and if that host becomes active, then that port are not considered as potential ICLs anymore till next satellite reload. Also, the two ports (42 and 43) when used as access ports, lose the link integrity when the satellite reloads. The remote devices connected to these two ports stay up momentarily just after the satellite reload and till the satellite gets connected to the host and receives the cross-link configuration again.

You cannot bundle Dynamic ICL ports that are of dissimilar types. For example, you cannot bundle 1GigE and 10GigE ports.

Dynamic ICL is supported on 10G ports on Cisco NCS 5000 Series satellite either as an access port or an ICL fabric port based on user configurations done through the configurable fabric port functionality. By default, the highest 2x10G ports on both Cisco NCS 5001 (38 and 39) and Cisco NCS 5002 (78 and 79) are turned on as the ICL ports in the satellite mode (only the highest 10G port turns on for plug and play in factory reset / default mode). When the links come up over these default fabric ports and the satellite is connected, you can configure a list of candidate fabric ports. This gives the flexibility of using ports other than the default fabric ports as additional ICLs.

The default fabric ports (highest 2x10G ports) and the configured fabric ports lose link-integrity even if they are later configured as access ports until the satellite discovery goes down next. The default fabric ports permanently lose link integrity as they come up as potential ICLs every time the cross-link mapping to an ICL goes away. This is a necessary trade off to have certain ports available for zero touch plug and play. Therefore, at instances where link integrity is crucial, the default fabric ports, either must not be connected to such peers or must be connected over non-revertive switch-over topologies.

Configurable fabric ports

All satellite hardware variants have default set of fabric ports (permanent as well as dynamic) that can be used to connect them as ICLs. These include 4x10G ports and the highest 2x1G ports (port 42,43) for Cisco ASR 9000v. For Cisco NCS 5000 series satellites, these include 4x100G ports and the highest 2x10G ports. The highest 1G/10G ports are dynamic ICLs that can either be used as access ports or as ICLs based on whether they have an active mapping to another ICL at a given time or not.

Apart from these ports, the host also allows additional 10G ports on Cisco NCS 5000 series satellites to be configured as fabric ports. These are configurable as Candidate Fabric Ports on the host under the nV satellite global configuration. When the configured list is accepted, the satellite reports these interfaces to the host (including the default/permanent candidate fabric ports) and subsequently sends change notifications when that set changes or during re-synchronization with the host. The satellite might choose to reject some of these

configured candidate fabric ports if they are already being used as active access ports. As in the case of the normal Satellite nV System fabric ports, these interfaces are deleted if they have no configuration presence and become inactive in the fabric topology, either as notified by the satellite over the topology channel, or if the satellite is entirely undiscovered by the host. The information about the nV fabric port, and specifically fabric port statistics, does not persist across deletion and re-creation of an nV Fabric interface with the same name.

An overlapping configuration using the same port as both access port and candidate fabric port is not recommended. However, when such a configuration is used, the behavior depends on the chronological ordering of whether the port was first discovered as an ICL or whether it got mapped as an access port to another active ICL. While a configuration is rejected to map the port in one mode when the port is already active in the other mode, there might be cases where the candidate fabric port is configured and is waiting to be discovered. If it is also mapped as an access port to another ICL, then the port becomes an active access port and cannot be discovered thereafter, until the access port mapping goes away.

The list of permanent and configured fabric ports, including their acceptance status is displayed in the status of show commands under *Monitoring the Satellite Software* section.



Note With respect to the existing dynamic ICLs, link integrity is not supported for ports configured as fabric ports. Therefore, if they are also used as access ports, the ports will come up as soon as the access port mapping goes away. This is required for listening to discovery when they act as dynamic fabric ports.

Multiple ICCP Groups for Satellite nV System Network Topologies

The Multiple ICCP group feature allows you to configure multiple redundancy groups on a single host. These redundancy groups function completely independently of each other. The primary function of this feature is to support multiple satellites that are dual-homed to different but overlapping pairs of Cisco ASR 9000 Series Routers.

In a multiple ICCP redundancy group setup, failovers can happen similar to the case of a basic Dual Home setup. See the [Dual Home Mode, on page 9](#) for information on support for seamless split brain handling.

Satellite nV System Access Port Bundles along with ICL Bundles

This feature provides redundancy for Inter-Chassis Links (ICLs) along with redundancy for satellite access ports. These are different deployment models for this feature:

- Bundles of Satellite interfaces with ICL links as non-redundant links.
- Satellite interfaces with ICL links as redundant links (ICL links in a bundle).
- Bundles of Satellite interfaces with ICL links as redundant links (ICL links in a bundle).

These are the topologies and features supported when bundles of satellite interfaces exist with redundant ICL bundles:


1. Bundles of Satellite interfaces with ICL redundant links to single host.
2. Bundles of Satellite interfaces with ICL redundant links to two hosts, which are connected by ICCP to provide node redundancy.
3. Bundles of Satellite interfaces with many ICL redundant links to single host.

4. Sub-interfaces over Bundles of Satellite interfaces with ICL redundant links to single host.
5. Sub-interfaces over Bundles of Satellite interfaces with ICL redundant links to two hosts.
6. LACP over virtual Satellite access member interfaces.
7. BFD configuration of BLB.
8. L2VPN support for both VPWS and VPLS connectivity.
9. IPv4 and IPv6 protocols with and without L3VPNs.
10. Netflow, ACL, and routing protocols such as BGP, OSPF, ISIS.
11. Dynamic Host Configuration Protocol(DHCP).
12. Host QoS and QoS offload are supported.
13. PBB-EVPN (Provider Backbone Bridging Ethernet Virtual Private Network) is supported with single homed, single active, and all active topologies with satellite access.

Use Cases for Satellite nV System Access Port Bundles

These are different use cases of satellite topologies with Satellite nV System access port bundles along with redundant ICL.

Figure 4: Multiple Access Bundles over single ICL Bundle to Single Host



In this case, different members of the access bundle can be configured over different ICL bundle links.

Figure 5: Multiple Access Bundles over many ICL bundles to Single Host






Figure 6: Multiple Access Bundles over many ICL bundles to Dual-Homed Hosts



Note In the figures 8, 9, and 10, the satellite can also be Cisco NCS 5000 Series Router in place of Cisco ASR 9000v.

In this case where MC-LAG solution is provided to CPE device, wherein each member in the CPE bundle is connected to different satellite box. The satellite is single homed to a Cisco ASR 9000 host.

Figure 7: Access Bundle over ICL Bundles to two different hosts in CPE dual homing



Port Partitioning

The Cisco ASR 9000 Series Satellite system allows you to split or partition the satellite ethernet ports across multiple ICL interfaces. You can split the satellite ports between 4 ICLs in Cisco ASR 9000v satellite. You can split the satellite ports between upto 10 ICLs that can be either 100G or 10G ICLs in Cisco NCS 500x Series satellites. See *Support of 10x10G ICLs on Cisco NCS500x satellites* section for additional information.

BFD over Satellite Interfaces

Bidirectional Forwarding Detection (BFD) over satellite interfaces feature enables BFD support on satellite line cards. Satellite interfaces are known as virtual (bundle) interfaces. BFD uses multipath infrastructure to support BFD on satellite line cards. BFD over satellite is a multipath (MP) single-hop session and is supported on IPv4 address, IPv6 global address, and IPv6 link-local address. BFD over Satellite is supported on Cisco ASR 9000 4th Generation QSFP28 based dense 100GE line cards, Cisco ASR 9000 5th Generation High-Density Multi-Rate line cards. BFD over satellite is not supported in echo mode.

**Note**

- The nV Satellite access port bundles do not support BFD over bundles (BoB) over physical or bundle ICLs
- The **bfdmultipath include location** *node-id* command is required for all the line cards that host ICL links towards the Satellite.

BNG Interoperability

Important points about BNG interoperability with nV system:

- The only topology that BNG supports with nV Satellite is bundle access on the CPE side with non-bundle ICL links, as shown below:

CPE --- Bundle --- [Satellite] --- Non Bundle ICL --- ASR9K

- BNG does not support nV Satellite with bundle ICL and bundle access links, as shown below:

CPE --- Bundle --- [Satellite] --- Bundle ICL --- ASR9K

Layer-2 and L2VPN Features

All L2 and L2VPN features that are supported on physical ethernet or bundle ethernet interfaces are also supported on Satellite Ethernet interfaces, except L2TPv3 over IPv4 and L2TPv3 over IPv6. The maximum number of bundles supported by Cisco ASR 9000 Series Router is increased to 510.

For more details on L2VPN features supported on the Cisco ASR 9000 Series Satellite System, see *Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Configuration Guide*.

Layer-3 and L3VPN Features

All MPLS L3VPN features supported on ethernet interfaces are also supported on the Cisco ASR 9000 Series Satellite nV System.

For more information on these features, see *Cisco ASR 9000 Series Aggregation Services Router MPLS Layer 3 VPN Configuration Guide* and *Cisco ASR 9000 Series Aggregation Services Router Netflow Configuration Guide*.

Layer-2 and Layer-3 Multicast Features

The nV Satellite interfaces support edge and access functionalities of Layer-2 and Layer-3 multicast features, similar to the support on normal Ethernet and Bundle Ethernet interfaces.

For more information on these features, see *Cisco ASR 9000 Series Aggregation Services Routers Multicast Configuration Guide*.

Multicast IRB

Multicast IRB provides the ability to route multicast packets between a bridge group and a routed interface using a bridge-group virtual interface (BVI). It can be enabled with multicast-routing. THE BVI is a virtual

interface within the router that acts like a normal routed interface. For details about BVI, refer *Interface and Hardware Component Configuration Guide for Cisco ASR 9000 Series Routers*

BV interfaces are added to the existing VRF routes and integrated with the replication slot mask. After this integration, the traffic coming from a VRF BVI is forwarded to the VPN.

Quality of Service

Most Layer-2, Layer-3 QoS and ACL features are supported on Satellite Ethernet interfaces that are similar to normal physical Ethernet interfaces, with the exception of any ingress policy with a queuing action. However, for QoS, there may be some functional differences in the behavior because, in the Cisco IOS XR Software Release 4.2.x, user-configured MQC policies are applied on the Cisco ASR 9000 Series Router, and not on the satellite switch interfaces.

For more detailed information on QoS offload and QoS policy attributes, features, and configuration, see [Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide](#).



Note User-configured QoS policies are independent of any default port level QoS that are applied in order to handle IC link congestion and over-subscription scenarios. In addition to the default port-level QoS applied on the satellite system ports, default QoS is also applied on the Cisco ASR 9000 Series Router Host side, to the ingress and egress traffic from and to the Satellite Ethernet ports.

Time of Day Synchronization

The Time of Day parameter on the satellite switch is synchronized with the time of day on the host. This ensures that time stamps on debug messages and other satellite event logs are consistent with the host, and with all satellite switches across the network. This is achieved through the SDAC Discovery Protocol from the host to the satellite switch when the ICLs are discovered.

Satellite Chassis Management

The chassis level management of the satellite is done through the host because the satellite switch is a logical portion of the overall virtual switch. This ensures that service providers get to manage a single logical device with respect to all aspects including service-level, as well as box-level management. This simplifies the network operations. These operations include inventory management, environmental sensor monitoring, and fault/alarm monitoring for the satellite chassis through the corresponding CLI, SNMP, and XML interfaces of the host Cisco ASR 9000 Series Router.



Note The satellite system hardware features, support for SFPs, and compatible topologies are described in the *Cisco ASR 9000 Series Aggregation Services Router Hardware Installation Guide*.



Note All the SNMP features supported on the Cisco ASR 9000 Series router is supported for satellite, provided the **SystemOwner** permissions are assigned to the snmp-server community. For more information, see *Cisco ASR 9000 Series Aggregation Services Router MIB Specification Guide*.

ARP Redundancy Support for Dual Head Mode

The Address Resolution Protocol(ARP) redundancy support feature allows you to synchronize the ARP table entries between the active and standby hosts in a dual head mode. In the Dual Head mode, all satellite traffic flows to one host router (active host) and therefore all ARP entries reside on the active host. When the active host goes down, the standby host needs to rebuild the ARP table. But with a large number of devices, this can take a significant amount of time, impacting the network uptime. The process of synchronizing the ARP database between the active and standby hosts can reduce this downtime. The synchronization happens as and when a new entry is learnt or when the active host goes down. There is no specific frequency for synchronization. You can also pause or restart the synchronization process whenever required.

The ARP database is distributed across all the nodes (RP/LC) in the router. The ARP entries are programmed for an interface only in the node where the data plane is replicated. In case of virtual interfaces (like Bundle-Ether, BVI and so on), where the control plane is located in RP, the ARP entries are programmed only in the line cards wherever the data plane is replicated. The RP would have entries only for the physical interfaces that have its control plane in RP like Management interfaces. The Txlist infrastructure library facilitates the ARP table entries synchronization process to happen. The ARP table has three key entries, which are the IPv4 address, MAC address, and the corresponding interface.

The current architecture can support only two host nodes in a redundancy group. But a specific host router can be part of multiple redundancy groups. This is a sample configuration for ARP synchronization.

```
group <id-number>
  peer <neighbor ipv4 address>
  source-interface <interface-name>
  interface-list
    interface <interface-name to be synced> id <unique-id>
  !
!
```

You can use the **show arp redundancy summary location** command to display the summary of entries.

Ethernet Link OAM

The Satellite nV system Ethernet interfaces support Ethernet Link OAM feature when ICL is a physical interface. Cisco IOS XR Software also supports Ethernet Link OAM feature over Satellite Ethernet interfaces when the ICL is a bundle interface.

802.3ah Loopback Support on Satellite nV System

The 802.3ah loopback feature allows you to loopback all the non OAMPDU packets from the remote subordinate node. The primary port is the one that controls the loopback status on the remote node. You can conduct various performance and quality testing on the remote subordinate port using this feature. This loopback mode is disruptive because the traffic does not flow through the router in subordinate state. The Loopback function is supported on the satellite nV system ports both in over protected and unprotected ICL modes. You can define the loopback function on the primary node using the **remote-loopback** command in the ethernet oam interface configuration mode.

You can use the **ethernet oam loopback enable <interface>** and **ethernet oam loopback disable <interface>** commands from the primary node to enable and disable the loopback respectively on the subordinate node.



Note You cannot loopback the Layer 2 protocol packets such as CFM, OAM, and LACP.

This configuration snippet shows how to configure loopback.

```
RP/0/RSP0/CPU0:router(config)# interface gigabitEthernet 0/1/0/9
RP/0/RSP0/CPU0:router(config-if)# ethernet oam
RP/0/RSP0/CPU0:router(config-if-eoam)# remote-loopback
RP/0/RSP0/CPU0:router(config-if-eoam)# commit
```

This is a sample output of the show ethernet oam discovery command, which indicates that loopback is supported.

```
RP/0/RSP0/CPU0:router# show ethernet oam discovery

Local client
-----
Administrative configuration:
  PDU revision:                9
  Mode:                        Active
  Unidirectional support:      Y
  Link monitor support:        Y
  Remote loopback support:      Y <<Indicates loopback is supported
  MIB retrieval support:       Y
  Maximum PDU size:            1500
Operational status:
  Port status:                 Operational
  Loopback status:             Remote <<Indicates primary
  Interface mis-wired:         N
```

For more information on Ethernet OAM configuration and commands, see the chapters *Configuring Ethernet OAM* and *Ethernet OAM Commands on the Cisco ASR 9000 Series Router* modules in the *Interface and Hardware Component Configuration Guide for Cisco ASR 9000 Series Routers* and *Interface and Hardware Component Command Reference for Cisco ASR 9000 Series Routers* respectively.

Native Image and SMU Push Support for Cisco NCS 5000 Series Satellites

The Cisco NCS 5000 Series Satellite runs virtualized IOS XR software and mainly consists of three Virtual Machines (VM), the XR VM, the System Admin VM and the Host VM.

Cisco ASR 9000 Series Router host supports the capability to push native Cisco NCS 5000 Series Satellite XR VM images and SMU fixes from the Cisco ASR 9000 Series Router host similar to satellite nV pie image upgrades from Cisco IOS XR Software Release 6.1.2 onwards.

This capability is expanded to native Cisco NCS 5000 Series Satellite System Admin VM and Host VM images and SMU fixes from Cisco IOS XR Software Release 6.3.1 onwards.

This provides the flexibility of managing the software images and pies on the satellite and also offers the zero touch provisioning model of nV satellites. The commands can be run on the Cisco ASR 9000 Series Router host for addition and activation of both upgrade images and mandatory SMUs and this allows a fast turnaround time for any critical fixes on Cisco NCS 5000 Series Satellite.

Support of 10x10G ICLs on Cisco NCS 5000 Series Satellites

A maximum of 10x10G ICL ports can be connected from the Cisco NCS 5000 Series Satellite to each of the Cisco ASR 9000 Series Router host in a Layer 1 Hub and Spoke topology (both single and dual head modes) using this feature. The ICLs could be either a 10G ICL or one of the 4x100G ICLs, amounting to a maximum of 10 physical links with each of the hosts.

The 10G ICLs have to be configured using the configurable fabric port feature. The 10x10G ICL links can be configured as members of a Bundle ICL or Physical ICL. If they are used as members of a Bundle ICL, then members must be distributed across NPs or line cards to achieve redundancy and overcome access port scale constraints. Furthermore, the bundle ICL can take up to 10x10G members and still have 100G links as ICLs in a mixed bundle and physical ICL connectivity for 10G to 100G fabric transitions for each of the hosts.

- Qos offload
- Fabric CFM for both physical and bundle ICLs and dual head modes.

Similarly, fabric CFM on bundle ICLs in Layer 1 Hub and Spoke topology is supported for Cisco ASR 9000v Satellites.

Support for 40G ICLs on Cisco NCS 5000 Series Satellites

40G ICL ports can be connected from the Cisco NCS 5000 Series Satellite to the Cisco ASR 9000 Series Router in a single host hub and spoke topology. This feature is supported only with the A9K-8X100G-SE and A9K-8X100G-TR variants of the Cisco ASR 9000 High Density 100GE Ethernet line cards and specifically with the optics CPAK-100G-SR10 breaking out to 2xQSFP-40G-SR4 ports. The breakout configuration will have to be applied to the CPAK in order to provision the 40G ports on the Cisco ASR 9000 Series Router host. For details on the breakout configuration, see *Managing the Router Hardware* chapter in *Cisco ASR 9000 Series Aggregation Services Router System Management Configuration Guide*.

ICL bundles can also be configured with 40G ICL links.

Support for MACsec Sessions on nV Satellite Access Interface

MACsec encryption is an IEEE 802.1AE standards based layer 2 hop-by-hop encryption that provides data confidentiality and integrity for media access independent protocols. MACsec Key Agreement Protocol (MKA) developed as part of 802.1X allows for peer discovery, confirmation of mutual authentication, and agreement of secret keys that may be used by MACsec to protect any data exchanged between peers. For details, see *Implementing MACsec Encryption* chapter in the *System Security Configuration Guide for Cisco ASR 9000 Series Routers*.

MACsec sessions are supported on 1G access interfaces over 10G ICL on the Cisco ASR 9000v Satellite and on both 1G and 10G access interfaces over either 10G candidate ICL or 100G ICL on the Cisco NCS 5000 Series Satellite.

Restrictions of the Satellite nV System

Software restrictions of the satellite system are:

- The nV Satellite interfaces are expected to be deployed as access or edge interfaces and hence do not support functionalities of core interfaces on multicast topologies.

- The inter-chassis link redundancy is supported only through the static EtherChannel, and not through LACP based link bundles. Minimum and maximum link commands are not applicable when ICL is a bundle.
- Multi-chassis Link Aggregation is supported if there are two independent Cisco ASR 9000 Series Routers acting as the POA (Point of Attachment), each with its own satellite switch, and the DHD (Dual Homed Device) connecting through each of the satellite switches. However, MC-LAG is not supported with a single satellite switch that connects two separate Cisco ASR 9000 Series Routers through an ICL LAG.
- Pseudowire Headend is not supported on the Satellite interfaces.
- Access bundles having both satellite and non-satellite interfaces is not supported.
- On the Cisco ASR 9000 Series 24-Port and 48-Port Dual-Rate 10GE/1GE line cards (A9K-48X10GE-1G-SE, A9K-48X10GE-1G-TR, A9K-24X10GE-1G-SE, A9K-24X10GE-1G-TR), only one interface from each PHY can be configured for SyncE output. So, when breakout interfaces are configured as ICLs, only one of them can support SyncE output for SyncE offload to the satellite.
- The nV Satellite access port bundles do not support BFD over bundles (BoB) over physical or bundle ICLs
- BFD hardware offload is not supported on satellite line cards because BFD hardware offload is supported for single-hop and single-path asynchronous packets only. For more information on BFD hardware offload, see the chapter *Implementing BFD* in the *Routing Configuration Guide for Cisco ASR 9000 Series Routers*.
- MACsec encryption over nV Satellite access sub-interface with dual VLAN tags is not supported
- When 100G ICL port is used with a breakout of 10x10G ports, the 10G session limit would be applied for each of the breakout ports. When 100G port is used as an ICL as is, then MACsec session limit of 100G is applied.
- MACsec is not supported on nV satellite access interfaces if the satellites are in Dual host mode.
- QoS offload feature is not supported when MACsec is configured on the nV satellite access interfaces.
- In the case of Cisco ASR 9000v satellites with bundle ICL, you should first shut down the ICL bundle before making any configuration changes on the ICL bundle or before changing the bundle members or the bundle member interface types
- Bundle ICL is not supported with MACsec feature.
- MACsec feature is not supported on nV satellite access interfaces that are mapped to the Cisco ASR 9000 Series host through 1G candidate ICL or the 40G ICL ports.
- MACsec feature is not ISSU aware.
- Whenever the first MACsec policy is applied on any of the nV satellite access interfaces or when the last MACsec policy is removed from any of the nV satellite access interfaces, there will be a ICL interface link flap and the satellite will be reconnected post this flap.
- MACsec policy can be applied either only on the main nV satellite access interface or only on the nV satellite access sub-interface that are mapped to a single Physical ICL.
- MACsec feature is not supported on nV satellite access interfaces if they are configured as core-facing interfaces of L2VPN or L3VPN services.

- MACsec policies in the Must-Secure and Should-Secure modes cannot be applied on same nV satellite access interface. But they can be applied on different access interfaces on same ICL.
- Mixing of interface types (1G, 10G, 100G, 40G) is not supported for bundle ICLs.
- The maximum number of satellites per system is 64.
- Link failure detection on nV satellite access ports can take seconds to propagate to the host, especially with traffic congestion on the nV satellite fabric links. Hence, OAM or BFD should be run over nV satellite access ports for faster detection.
- The nV satellite access interfaces do not support IP Fast Reroute (FRR) and Topology-Independent Loop Free Alternate (TI-LFA) protection.
- Dual home mode does not support aggressive-mode UDLD as the standby host causes the active host to time out and the port is error-disabled when the active host recovers. Hence, you should configure normal UDLD timer for the dual home mode.
- With aggressive UDLD timers in hub and spoke topology, UDLD takes the far end down, if the host is ungracefully shutdown or with host configuration operations such as **commit replace**.
- A satellite can be connected to only one Host in the Hub and Spoke topology model and can be connected to only two hosts in a Dual-homed network architecture.
- During configuration changes that removes an ICL from a satellite, there is no guarantee that a reject packet will be transmitted. Hence, it is recommended that you shut down the ICL port before you change or remove a configuration on an ICL interface or wait for an idle time out (which is 30 seconds) to bring down sdac discovery.
- The nV system supports hub and spoke topology with both bundle and physical (non-bundle) links as access-ports and ICL, including access bundle over ICL bundle cases.
- The Cisco ASR 9000 4th Generation QSFP28 based dense 100GE line cards support a maximum of 400 nV satellite access ports per NPU and 1600 nV satellite access ports per line card on both -TR and -SE line card variants. When using ICLs in breakout mode or for lower speed ICLs (<100G), a single QoS resource group within an NPU, called a “chunk” can support a maximum of up to 200 nV satellite access ports. The actual possible scale on a given chunk depends on the port density, and other users of these chunk resources. This can be monitored using **show qoschal resource summary [npid | locationnode-id]** command.

If a chunk is full, you can still distribute the satellites across ICLs on other chunks. There are 4 chunks per NPU, and the port to chunk mapping can be checked using **show qoschal ep [npid | locationnode-id]** command to achieve the full NPU scale.



Note After RSP Failover, it is expected to see the satellite in Connecting state for about a min and OIR messages for satellite and sat-ether ports like below:

```
RP/0/RSP0/CPU0:Oct 24 05:19:43.278 : invmgr[252]: %PLATFORM-INV-6-OIRIN : OIR: Node 100/
inserted
RP/0/RSP0/CPU0:Oct 24 05:19:43.484 : invmgr[252]: %PLATFORM-INV-6-IF_OIRIN : xFP OIR:
SAT100/0/0 port_num: 0 is inserted, state: 1
```

However, the data plane forwarding is expected to be up throughout.



Note E-LMI is not supported over access LAG. OAM / EDPL are not supported over bundle ICL when access is also a bundle.

Refer to the *Cisco ASR 9000 Series Aggregation Services Router Release Notes* for additional software restrictions.

Restrictions of the Cisco NCS 5000 Series Satellite

These are the restrictions and limitations of the Cisco NCS 5000 Series Satellite:

Hardware Limitations

- Satellite hosting line cards must be second or third generation Ethernet cards on Cisco ASR 9000 Series Router. The line card can be –TR or –SE. If –TR, each satellite access gets 8 TM queues and overall nV and other scale limitations persist.
- The native Cisco NCS 5000 Series satellite SMU push feature, supported from Cisco IOS XR Software Release 6.1.1, stores Cisco NCS 5000 Series satellite images and pies on the disk.
- 100G ICLs are only supported on the 8x100G or 4x100G Cisco ASR 9000 High Density 100GE Ethernet line cards and the 1x100GE and 2x100GE MPAs..
- On the Cisco ASR 9000 Series Router host, 40GE ICL is supported only on the A9K-8X100GE-SE and A9K-8X100GE-TR variants of the Cisco ASR 9000 High Density 100GE Ethernet line cards with the optic pluggable CPAK-100G-SR10 breaking out into 2xQSFP-40G-SR4.
- Cisco NCS 5000 Series satellites require manual disk cleanup of inactive packages to continue having space for subsequent software upgrades. You can get the inactive package list from the output of **show nv satellite install package** on the host and then, remove the inactive/old packages through **install nv satellite <satellite id> remove package <package name>**.

General Limitations

- Cisco NCS 5000 Series satellite is supported as a satellite on Layer 1 Hub and Spoke topology in both single and dual host connectivity. The other advanced topologies are not supported.
- Unlike Cisco ASR 9000v satellite, Cisco NCS 5000 Series satellite does not reboot on being disconnected from the Host for more than 30 minutes as there are internal mechanisms to auto correct/recover the driver issues without reloads on this device.
- Auto QoS on Cisco NCS 5000 Series satellite follows a trusted CoS model and only differentiates incoming frames based on their CoS/DSCP values. No specific bandwidth reservation is done for any other protocol packets implicitly.
- Auto QoS on Cisco NCS 5000 Series satellite cannot classify incoming frames on the access ports based on MPLS exp. So, any tunneled high priority packet based on just MPLS exp may not get the right scheduling weights.
- Fan RPM thresholds are adaptive and driven using algorithms for optimum speed thresholds. Only fan tray online insertion and removal and fan failure alarms are reported.
- 10GE ports can operate at restricted speeds of 10/100/1000 M on qualified optics. While the software does not support changing the speed and auto negotiation parameters for such ports using the configuration

commands, the port will always auto negotiate to the maximum peer capability. Hence, the restricted speed operations are only supported with peers which can also auto negotiate. The interfaces continue to be named as 10GE interfaces with 10/100/1000 M only reflecting in the operating speed, bandwidth, and route metrics

- Due to the remote port configuration being mapped to a satellite id rather than a specific satellite type internally, 1G remote port configuration from legacy satellites may also be allowed for Cisco NCS 5000 Series satellite id. However, this is a misconfiguration and the satellite ports must always be configured as 10G ports even when 1G optics are used. This is in line with the previous restriction on 10G ports being used for 1G optics but still being 10G ports operating at restricted speeds. 1G remote port configuration is rejected by the satellite otherwise.
- Existing caveats for the features – on loaded on the Host and offloaded, if any to the satellite remain. The addition of this new satellite variant does not implicitly add enhanced capabilities. None of the offload features on Cisco ASR9000v are supported on Cisco NCS 5000 Series satellites unless explicitly stated.
- Only a maximum of 40 Cisco NCS 5000 Series satellites for each Cisco ASR 9000 Series Host are allowed.
- Scale restrictions for QoS offload are same as regular QoS on the Cisco NCS 5000 Series routers.
- When upgrading the Cisco NCS 5000 Series satellite from an older image to Cisco IOS XR Software Release 6.1.x or later, the auto FPD upgrade feature has a caveat which requires a one time manual upgrade of FPD images during this upgrade on each of the Cisco NCS 500x satellite. This upgrade procedure is similar to that of a standalone Cisco NCS 500x device where an upgrade **hw-module fpd** <> **CLI** is used in the admin mode of the satellite console. For more details on the standalone procedure, refer the *Upgrading FPD* chapter in the *System Management Configuration Guide for Cisco NCS 5000 Series Routers*.
- QoS offload on Cisco NCS 5000 Series satellites cannot support multiple marking action. For more details on remaining action support, please see the QoS offload supported capability matrix in the Supported Platform-Specific Information for QoS Offload section of *Modular QoS Configuration Guide for Cisco ASR 9000 Series Routers*.
- While Cisco NCS 5000 Series satellites can support up to 10 members in a bundle ICL, the Cisco ASR 9000 Series Host has hardware restrictions on the number of access ports that can be supported for each NP. As the access ports are instantiated for each member of a bundle ICL, the maximum number of access ports supported in Layer 1 Hub and Spoke topology(200), can be easily exceeded if the bundle ICL members are on the same NP. For example, in a case where all 80 ports of a Cisco NCS 5002 satellite are mapped to a bundle ICL, only two members would consume up to 160 access port resources. Hence, the expectation would be to spread out such members across NPs and line cards to achieve the maximum supported scale as well as redundancy benefits.
- For the native SMU push feature, combo SMUs, tar balls, and non-XR SMUs are not supported. Only Cisco IOS XR SMUs(process restart or reload) in rpm format is supported by this infrastructure.
- When 40G ports are used as 40G ICLs on the Cisco NCS 5000 Series Satellite, the 40G ports are nothing but the 100G ports operating at restricted speed with no auto negotiation. The interfaces will continue to be named as HundredGigE interfaces and just the operating speed, bandwidth and route metrics will be reflected as 40G. The corresponding fabric ports created on the Cisco ASR 9000 Series Router host will also be displayed as nVFabric-HundredGigE ports even though they operate at 40G speed.
- Only Hub and Spoke single host topology is supported with 40G ICL ports.

- We do not recommend you to perform an RSP switchover in between a satellite install operation through reference lists or the file option (For more information on these install operation variants, see the topic [Install Reference Command, on page 57](#)). This is because the install operation remains incomplete even after the RSP switchover is over. To recover from this condition, you should terminate the incomplete install operation manually after the RSP switchover and then start a new install operation.

Satellite nV Usability Enhancements

Satellite nV usability enhancements introduce the following functionalities:

nV Satellite Auto Image Upgrade

nV satellite auto image upgrade feature introduces support for automated upgrade of the satellite image in these two scenarios:

- Following the Cisco IOS XR software release upgrade
- Satellite is connected to a host and the image on the satellite is not matching the one packaged in Cisco IOS XR software

To configure the nV satellite auto image upgrade feature, use the **upgrade on-connect** command in nV satellite configuration.



Note Starting with Cisco IOS XR Software Release 6.6.1, the command for configuring nV satellite auto image upgrade feature is **upgrade [on-connect | on-first-connect] <reference ref_name >**.



Note For dual-head satellite systems, the auto upgrade configuration must only be applied on one host, that is, the host that the satellite will install from. There is currently no checking for this, and if the user applies the configuration on both hosts and the hosts have different satellite versions installed, then the satellite may go into a loop of installing, until either:

- the configuration is removed from one host OR
- the satellite version activated on both hosts is the same



Note The auto upgrade feature downgrades the satellite image automatically if the host software is downgraded to a release that still supports this feature. This is unintentional as the downgrade of some of the satellite firmware like ROMMON is not supported. It is recommended to turn off the auto upgrade feature before such downgrade operations.

Auto Upgrade Behavior

Auto upgrade *on-connect* variant means that a satellite is auto upgraded to the current version installed on the host the next time the TCP connection to the satellite is established. On establishing the TCP connection, the configuration is checked, and an upgrade is started if the version is not the latest and the user has configured auto upgrade on-connect.



Note Starting with Cisco IOS XR Software Release 6.3.1, auto upgrade *on-first-connect* variant is similar but the configuration check and upgrade happens only during the first time when the TCP connection is established. If the satellite is in a discovered state, TCP connection may reset but auto upgrade will not take place.

on-connect or *on-first-connect* does not trigger an auto upgrade at the time of applying the configuration (and so does not impact a satellite that is already connected). However, given that a satellite upgrade impacts traffic, a notification *syslog* is generated any time the configuration is applied to a connected satellite that currently does not have the latest software version installed.



Note The auto upgrade feature stops a satellite from progressing further into the control state machine, until the operation is complete.



Note Due to the above mentioned behavior of auto upgrade and because the Cisco ASR 9000v satellite does not support ROMMON downgrade, a downgrade to Cisco IOS XR Software Release 5.3.x based releases on the host will end up with the satellites being stopped until the downgrade completes, which would never happen because the ROMMON cannot be downgraded. This can result in the Cisco ASR 9000v satellites stuck in auto upgrade and not coming up. The recommendation is to turn off auto upgrade for any downgrade to Cisco IOS XR Software Release 5.3.x on the host from any other future release that has a higher ROMMON version for Cisco ASR 9000v.



Note For the Cisco NCS 5000 Series Satellite, auto upgrade to a reference list of software packages is possible with an optional reference keyword and reference name along with upgrade [*on-connect* | *on-first-connect*] configuration. For more details on defining a named reference list, refer to the section [Install Reference Command, on page 57](#). On Cisco ASR 9000 Series Router host running IOS XR 64 Bit operating system, auto upgrade for the Cisco NCS 5000 Series satellite is only possible with the reference list due to restrictions in supporting the satellite nV pie.

Simplified Access to Satellites

The provision of new variants to the telnet command simplifies the process of accessing satellites. These new variants of the telnet command make it easier to the users with specific task groups to be able to telnet to the satellite.

There are two methods of using the command:

- If the satellite is connected, the user can specifically use the satellite ID to connect.

Example: telnet satellite 100 ----> telnet to satellite with ID 100.

An error message is generated if this method does not work.

- If the satellite is not connected, the user can use the VRF and IP address of the satellite to connect.

Example: telnet satellite vrf default 1.2.3.4 ----> telnet to a satellite in the default VRF with IP address 1.2.3.4



Note The VRF must always be specified by the user. If the user has configured any IP address configuration, then the VRF is either *default* or the configured VRF name. Else, the VRF is the hidden satellite VRF name (**nVSatellite). The VRF is visible in the status command when applicable.

MPP Check Skip for Satellite Image Download with Auto IP

The Management Plane Protection (MPP) feature in Cisco IOS XR software provides the capability to restrict the interfaces on which network management packets are allowed to enter a device.

User MPP configuration requiring to allow Trivial File Transfer Protocol (TFTP) if any other MPP configuration is present on the router, is complicated, especially when users do not know the protocol being used to download image for nV Satellites. The upgrade of the satellite image fails unless an entry is added to MPP configuration for the satellite VRF (if any) to allow TFTP through the Inter-Chassis Link (ICL) interface .

The enhancement to skip MPP checks for nV satellites configured with auto IP, automatically allows TFTP over the satellite ICL interface (for satellite image download) by skipping MPP checks for authenticated satellites. TFTP is only allowed under the hidden satellite VRF name (**nVSatellite). So, any satellites configured with manual IP or VRF will still need MPP configuration to be modified to allow TFTP for the configured VRF on the ICL .

Implementing a Satellite nV System

The Interface Control Plane Extender (ICPE) infrastructure has a mechanism to provide the Control Plane of an interface physically located on the Satellite device in the local Cisco IOS XR software. After this infrastructure is established, the interfaces behave like other physical ethernet interfaces on the router.

The ICPE configuration covers these functional areas, which are each required to set up full connectivity with a Satellite device:

Defining the Satellite nV System

Each satellite that is to be attached to Cisco IOS XR Software must be configured on the host, and also be provided with a unique identifier. In order to provide suitable verification of configuration and functionality, the satellite type, and its capabilities must also be specified.

Further, in order to provide connectivity with the satellite, an IP address must be configured, which will be pushed down to the satellite through the Discovery protocol, and allows Control protocol connectivity.

This task explains how to define the satellite system by assigning an ID and basic identification information.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	nv Example: RP/0/RSP0/CPU0:router(config)# nv	Enters the nV configuration submode.
Step 3	satellite <i>Sat ID</i> Example: RP/0/RSP0/CPU0:router(config-nV)# satellite <100-65534>	Declares a new satellite that is to be attached to the host and enters the satellite configuration submode. . Note Adding a new satellite to a host does not impact any other working satellites or line cards in the host router.
Step 4	serial-number <i>string</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# serial-number CAT1521B1BB	Serial number is used for satellite authentication.
Step 5	description <i>string</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# description Milpitas Building12	(Optional) Specifies any description string that is associated with a satellite such as location and so on.
Step 6	type <i>type_name</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# satellite 200 type ? asr9000v Satellite type	Defines the expected type of the attached satellite. The satellite types are asr9000v, asr9000v2, Cisco NCS 5001, and Cisco NCS 5002.. For other supported satellite types, see Prerequisites for Configuration .
Step 7	ipv4 address <i>address</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# ipv4 address 10.22.1.2	Specifies the IP address to assign to the satellite. ICPE sets up a connected route to the specified IP address through all configured ICLs.
Step 8	secret <i>password</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# secret <password>	Specifies the secret password to access the satellite. In order to login you must use root as the user name and password as the secret password.
Step 9	candidate-fabric-ports <i>interface-type slot/subslot/port-range</i> Example: RP/0/RSP0/CPU0:router(config-satellite)# candidate-fabric-ports nVFabric-GigE 0/0/21-25	Specifies the additional ports on the satellite which can be used as dynamic ICLs/ fabric ports along with the default set of ports. The interface type can be nVFabric-GigE, nVFabric-TenGigE or nVFabric-HundredGigE. The value of slot and subslot can be in the range of 0-5. The port ranges are specified as a comma-separated list of

	Command or Action	Purpose
		sub-ranges. Each subrange consists of either a single number, or a low and high (both inclusive) hyphen-separated range. The sub-ranges must be in a non-overlapping, and strictly in an ascending order.
Step 10	end or commit Example: RP/0/RSP0/CPU0:router(config)# end or commit	Saves configuration changes. <ul style="list-style-type: none"> • When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Auto-IP

The Auto IP feature improves the plug-and-play set up of an nV satellite system. With the Auto IP feature, IP connectivity to the satellite is automatically provisioned. As a result:

- The nV Satellite Loopback interface is created on the host
- Loopback interface is given an IP address from a private satellite VRF
- Satellite fabric links are unnumbered to the loopback interface
- The IP address assigned to satellite is auto-generated from the satellite VRF

In the case of Auto IP, you need not specify any IP addresses (including the IP address on the Satellite itself, under the satellite submode), and the nV Satellite infrastructure assigns suitable IP addresses, which are taken from the 10.0.0.0/8 range within a private VRF to both the satellites and the satellite fabric links. All such Auto IP allocated satellites are in the same VRF, and you must manually configure IP addresses, if separate VRFs are required.



Note You cannot combine auto-configured Satellites with manually configured Satellites within the same satellite fabric.

The auto-IP feature assigns an IP address in the format 10.x.y.1 automatically, where:

- x is the top (most significant) 8 bits of the satellite ID
- y is the bottom 8 bits (the rest) of the satellite ID



Note The Auto-IP configuration is the recommended mode of configuration over the manual Host IP address configuration. You can also override the Auto IP feature by using the standard IP configuration.

For examples on configuration using the Auto-IP feature, see [Configuration Examples for Satellite nV System](#).

Configuring the Host IP Address

This procedure gives you the steps to configure a host IP address on a loopback interface.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/0RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface loopback0 Example: RP/0/0RSP0/CPU0:router(config)# interface loopback0	Specifies the loopback address for the interface.
Step 3	ipv4 address Example: RP/0/0RSP0/CPU0:router(config-int)# ipv4 address 8.8.8.8 255.255.255.255	Configures the host IP address on a loopback interface.
Step 4	end or commit Example: RP/0/0RSP0/CPU0:router(config)# end or RP/0/0RSP0/CPU0:router(config)# commit	Saves configuration changes. <ul style="list-style-type: none"> • When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

	Command or Action	Purpose
		<ul style="list-style-type: none"> - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring Inter-Chassis Links and IP Connectivity

Inter-Chassis Links (ICLs) need to be explicitly configured, in order to indicate which satellite is expected to be connected. You must also specify the access port, that is down-stream 10GigE ports, which cross-link up to the Host through the configured ICL. In order to establish connectivity between the host and satellite, suitable IP addresses must be configured on both sides. The satellite IP address is forwarded through the Discovery protocol. The configuration is described in the section, [Defining the Satellite nV System, on page 31](#).



Note This configuration shows the use of the global default VRF. The recommended option is to use a private VRF for nV IP addresses as shown in the [Satellite Management Using Private VRF, on page 70](#) subsection under *Configuration Examples for Satellite nV System*.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface <i>interface-name</i> Example: RP/0/RSP0/CPU0:router(config)# interface TenGigE0/2/1/0	The supported inter-chassis link interface types are limited by the connectivity provided on the supported satellites. GigabitEthernet, TenGigE, FortyGigE, HundredGigE and Bundle-Ether interfaces are the supported ICL types.
Step 3	description Example: RP/0/RSP0/CPU0:router(config-interface)# description To Sat5 1/46	Specifies the description of the supported inter-chassis link interface type.
Step 4	ipv4 point-to-point Example:	(Optional) Configures the IPv4 point to point address.

	Command or Action	Purpose
	RP/0/RSP0/CPU0:router(config-interface)# ipv4 point-to-point	
Step 5	ipv4 unnumbered loopback0 Example: RP/0/RSP0/CPU0:router(config-interface)# interface unnumbered loopback0	(Optional) Configures the IPv4 loopback address on the interface.
Step 6	nV Example: RP/0/RSP0/CPU0:router(config-if)# nv	Enters the Network Virtualization configuration mode.
Step 7	satellite-fabric-link satellite <id> Example: RP/0/ORSPO/CPU0:router(config-int-nv)# satellite-fabric-link satellite 200	Specifies that the interface is an ICPE inter-chassis link.
Step 8	remote-ports interface-type Example: RP/0/RSP0/CPU0:router(config-int-nv)# remote-ports GigabitEthernet 0/0/0-30	Configures the remote satellite ports 0 to 30.
Step 9	end or commit Example: RP/0/RSP0/CPU0:router(config)# end or RP/0/ORSPO/CPU0:router(config)# commit	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session. <p>For information on QoS configuration on ICLs, see <i>Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide</i>.</p>

Configuring the Inter-Chassis Links in a Dual Home Mode

These are the steps for configuring Inter-chassis links in the case of a dual home mode.

Before you begin

MPLS LDP needs to be up and running between the two hosts for the dual home configuration.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface <i>interface-name</i> Example: RP/0/RSP0/CPU0:router(config)# interface TenGigE0/2/1/0	The supported inter-chassis link interface types are limited by the connectivity provided on the supported satellites. GigabitEthernet, TenGigE, and Bundle-Ether interfaces are the only support ICL types.
Step 3	satellite-fabric-link satellite <id> Example: RP/0/RSP0/CPU0:router(config-interface)# satellite-fabric-link satellite 100	Configures the ICPE inter-chassis link for the specified satellite.
Step 4	ipv4 point-to-point Example: RP/0/RSP0/CPU0:router(config-interface)# ipv4 point-to-point	(Optional) Configures the IPv4 point to point address.
Step 5	ipv4 unnumbered loopback0 Example: RP/0/RSP0/CPU0:router(config-interface)# interface unnumbered loopback0	(Optional) Configures the IPv4 loopback address on the interface.
Step 6	redundancy iccp-group Example: RP/0/RSP0/CPU0:router(config-interface)# redundancy iccp-group 1	Configures the ICCP redundancy group. In order to configure multiple ICCP redundancy groups, repeat steps 6 through 10 with a different redundancy group number.
Step 7	minimum preferred links <i>num</i> Example: RP/0/RSP0/CPU0:router(config-satellite-fabric-link)# redundancy minimum preferred links 2 (or) RP/0/RSP0/CPU0:router(config-satellite-fabric-link)#	(Optional) Configures the minimum number of preferred satellite fabric bundle-ether links. To configure this parameter, the interface must be a bundle-ether interface. If you do not enable this parameter on any host, then it is assumed as 0 by default on that host. Hence, the decision of failover is judged using that value. The range is 1 to 64.

	Command or Action	Purpose
	<pre> redundancy RP/0/RSP0/CPU0:router(config-nv-red)# minimum preferred links 2 </pre>	<p>Note You can either go to the redundancy mode to configure minimum preferred links or you can specify redundancy keyword before configuring the minimum preferred links.</p>
Step 8	<p>member neighbor 9.9.9.9</p> <p>Example:</p> <pre> RP/0/RSP0/CPU0:router(config-interface)# member neighbor 9.9.9.9 </pre>	Configures the LDP neighbor.
Step 9	<p>backbone interface <i>interface_type</i></p> <p>Example:</p> <pre> RP/0/RSP0/CPU0:router(config-interface)# backbone interface TenGigE0/1/0/3 </pre>	(Optional) Configures the backbone interface for the PE isolation.
Step 10	<p>remote-ports <i>interface-type</i></p> <p>Example:</p> <pre> RP/0/RSP0/CPU0:router(config-int-nv)# remote-ports GigabitEthernet 0/0/0-30 </pre>	Configures the remote satellite ports 0 to 30.
Step 11	<p>end or commit</p> <p>Example:</p> <pre> RP/0/RSP0/CPU0:router(config)# end or RP/0/RSP0/CPU0:router(config)# commit </pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring the Satellite nV Access Interfaces

The access 1Gigabit Ethernet/10GigE interfaces on the satellite are represented locally in Cisco IOS XR Software using interfaces named Gigabit Ethernet similar to other non-satellite 1Gigabit Ethernet/10GigE

interfaces. The only difference is that the rack ID used for a satellite access 1Gigabit Ethernet/10GigE interface is the configured satellite ID for that satellite.

These interfaces support all features that are normally configurable on 1Gigabit Ethernet/10GigE interfaces (when running over a physical ICL), or Bundle-Ether interfaces (when running over a virtual ICL).



Note With respect to the dual home mode, the satellite access port configuration needs to be done on both the active and standby hosts. The administrator needs to make sure that the same configuration is applied for a particular access port on both the active and standby hosts. In addition, any feature configurations on satellite-access ports needs to be configured identically on both the hosts. Also, the configuration synchronization between the hosts is not currently supported. See [Satellite Interface Configuration](#).

Configuring the Fabric CFM for L1 Physical ICL

This procedure gives you the steps to configure CFM on a Satellite nV Fabric link interface. CFM can be configured with L1 Physical ICL architecture. It can also be configured on Dual Home modes.



Note While configuring Ethernet CFM, if no level or interval values are selected, then level 0 and interval of 1 second are selected by default.



Note Fabric CFM is supported on both physical and bundle ICLs with minimum CCM interval being 1 second for Cisco ASR 9000v and 100ms for Cisco NCS 500x satellites. When you configure CCM intervals lower than the minimum interval supported by the satellite, then it may be auto corrected while programming the satellite. But when you configure intervals below the minimum interval supported by the host, then it would be rejected during configuration itself. It is strongly recommended to configure only up to the minimum interval as specified.

SUMMARY STEPS

1. **configure**
2. **interface** *interface-name*
3. **nv**
4. **satellite-fabric-link satellite** *<id>*
5. **ethernet cfm**
6. **level** *value*
7. **continuity-check interval** *time*
8. **end** or **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface interface-name Example: RP/0/RSP0/CPU0:router(config)# interface Gigabit 0/1/0/0	Specifies the supported inter-chassis link interface types limited by the connectivity provided on the supported satellites.
Step 3	nv Example: RP/0/RSP0/CPU0:router(config-if)# nV	Enters the Network Virtualization configuration mode.
Step 4	satellite-fabric-link satellite <id> Example: RP/0/RSP0/CPU0:router(config-int-nv)# satellite-fabric-link satellite 200	Specifies that the interface is an ICPE inter-chassis link.
Step 5	ethernet cfm Example: RP/0/RSP0/CPU0:router(config-int-nv)# ethernet cfm	Enters Ethernet Connectivity Fault Management (CFM) configuration mode.
Step 6	level value Example: RP/0/RSP0/CPU0:router(config-int-nv-cfm)# leve 0	Specifies the CFM level. It ranges from 0 to 7.
Step 7	continuity-check interval time Example: RP/0/RSP0/CPU0:router(config-int-nv-cfm)# continuity-check interval 100ms	(Optional) Enables Continuity Check and specifies the time interval at which CCMs are transmitted. Note Level values of 0 to 7 and interval values of 3.3 ms, 10 ms, 100 ms, 1 sec, 10 sec, 1 min and 10 min are allowed while configuring Ethernet CFM.
Step 8	end or commit Example: RP/0/RSP0/CPU0:router(config)# end or RP/0/RSP0/CPU0:router(config)# commit	Saves configuration changes. • When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

	Command or Action	Purpose
		<ul style="list-style-type: none"> - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Plug and Play Satellite nV Switch Turn up: (Rack, Plug, and Go installation)

1. Unpack the satellite rack, stack, and connect to the power cord.
2. Plug in the qualified optics of correct type into any one or more of the SFP+ slots and appropriate qualified optics into SFP+ or XFP slots on the host. Connect through the SMF/MMF fiber.

When Cisco NCS 5000 Series Satellite is used as satellite, plug in the Cisco NCS 5002/NCS 5001 qualified 100G optics into the first QSFP28 slot (100GigE 0/0/1/0) and a qualified CPAK optics on the host. Connect through a regular SMF fiber. Only one time bootstrap requires the 1st port. Otherwise, any of the 100GigE ports on the satellite can be used.

For 10G and 40G ICLs, a similar bootstrap on fixed ports is needed for plug and play operation in case of first time boot up after a factory reset. The highest 10G port turns on for plug and play in factory reset mode.



Note The Cisco NCS 5000 Series Satellite is shipped in a factory mode. When the satellite is discovered for the first time, it would automatically detect itself as a satellite on receiving probes from the host and would go for a reset and come up in satellite mode. This would be transparent to the user and happens only on a first time discovery of a Cisco NCS 5000 Series satellite. Subsequent reloads or image upgrades will not trigger or require an extra reload.



Note The nV system can use Cisco ASR 9904, Cisco ASR 9006, Cisco ASR 9010, Cisco ASR 9901, Cisco ASR 9906, Cisco ASR 9910, Cisco ASR 9912 and Cisco ASR 9922 Series Routers as hosts. The Cisco ASR 9000v Satellite or the Cisco NCS 5000 Series Satellite can be used as satellite devices.

To configure wavelength on DWDM SFP+, use the following CLI command on satellite console:

```
test dwdm wavelength set ppmId wavelength_channel_number
```



Note ppmId = port number -1

The following example shows how to configure wavelength channel 20 on port 45.

```
Satellite#test dwdm wavelength set 44 20
```

To see the configured wavelength, use the following CLI command on satellite console:

- **show satellite dwdm-dump ppmId**
- **show satellite inventory port 45**



Note It is mandatory to configure the same wavelength on both hosts and satellite, you can follow the same steps above on the hosts.

3. Configure the host for nV operations as described in the sections [Defining the Satellite nV System, on page 31](#), [Configuring the Host IP Address, on page 34](#) and [Configuring Inter-Chassis Links and IP Connectivity, on page 35](#). Configure the satellite nV system through CLI or XML on the host on 100GigE ports for Cisco NCS 5002 Satellite.
4. Power up the chassis of the satellite device.



Note For power supply considerations of ASR 9000v, refer to the *Appendix C, Cisco ASR 9000 and Cisco CRS Satellite Systems* of the *Cisco ASR 9000 Series Aggregation Services Router Hardware Installation Guide* online.

5. You can check the status of the satellite chassis based on these chassis error LEDs on the front face plate.
 - If the Critical Error LED turns ON, then it indicates a serious hardware failure.
 - If the Major Error LED turns ON, then it indicates that the hardware is functioning well but unable to connect to the host.
 - If the Critical and Major LEDs are OFF, then the satellite device is up and running and connected to the host.
 - You can do satellite ethernet port packet loopback tests through the host, if needed, to check end to end data path.



Note When the satellite software requires an upgrade, it notifies the host. You can do an inband software upgrade from the host, if needed. Use the **show nv satellite status** on the host to check the status of the satellite.



Note For the satellite image upgrade to work, you must ensure that the management-plane CLI is not configured on the Cisco ASR 9000 Series Router. If it is configured, then you need to add this exception for each of the 10GigE interfaces, which are the satellite ICLs. This is not required for auto IP configurations from Cisco IOS XR Software Release 5.3.2.

Synchronous Ethernet (SyncE) Offload in nV System

Cisco IOS XR Software Release 5.2.0 supports Synchronous Ethernet (SyncE) offload, a physical layer frequency protocol used to provide frequency synchronization in an nV satellite system both from the host to the satellites, and from the satellites to downstream devices.

Automatic syncE offload configuration is available on the host and syncE is disabled on the satellite until Ethernet Synchronization Messaging Channel (ESMC) messages are received from the host. When an ESMC message is received, the satellite enables syncE on the fabric and the access interfaces.

From Cisco IOS XR Software Release 6.2.1 onwards, per-fabric syncE offload configuration is available on the host for the Cisco NCS 5000 Series Satellite, to explicitly enable syncE on a per-fabric basis. The Cisco NCS 5000 Series Satellite inherits the per-fabric syncE offload configuration from the host and applies it to the corresponding satellite ICL port. Following this, all the access interfaces on the satellite which are mapped to that specific ICL port start to transmit SyncE signals.

The **show frequency synchronization interfaces** and **clear frequency synchronization esmc statistics** commands on the host are extended to include satellite access ports.

From Cisco IOS XR Software Release 6.3.1 onwards, per-access-interface syncE offload configuration is available on the host for the Cisco NCS 5000 Series Satellite. The syncE offload configuration on the host is no longer automatically fanned out to all the access interfaces which are mapped to the specific ICL port. Instead, the feature is configured explicitly on the required access interfaces and the ICL ports depending on the configuration on the host.

Different phases of the syncE offload feature over the releases

Table 3: Comparison of the syncE offload feature over the Cisco IOS XR Software Releases

Cisco IOS XR Software Release	Prior to Cisco IOS XR Software Release 6.2.1	Cisco IOS XR Software Release 6.2.x	Cisco IOS XR Software Release 6.3.1 onwards
SyncE offload type	Automatic - SyncE offload is enabled on the satellite on receipt of the ESMC messages	Per-fabric - SyncE offload is enabled on the ICL and all cross-linked access interfaces.	Per-access-interface - SyncE offload is enabled on the ICLs and/or access ports depending on configuration used.
Cisco ASR 9000v Satellite SyncE offload support	Automatic syncE offload is supported	Per-fabric syncE offload is not supported. Only Automatic SyncE offload is supported.	Per-access-interface syncE offload is not supported. Only Automatic SyncE offload is supported.
Cisco NCS 5000 Series Satellite SyncE offload support	Automatic syncE offload is not supported	Per-fabric syncE offload is supported.	Per-access-interface syncE offload is supported.

In the case when the host is running an image prior to Cisco IOS XR Software Release 6.3.1 and the Cisco NCS 5000 Series Satellite has the Cisco IOS XR Software Release 6.3.1 or later image, the satellite maintains full backwards compatibility and responds to syncE offload messages from the host by applying the appropriate

configuration. For example, syncE configure messages sent for an ICL will automatically fan-out the configuration to any of the cross-linked satellite nV access-interfaces.

Alternately, in the case when the host is running the Cisco IOS XR Software Release 6.3.1 or later image and the Cisco NCS 5000 Series Satellite has an image that is prior to Cisco IOS XR Software Release 6.3.1, the host determines the satellite's capability upon establishing a connection and the satellite drops any syncE specific messages from the host that it cannot interpret



Note SyncE offload is supported on the nV system with the Cisco ASR 9000 Series Router host running Cisco IOS XR 64-bit operating system from the Cisco IOS XR Software Release 6.5.1 onwards.

Restrictions of SyncE Offload on the nV System

The following are some of the restrictions of Synchronous Ethernet offload within an nV system:

- Only limited syncE offload support is provided to customers in release 5.2.0.
- ASR9000v and NCS 5000 are the supported satellites.
- Hub-and-spoke (Dual-home) is the only supported topology.
- Physical and Bundle ICLs are supported.
- Host cannot synchronize to the satellite.
- Minimal configuration support.
- Minimum show command support.
- Only supported for directly connected satellites in hub-and-spoke topologies.
- Commands to configure the host's fabric interfaces as syncE inputs are not permitted, as synchronizing the host to one of its sites is not supported.
- Application of Frequency Synchronization configuration commands on the satellite access ports is not permitted.
- SyncE offload features will not work on Copper SFPs and GLC-GE-100FX SFPs.
- SyncE offload is supported only on 10G ICL ports of ASR9000v1/v2 satellites operating in 10G mode.
- SyncE offload is not supported on 1G ICL ports and 10G ICL ports operating in 1G mode of ASR9000v1/v2 satellites.

Restrictions of SyncE Offload on the Cisco NCS 5000 Series Satellite

The following are some of the restrictions of syncE offload on Cisco NCS 5000 Series Satellite:

- SyncE offload on Cisco NCS 5000 Series satellite is only supported for a directly connected Hub-and-Spoke topology with no intermediate devices between the host and the satellite. The mode can either be a single or a dual head nV system.
- EDC ports cannot be used to receive SyncE signals on Cisco NCS 5000 series satellite. There should be at least one 100G ICL or a 10G ICL other than the EDC ports to receive SyncE from the Host. On Cisco

NCS 5002 satellite, EDC ports are from ports 40-79 and on Cisco NCS 5001 satellite, they are from ports 16-39.

- There is no support to receive clocking from the satellite or satellite access towards the host. While the host side configuration for receive isn't blocked, only transmit is supported.
- Clock selection is done only based on the best QL value. Priority is not considered. If two clock sources are received with same QL value, then one clock source is selected arbitrarily.
- Logs from Cisco NCS 5000 Series Satellite frequency synchronization event changes or errors will not be sent to the host logging buffers.
- SyncE input configuration options, such as selection input or quality receive, are not supported on satellite nV access-interfaces.
- The only syncE offload configuration options that are available on the satellite ICL ports are **enable** and **disable**.
- SyncE offload is not supported with 40GE ICL ports between the Cisco ASR 9000 Series Router host and the Cisco NCS 5000 Series satellite.

Hub-and-Spoke Topology for SyncE Offload

A hub-and-spoke topology is the simplest topology where a centralized hub is connected to the peripherals called "spokes" and it can be used for syncE offload in an nV system. In this topology, the satellites are directly connected to the host by using several links or a single link bundle. The methods are as follows:

- **Satellite connected with several devices**—The satellites are connected to the host through several links.
- **Satellites connected with a link bundle**— The satellites are connected to the host through a single link bundle.

The Inter-Chassis Links (ICLs) appear as point-to-point links for the host and the satellite. When you set up a dual-host satellite, ensure that you apply identical configuration on each host to avoid synchronizing the configuration between the two hosts.

Configuring SyncE Offload on Cisco ASR 9000 Series Router Hosts

To implement syncE offload:

- The host is configured to derive a frequency signal from the external clock source and to provide frequency synchronization to the satellites using the host's fabric interfaces.
- Satellite based configuration is not required. The user must configure the syncE offload feature considering the host as a standalone or non-nV device. The user must configure syncE offload on the fabric interfaces using the existing syncE configuration.



Note In a dual host system, the syncE offload configuration on both the hosts must be identical to ensure consistency if one of the hosts fail.

- After receiving the valid SyncE signal and the Quality Level (QL) that the associated Ethernet Synchronization Message Channel Synchronization Status Messages (ESMC SSMs) over one or more

of its fabric links, the satellite selects one of the fabric links to synchronize to, and uses this timing stream to drive SyncE output on its access interfaces.



Note There is no configuration that controls the behavior of the satellite device and synchronization is enabled automatically on both its fabric interfaces and access interfaces after SyncE information is received.

- You can configure any downstream device to synchronize with this syncE signal.

Basic Configuration:

- Configuration of syncE offload on the host(s) includes:
 - Configuring external clock source as frequency input.
 - Configuring ICL interfaces as syncE outputs, with ESMC output.
- SyncE is disabled by default. It will be automatically enabled on the satellite after ESMC packet is received.
- Use syncE specific show commands on the host to verify the state of the satellite. The **show frequency synchronization interfaces** command is extended to show satellite access port information.

Configuration Examples of SyncE Offload on Cisco ASR 9000 Series Router Hosts

Dual-home Satellite Configuration:

The following are configuration examples for the dual-home satellite configuration.

Configuring the Cisco ASR 9000v Satellite in Host A:

```
nv
 satellite 100
   type asr9000v
   ipv4 address 100.0.0.3
   description sat100
   serial-number CAT1708U0NA
 !
 !
```

Configuring the Cisco ASR 9000v Satellite in Host B:

```
nv
 satellite 100
   type asr9000v
   ipv4 address 100.0.0.3
   description sat100
   serial-number CAT1708U0NA
 !
 !
```

Configuring the fabric interface for the satellite on Host A:

```

interface Loopback100
  ipv4 address 100.0.0.1 255.255.255.0
!
interface Bundle-Ether100
  ipv4 point-to-point
  ipv4 unnumbered Loopback100
  nv
    satellite-fabric-link satellite 100
    redundancy
      iccp-group 1
    !
    remote-ports GigabitEthernet 0/0/0-43
  !
!
!
interface TenGigE0/0/0/0
  bundle id 100 mode on
!

```

Configuring the fabric interface for the satellite on Host B:

```

interface Loopback100
  ipv4 address 100.0.0.2 255.255.255.0
!
interface TenGigE0/0/0/0
  ipv4 point-to-point
  ipv4 unnumbered Loopback100
  nv
    satellite-fabric-link satellite 100
    redundancy
      iccp-group 1
    !
    remote-ports GigabitEthernet 0/0/0-43
  !
!
!

```

Configuring ICCP Redundancy and MPLS LDP on Host A:

```

redundancy
  iccp
    group 1
      member
        neighbor 70.0.0.2
    !
  !
!
!
mpls ldp
  router-id 70.0.0.1
  address-family ipv4
    discovery targeted-hello accept
  !
  interface TenGigE0/0/0/1
  !
!

```

Configuring ICCP Redundancy and MPLS LDP on Host B:

```

redundancy
  iccp
    group 1
      member
        neighbor 70.0.0.1
    !
  !
!

```

```

!
!
mpls ldp
router-id 70.0.0.2
address-family ipv4
discovery targeted-hello accept
!
interface TenGigE0/0/0/1
!
!

```

Configuring ICCP Interface IP Address on Host A:

```

interface TenGigE0/0/0/1
ipv4 address 70.0.0.1 255.255.255.0
!

```

Configuring ICCP Interface IP Address on Host B:

```

interface TenGigE0/0/0/1
ipv4 address 70.0.0.2 255.255.255.0
!

```

Configuration to enable SyncE offload with Cisco NCS 5000 Series Satellite

The first step is setting up the host to transmit syncE signals including ESMC to the satellite. This is same as configuring syncE offload on Cisco ASR9000 Series Router hosts. The configuration involves turning on syncE transmit on the ICL ports outside the nV mode and on the members in case of a bundle ICL.

The second step is to enable Cisco NCS5000 Series satellite to receive SyncE signals from the ICL ports and transmit them over the access ports mapped to that ICL port. A new configuration on the host to enable syncE offload on Cisco NCS 5000 Series Satellite on a per-fabric basis is added.

Below configuration turns on syncE receive on the ICL port and syncE transmit on all the access ports mapped to this ICL port on the Cisco NCS 5000 Series Satellite.

```

Router(config)# interface <interface name>
Router(config-if)# nv satellite-fabric-link satellite <satellite ID>
Router(config-satellite-fabric-link)# frequency synchronization

```



Note If the first step is skipped, then the clock in Cisco NCS 5000 Series Satellite will be in free run mode. If the first step is done after second step, then the Cisco NCS 5000 Series Satellite will start in free run mode but will lock to the ICL as input source when the host starts transmitting the clock output.

If the second step is skipped, then the Cisco NCS 5000 Series Satellite will remain SyncE opaque and there won't be any clock output on its access ports.

For bundle ICL, syncE offload is configured per member, although the offload configuration in the nV mode is still configured on the bundle ICL nV sub mode configuration.

Running Configuration for enabling per-fabric syncE offload with Cisco NCS 5000 Series Satellite:

```

nv satellite 100
type ncs5002

frequency synchronization
quality itu-t option 1

```



```

interface loopback0
  ipv4 address 10.2.3.5

interface TenGigE 0/0/0/1
  frequency synchronization
  ipv4 point-to-point
  ipv4 unnumbered loopback 0
  nv satellite-fabric-link satellite 100
  frequency synchronization
  remote-ports TenGigE 0/0/0

clock-interface sync 0 location 0/RSP0/CPU0
  port-parameters bits-in 2m
  frequency synchronization
  selection input
  ssm disable
  quality-level receive exact itu-t option 1 PRC

```

Configuring Per-Access-Interface SyncE Offload on the Cisco NCS 5000 Series Satellite

SyncE offload per-access-interface configuration allows the user to enable syncE offload explicitly on specific satellite nV access interfaces. This enhancement is available from Release 6.3.1 onwards. Unlike prior releases, syncE offload configuration on the ICL ports is no longer automatically configured on all the cross-linked satellite access interfaces.

In addition to the syncE offload configurations in the previous section, the below configuration is also required to enable per-access-interface syncE offload on the Cisco NCS 5000 Series Satellite.

```

Router(config)# interface <satellite nV access interface name>
Router(config-if)# frequency synchronization
Router(config-if-freqsync)# ssm disable
Router(config-if-freqsync)# quality transmit { exact | highest | lowest }
  itu-t option { 1 | 2 generation 1 | 2 generation 2 } { DNU | PRC | ... } ...

```

Running Configuration for enabling per-access-interface syncE offload with Cisco NCS 5000 Series Satellite:

```

nv satellite 100
  type ncs5002

frequency synchronization
  quality itu-t option 1

interface loopback0
  ipv4 address 10.2.3.5

interface TenGigE 0/0/0/1
  frequency synchronization
  ipv4 point-to-point
  ipv4 unnumbered loopback 0
  nv satellite-fabric-link satellite 100
  frequency synchronization
  remote-ports TenGigE 0/0/0

clock-interface sync 0 location 0/RSP0/CPU0
  port-parameters bits-in 2m
  frequency synchronization
  selection input
  ssm disable

```

```

quality-level receive exact itu-t option 1 PRC

interface TenGigE 100/0/0/0
frequency synchronization
ssm disable
quality transmit exact itu-t option 1 PRC

```

Minimum Active Links on nV System

The Minimum Active Links feature allows the user to configure the minimum number of bundle members which must be active in a bundle satellite-fabric-link in order for traffic to go through that bundle.

There are mainly 2 types of Minimum Active Links on the nV System.

- Soft Minimum Active Links for Dual Home mode
- Hard Minimum Active Links for Single Host mode



Note It is not recommended to configure both Soft and Hard Minimum Active Links together on the same ICL.

Soft Minimum Active Links for Dual Home Mode

The Soft minimum active links feature for dual home mode allows you to configure a minimum number of active links in a bundle satellite-fabric-link. Hence, in a Dual Home setup, if the number of active links to the active host is less than the configured minimum due to any failure on member links, then the satellite failover to the standby host. When soft minimum active links is configured, failover is executed if another suitable host is available or else the traffic is left unaffected. In the case where both the active and standby hosts have less active links than the configured values, then no switch over occurs.



Note If a split brain event occurs followed by the number of active links on the active host dropping below the configured amount, then the satellite does not failover unless there are no active links remaining.

Configuring Soft Minimum Active Links on Dual Home Mode

The Soft Minimum Active links can be configured on the host as below:

```

RP/0/RSP0/CPU0:router(config)# interface <Bundle-ether ICL interface name>
RP/0/RSP0/CPU0:router(config)# nv satellite-fabric-link satellite <satellite ID>
RP/0/RSP0/CPU0:router(config)# redundancy minimum preferred links <value>

```

Configuration example for Soft Minimum Active Links for Dual Home Mode

```

RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether1
RP/0/RSP0/CPU0:router(config-if)# nv satellite-fabric-link satellite 100
RP/0/RSP0/CPU0:router(config-satellite-fabric-link)# redundancy minimum preferred links 2
RP/0/RSP0/CPU0:router(config-satellite-fabric-link)# commit

```

Hard Minimum Active Links for Single Host Mode


The Hard Minimum Active links feature for Single Host mode allows the user to configure the number of bundle members which must be active in order for traffic to go through that route. This is achieved by shutting down the access ports on the satellite should the threshold not be satisfied. This allows the Customer Edge devices connected to the satellite to immediately detect the degradation and reroute traffic accordingly.

Detection of active links falling below the configured value is done on both the host and the satellite.

• Host Side Detection

The host monitors the state of bundle members and when the number of active links falls below the configured value, it requests the satellite to bring down the access ports, over the L1 channel. This is implemented on both Cisco ASR 9000v Satellite and Cisco NCS 5000 Series Satellite.

Figure 8: Host Side Detection of Active Links falling below the configured value




In the above figure, the traffic can pass through 2 possible routes from the Core router to the Customer Edge (CE) device. The route taken by the initial traffic flow is indicated in the above figure. Host A has a hard minimum active links configuration of 3 on its bundle ICL with Satellite A. If one of the bundle members of the Bundle ICL goes down, the following sequence of events take place:

1. Host A detects that the bundle member of the bundle ICL has gone down. Now, the number of active links are below the hard minimum active links configuration.
2. Host A updates Satellite A and the satellite brings down the satellite nV access-interfaces.
3. Loss of signal is detected on the CE device and the traffic takes the alternate route as indicated by the "rerouted traffic flow" in the above figure.

• Satellite Side Detection

The host sends the hard minimum active links configuration to the satellite. The satellite monitors the state of the ICL bundle member links and brings down access ports on crossing threshold for the number of hard minimum active links. Access port state is then sent to the host. This improves convergence times, and ensures that the access ports are still shut down on the satellite if the TCP connection is down. The satellite side detection is implemented only on the Cisco NCS 5000 Series Satellite.

Figure 9: Satellite Side Detection of Active Links falling below the configured value



For example, in the same topology as before, these events take place during the satellite side detection.

1. Host A informs the Satellite A of the hard minimum active links configuration when the satellite is in connected state.
2. When the bundle member goes down, Satellite A detects it. Now, the number of ICL bundle member links is below the hard minimum active links value.
3. Satellite A immediately brings down the satellite nV access interfaces and updates the Host A.
4. Host A detects that the satellite nV access interfaces have gone down and makes the routing update. Though it may have already done this during the host side detection. And as in host side detection, when the satellite nV access interfaces are down, the interfaces on the CE device also go down.



Note In the event that the number of ICL bundle members links becomes lower than the hard minimum active links value, the satellite side detection is faster compared to the host side detection. And thus satellite side detection results in faster traffic convergence. Since Cisco NCS 5000 Series Satellite supports satellite side detection, the traffic convergence will be faster than with a Cisco ASR 9000v Satellite where only host side detection is supported.

Configuring Hard Minimum Active Links on Single Host Mode

The Hard Minimum Active links can be configured on the host as below:

```
Router(config)# interface Bundle-Ether1
Router(config-if)# nv satellite-fabric-link satellite 100
Router(config-satellite-fabric-link)# minimum required links 2
Router(config-satellite-fabric-link)# commit
```

Verification of Hard Minimum Active Links for Single Host Mode

The Hard Minimum Active Links can be verified in the output of the command **show nv satellite status**

```
Router# show nv satellite status satellite 100
Satellite 100
-----
Status: Connected (Stable)
Type: ncs5002
```

```

Displayed device name: Sat100
MAC address: 00ba.c000.000b
IPv4 address: 10.0.100.1 (auto, VRF: **nVsatellite)
Serial Number: INS1719U1Mhsatellite
Remote version: Compatible (latest version)
  b2.3.1
  Satellite Emulator
Received candidate fabric ports:
  nVFabric-GigE0/0/0-3,6-9 (permanent)
  nVFabric-TenGigE0/0/0-3,6-9 (permanent)
Configured satellite fabric links:
Bundle-Ether1
-----
Status: Satellite Ready
Required Links: Satisfied (Configured 2, Active 3)
No Port Ranges
Discovered satellite fabric links:
  TenGigE0/1/0/0: Satellite Ready; No conflict
  TenGigE0/1/0/1: Satellite Ready; No conflict
  TenGigE0/1/0/2: Satellite Ready; No conflict

```

Upgrading and Managing Satellite nV Software

Satellite software images are bundled inside a PIE and the PIE name is dependent on the type of satellite, such as **asr9k-9000v-nV-px.pie** within the Cisco ASR 9000 Series Router package. The Cisco IOS XR software production SMU tool can be used to generate patches for the satellite image in the field to deliver bug fixes or minor enhancements without requiring a formal software upgrade.

For more information on auto image upgrade supported from Cisco IOS XR Software Release 5.3.2 and later, see [nV Satellite Auto Image Upgrade, on page 29](#).

Image Upgrade for nV Satellite

Image Upgrade for Cisco ASR 9000v Satellite

The **asr9k-asr9000v-nV-px.pie** contains two sets of binaries, namely, the intermediate binaries and the final binaries. When a Satellite nV system running Cisco IOS XR Software prior to Cisco IOS XR Software Release 5.1.1 is upgraded to Cisco IOS XR Release 5.1.1, the satellite downloads the intermediate binaries and reloads as per the instructions of the operator. These intermediate binaries include the logic to request the file name from the host rather than hard coding the file name. Also, they automatically trigger the second download (final binaries) without requiring manual intervention. When a Satellite nV system running Cisco IOS XR Software Release 5.1.1 or later is upgraded to Cisco IOS XR Software Release 5.3.x, the satellite does not require intermediate image transfers and reloads once with the latest image.



Note The **show nv satellite status** command does not display the intermediate version. However, it displays the final Cisco IOS XR Software Release 5.1.1 and prompts for any further upgrade. But, internally two reloads happen. On the other hand, when you upgrade from Cisco IOS XR Release 5.1.x to future releases, two reloads do not occur. When you downgrade, the system does not downgrade two releases. CFM needs to be disabled if image upgrade is done over Bundle ICL for Cisco ASR9000v satellite.



Note An auto transfer internal message comes up when the second software reload happens, which requires no explicit user-intervention.



Note For upgrading from a release prior to Cisco IOS XR Software Release 5.1.1, the Satellite nV System must be connected in the Hub and Spoke topology as the previous releases do not support the dual head mode.

Image Upgrade for Cisco NCS 5000 Series Satellite

The existing nV pie based satellite image packaging model is extended to Cisco NCS 5000 series satellites from Cisco IOS XR Software Release 6.0.1. The image binary including all device FPD upgrade is done through **asr9k-ncs500x-nV-px.pie** as in the case of Cisco ASR9000v. The upgrade image for Cisco NCS 5000 Series satellites including firmware is available as an nV pie on the Cisco ASR 9000, which once activated can be used to push images to the relevant satellites in the connected state.



Note A message indicating a successful image upgrade is displayed in **show nv satellite status** command even if the image upgrade has failed on the satellite. However the **show nv satellite status** command will show that satellite does not have the latest image even after the upgrade.

Whereas, in Cisco IOS XR Software Release 5.3.0, in case of a failure, the host tries to transfer the image indefinitely; the upgrade will recover with subsequent retries if the failure is transient. A syslog message on the host indicates that retry is in progress.

On the other hand, for Cisco IOS XR Software Release 5.3.1 and later, in case of a failure, the host tries to transfer the image indefinitely. Additionally, **install nv satellite <sat-id> abort** command can be used to stop any active image upgrade operations on a particular satellite.



Note For the Cisco NCS 5000 Series Satellite, the native satellite image cannot be added and activated on the Cisco ASR 9000 Series Router host running the IOS XR 64 bit operating system due to its large size. Hence for loading a new image on the Cisco NCS 5000 Series Satellite, the procedure mentioned in the section [Installing native images and SMUs for Cisco NCS 5000 Series Satellite, on page 56](#) needs to be followed. Transfer and activation of the Cisco NCS 5000 Series Satellite image can be initiated either directly from the TFTP or from the hard disk of the Cisco ASR 9000 Series Router host.

Prerequisites

- You must have installed the satellite installation procedure using the Plug and Play Satellite installation procedure. For more information, check the topic [Plug and Play Satellite nV Switch Turn up: \(Rack, Plug, and Go installation\)](#), on page 41 in this chapter.

Installing a Satellite

To download and activate the software image on the satellite, use the **install nv satellite** satellite ID / all **transfer/activate** commands. The **transfer** command downloads the image to the satellite.

When the **transfer** command is followed by the **activate** command, the software is activated on the satellite.



Note For ASR9000v1/v2 satellites with ICL bundle links, the CFM interval should be more than or equal to 1 minute for successful image transfer and satellite upgrade.

```
RP/0/RSP0/CPU0:sat-host# install nv satellite 100 transfer

Install operation initiated successfully.
RP/0/RSP0/CPU0:sat-host#RP/0/RSP0/CPU0:May  3 20:12:46.732 : icpe_gco[1146]:
%PKT_INFRA-ICPE_GCO-6-TRANSFER_DONE : Image transfer completed on Satellite 100

RP/0/RSP0/CPU0:sat-host# install nv satellite 100 activate

Install operation initiated successfully.
LC/0/2/CPU0:May  3 20:13:50.363 : ifmgr[201]: %PKT_INFRA-LINK-3-UPDOWN : Interface
GigabitEthernet100/0/0/28, changed state to Down
RP/0/RSP0/CPU0:May  3 20:13:50.811 : invmgr[254]: %PLATFORM-INV-6-OIROUT : OIR: Node 100
removed
```



Note If the **activate** command is run directly, then the software image is transferred to the satellite and also activated.

```
RP/0/RSP0/CPU0:sat-host# install nv satellite 101 activate

Install operation initiated successfully.

RP/0/RSP0/CPU0:sat-host#RP/0/RSP0/CPU0:May  3 20:06:33.276 : icpe_gco[1146]:
%PKT_INFRA-ICPE_GCO-6-TRANSFER_DONE : Image transfer completed on Satellite 101
RP/0/RSP0/CPU0:May  3 20:06:33.449 : icpe_gco[1146]: %PKT_INFRA-ICPE_GCO-6-INSTALL_DONE :
Image install completed on Satellite 101
RP/0/RSP0/CPU0:May  3 20:06:33.510 : invmgr[254]: %PLATFORM-INV-6-OIROUT : OIR: Node 101
removed
```



Note For the satellite image upgrade to work, you must ensure that the management-plane CLI is not configured on the Cisco ASR 9000 Series Router . If it is configured, then you need to add this exception for each of the satellite ICLs. This is not required for Auto IP configurations from Cisco IOS XR Software Release 5.3.2.

Ensure that the `tftp homedir`, `tftp vrf default ipv4 server homedir disk0` is not configured on the host when using manual IP default configuration, because this may cause the image transfer to fail.

You can include the exception using this CLI:

```

control-plane
management-plane
inband
!
!
interface TenGigE0/0/0/5 <=== To enable TFTP on nV satellite ICL
allow TFTP

```

If you do not include this exception, then the image download and upgrade fails.

Installing native images and SMUs for Cisco NCS 5000 Series Satellite

You can see the [Native Image and SMU Push Support for Cisco NCS 5000 Series Satellites, on page 23](#) section under the **Features Supported in the Satellite nV System** for description of this feature. These are the SMU/image install command syntax:

```

Router# install nv satellite {all | id-range | id } transfer file file-name

Router# install nv satellite {all | id-range | id } activate | deactivate | remove package
package-name
activate Install a new image or package on the satellite
deactivate Deactivate a package on the satellite
remove Remove a deactivated package on the satellite

```

The *file-name* argument specifies the name of the file to be transferred, which may be in a remote location (to be copied through TFTP, FTP or SFTP) or may be present on the host device. Note that this may only be a single package (batched transfer of tar balls and multiple packages stated consecutively may be supported in future releases). The *package-name* argument specifies the name of the package on the satellite that is to be activated, deactivated, or removed. It is reported in the logs for the transfer operation and in the satellite software version show commands (for example, after transferring a package, the **inactive** command can be used to find the package name to activate).

From Cisco IOS XR Software Release 6.3.1 onwards, native Cisco NCS 5000 Series Satellite System Admin VM and Host VM SMUs appear in the list of packages in the command **show nv satellite install packages** and also as one of the available package options in the command **install nv satellite satellite-id activate package**.

```

Router# show nv satellite install packages satellite 200
Active Packages: 2
    ncs5k-xr-6.3.1 version=6.3.1 [Boot image]
    ncs5k-sysadmin-6.3.1 version=6.3.1 [Boot image]
Inactive Packages: 1
    ncs5k-sysadmin-6.3.1.CSCuv34432.rpm
Committed Packages: 2
    ncs5k-xr-6.3.1 version=6.3.1 [Boot image]
    ncs5k-sysadmin-6.3.1 version=6.3.1 [Boot image]

```

The base package contains two entries, one for the native Cisco NCS 5000 Series Satellite Host VM and the other one which contains the string "sysadmin" for the native Cisco NCS 5000 Series Satellite System Admin VM.

Install Update Command

Satellite nV SMUs and base packages can be transferred and activated in a single operation using the install update command. The command supports a list of files as well as references (refer section [Install Reference](#)

Command, on page 57 in order to create references). And the files can be located on a local disk or the tftp server.

The SMU/image install update command syntax is as follows:

```
Router# install nv satellite {all | <id range> | <id>} update {file <file> | reference <reference>}
```

The example for install update command is as follows:

```
Router# install nv satellite 100 update file disk0:ncs5k-mps-2.2.0.0-r631.x86_64.rpm
tftp://172.27.18.8/ncs5k-mgbl-2.2.0.0-r631.x86_64.rpm
```

```
Router# install nv satellite all update reference MyRef
```

Install Reference Command

A reference name can be given for a set of files using the install references command. When a reference is created the specified files are copied to the host. Other install operations such as transfer, update or replace, can then be initiated using the reference name.

The SMU/image install references command syntax is as follows:

```
Router# install nv satellite reference <name of reference> create file <file1> <file 2>
```

When a reference is being created, the full list of files which should be included in the reference should be specified. Recreating a reference with the same name will delete the previous files and create a reference with the new files.

The example for install references command is as follows:

```
Router# install nv satellite reference MyRef create file
tftp://172.27.18.8/ncs5k-mps-2.2.0.0-r631.x86_64.rpm
tftp://172.27.18.8/ncs5k-mgbl-2.2.0.0-r631.x86_64.rpm
```

It is also possible to create empty references as follows:

```
Router# install nv satellite reference AnotherRef create empty
```

Verification

The command **show nv satellite install references** command can be used to verify the references that were created.

The command syntax is as follows:

```
Router# show nv satellite install references [name <name of reference>]
```

The examples are as below:

```
Router# show nv satellite install references
Reference: MyRef, file count: 2
    ncs5k-mps-2.2.0.0-r631.x86_64.rpm
    ncs5k-mgbl-2.2.0.0-r631.x86_64.rpm
Reference: AnotherRef, file count: 0
    Empty
Router# show nv satellite install references name MyRef
Reference: MyRef, file count: 2
    ncs5k-mps-2.2.0.0-r631.x86_64.rpm
    ncs5k-mgbl-2.2.0.0-r631.x86_64.rpm
```

Monitoring the Satellite Software

Status Check

To perform a basic status check, use the **show nv satellite status brief** command.

```
RP/0/RSP0/CPU0:router# show nv satellite status brief
```

Sat-ID	Type	IP Address	MAC address	State
100	asr9000v	101.102.103.105	dc7b.9426.1594	Connected (Stable)
200	asr9000v	101.102.103.106	0000.0000.0000	Halted; Conflict: no links configured
400		194.168.9.9	0000.0000.0000	Halted; Conflict: satellite has no type configured

These show commands display the status of Cisco NCS 5000 Series satellite:

```
RP/0/RSP0/CPU0:TARDIS# show nv satellite status brief
```

Sat-ID	Type	IP Address	MAC address	State
100	ncs5002	10.0.100.1	c472.95a6.2003	Connected

Check if Upgrade is Required

To check if an upgrade is required on satellite, run the **show nv satellite status satellite *satellite_id***.

The following code block shows the sample output for ASR9000v satellite:

```
Router# show nv satellite status satellite 100
```

```
Satellite 100
-----
State: Connected (Stable)
Type: asr9000v
Description: sat-test
MAC address: dc7b.9427.47e4
IPv4 address: 100.1.1.1
Configured Serial Number: CAT1521B1BB
Received Serial Number: CAT1521B1BB
Remote version: Compatible (latest version)
  ROMMON: 125.0 (Latest)
  FPGA: 1.13 (Latest)
  IOS: 200.8 (Latest)
Configured satellite fabric links:
  TenGigE0/2/0/6
  -----
    State: Satellite Ready
    Port range: GigabitEthernet0/0/0-9
  TenGigE0/2/0/13
  -----
    State: Satellite Ready
    Port range: GigabitEthernet0/0/30-39
  TenGigE0/2/0/9
  -----
    State: Satellite Ready
    Port range: GigabitEthernet0/0/10-19
```

The following code block shows the sample output for Cisco NCS 5000 Series Satellite :

```

Router# show nv satellite status satellite 100
-----
Satellite 100
-----
Status: Connected (Installing new image)
Type: ncs5002
Displayed device name: Sat100
MAC address: c472.95a6.0f25
IPv4 address: 10.0.100.1 (auto, VRF: **nVSatellite)
Serial Number: FOC1919R0WA
Remote version: Compatible (older version)
  IOFPGA: 0.16
  MB_MIFPGA: 0.15
  DB_MIFPGA: 0.15
  BIOS: 1.07
  XR: 6.1.1.08I (Available: 6.1.1.09I)
Received candidate fabric ports:
  nVFabric-TenGigE0/0/78-79 (permanent)
  nVFabric-HundredGigE0/1/0-3 (permanent)
Configured satellite fabric links:
  TenGigE0/0/0/3
-----
Status: Satellite Ready
No Port Ranges

```



Note If the satellite pie is not installed on the Cisco ASR 9000 Series Router Host, then the compatibility status will be shown as unknown as there is no local version to compare against. For the Cisco NCS 5000 Series satellites, in the case where there is no satellite pie on the host and the native image is pushed onto the satellite, this will always be the case. However, in such situations, a "Recommended" version will be displayed for guidance.

Output of Status Check

This example shows the output of `show nv satellite status` command for a Satellite configured in dual home mode.

```

RP/0/RSP0/CPU0:router# show nv satellite status

Satellite 100
-----
Status: Connected (Stable)
Redundancy: Active (Group: 10)
Type: asr9000v
Description: sat100
MAC address: 4055.3958.61e4
IPv4 address: 100.100.1.2 (VRF: default)
Serial Number: CAT1604B1AN
Remote version: Compatible (not latest version)
  ROMMON: 126.0 (Latest)
  FPGA: 1.13 (Latest)
  IOS: 322.5 (Available: 322.3)
Configured satellite fabric links:
  TenGigE0/1/0/0
-----
Status: Satellite Ready
Remote ports: GigabitEthernet0/0/30-43

```

```
TenGigE0/1/0/1
-----
Status: Satellite Ready
Remote ports: GigabitEthernet0/0/0-29
```

This example shows the sample output of a satellite interfaces over redundant ICL.

```
RP/0/RSP0/CPU0:router#show sits interface gig 300/0/0/0

Interface Name           : GigabitEthernet300/0/0/0
Interface IFHandle       : 0x020058a0
Base Interface           : TenGigE0/0/0/0
Base Interface IFHandle  : 0x000001c0
LAG Hash Index          : 2
Published Base interface : TenGigE0/0/0/0
Published Base ifh       : 0x000001c0
Published LAG Hash       : 2
Sat Interface State in DB : Republish Success
```

This command shows a sample output whether the bundle over bundle satellite interface is replicated.

```
RP/0/RSP1/CPU0:router# show bundle infrastructure ea bundle-ether 101 detail

Bundle-Ether101

Node           Platform Information
-----
0/0/CPU0      Ifhandle       : 0x02004d60
               Channel Map  : 0x3

Node: 0/0/CPU0

Member         Platform Information
-----
Gi300/0/0/3   Ifhandle       : 0x020057e0
               Channel Map  : 0x3
               UL Id       : 0
               Base Interfaces
               Count       : 2
               Ifhandle   : 0x000001c0 0x00000140 0x00000000 0x00000000
```



Note In this example output, **Remote version**, **ROMMON**, **FPGA**, and **IOS** must show the latest version. If it does not, an upgrade is required on the satellite. The version numbers displayed are the installed version on the ASR 9000v. If a version number is displayed, instead of latest key word in the above output, that would correspond to the ASR9000v image bundles in the satellite pie.



Note show tech from satellite devices can be pulled out remotely using **show tech-support satellite remote satellite [sat id] file disk0:[filename]** option for offline analysis of the states on the satellite device. This works for Cisco ASR 9000v and Cisco NCS 5000 Series satellites.

Show Commands for Dual Home Mode

The following example show commands used for Dual Home mode.

Dual Home Network Mode

```
RP/0/RSP1/CPU0:Router# show iccp group 10

Redundancy Group 10
  member ip:1.1.1.1 (vkg1), up (connected)
  monitor: route-watch (up)
  No backbone interfaces.
  enabled applications: SatelliteORBIT
  isolation recovery delay timer: 30 s, not running

RP/0/RSP1/CPU0:Router#show nv satellite protocol redundancy

ICCP Group: 10
-----
Status: Connected since 2014/01/22 15:47:35.845
Role: Primary (System MAC: 0000.0001.1234)
Channels:
  Control (0)
  -----
    Channel status: Open
    Messages sent: 8 (4 control), received: 6 (3 control).

  Topology (14)
  -----
    Channel status: Open
    Messages sent: 4 (3 control), received: 11 (0 control).

## active host:

RP/0/RSP1/CPU0:Router# show nv satellite status satellite 200
Satellite 200
-----
Status: Connected (Stable)
Redundancy: Active (Group: 10)
Type: asr9000v
MAC address: 8478.ac01.d2d8
IPv4 address: 192.1.1.200 (VRF: default)
Serial Number: CAT1708U0LV
Remote version: Compatible (latest version)
  ROMMON: 126.0 (Latest)
  FPGA: 1.13 (Latest)
  IOS: 322.6 (Latest)
Configured satellite fabric links:
  TenGigE0/0/0/1
  -----
    Status: Satellite Ready
    Remote ports: GigabitEthernet0/0/0-10
```

##Standby host:

```
RP/0/RSP1/CPU0:Router# show nv satellite status satellite 200

Satellite 200
```

```

-----
Status: Connected (Stable)
Redundancy: Standby (Group: 10)
Type: asr9000v
MAC address: 8478.ac01.d2d8
IPv4 address: 192.1.1.200 (VRF: default)
Serial Number: CAT1708U0LV
Remote version: Compatible (latest version)
  ROMMON: 126.0 (Latest)
  FPGA: 1.13 (Latest)
  IOS: 322.6 (Latest)
Configured satellite fabric links:
  TenGigE0/3/0/6
-----
Status: Satellite Ready
Remote ports: GigabitEthernet0/0/0-10

```

Monitoring the Satellite Protocol Status

To check the status of the satellite discovery protocol, use the **show nv satellite protocol discovery** command.

```
RP/0/RSP0/CPU0:router# show nv satellite protocol discovery brief
```

Interface	Sat-ID	Status	Discovered links
Te0/1/0/0	100	Satellite Ready	Te0/1/0/0
Te0/1/0/1	100	Satellite Ready	Te0/1/0/1

(Or)

```
RP/0/RSP0/CPU0:router# show nv satellite protocol discovery interface TenGigE 0/1/0/0
```

```

Satellite ID: 100
Status: Satellite Ready
Remote ports: GigabitEthernet0/0/0-15
Host IPv4 Address: 101.102.103.104
Satellite IPv4 Address: 101.102.103.105
Vendor: cisco, ASR9000v-DC-E
Remote ID: 2
Remote MAC address: dc7b.9426.15c2
Chassis MAC address: dc7b.9426.1594

```

```
RP/0/RSP0/CPU0:TARDIS# show nv satellite protocol discovery brief
```

Interface	Sat-ID	Status	Discovered links
Hu0/1/0/0	100	Satellite ready	Hu0/1/0/0

To check the status of the satellite control protocol status, use the **show nv satellite protocol control** command.

```
RP/0/RSP0/CPU0:router# show nv satellite protocol control brief
```

Sat-ID	IP Address	Protocol state	Channels
101.102.103.105	101.102.103.105	Connected	Ctrl, If-Ext L1, If-Ext L2, X-link, Soft Reset, Inventory, EnvMon, Alarm

```
RP/0/RSP0/CPU0:shanghai# sh nv satellite protocol control
```

```

Satellite 100
-----
IP address: 101.102.103.105
Status: Connected
Channels:
  Control
  -----
    Channel status: Open
    Messages sent: 24 (24 control), received: 23 (23 control).
  Interface Extension Layer 1
  -----
    Channel status: Open
    Messages sent: 7 (3 control), received: 14 (2 control).
  Interface Extension Layer 2
  -----
    Channel status: Open
    Messages sent: 11 (3 control), received: 10 (2 control).
  Interface Extension Cross-link
  -----
    Channel status: Open
    Messages sent: 4 (3 control), received: 3 (2 control).

```

```
RP/0/RSP0/CPU0:TARDIS#show nv satellite protocol control brief
```

```

Sat-ID  IP Address  Protocol state  Channels
-----
100     10.0.100.1  Connected      Ctrl, If-ExtL1, If-Ext L2, X-link,
                                         VICL, DevMgmt, Inventory, EnvMon,
                                         Alarm, Password, Topology,

```

To check the status of satellite protocol redundancy, use the **show nv satellite protocol redundancy** command.

```
RP/0/RSP0/CPU0:router# show nv satellite protocol redundancy
```

```

ICCP Group: 10
-----
Status: Connected since 2014/01/11 08:44:58.764
Role: Secondary (System MAC: 6c9c.ed23.c4e6)
Channels:
  Control (0)
  -----
    Channel status: Open
    Messages sent: 18 (9 control), received: 24 (12 control).

  Topology (14)
  -----
    Channel status: Open
    Messages sent: 88 (10 control), received: 60 (0 control).

```

Monitoring the Satellite Inventory

You can use the **show inventory | beg SAT** and **show environment nv satellite all** commands in the exec mode to monitor the status of satellite inventory.

```

RP/0/RSP0/CPU0:router# show inventory | beg SAT
NAME: "SAT201/RP0", DESCR: "Cisco NCS 5002 Route processor"
PID: NCS-5002          , VID: V00, SN: FOC19390B85

NAME: "SAT201/PORT-40", DESCR: "Cisco 10GE SFP+ SR Pluggable Optics Module"

```

```

PID: SFP-10G-SR      , VID: V03, SN: FNS17021J11

NAME: "SAT201/PORT-57", DESCR: "Cisco 10GE SFP+ SR Pluggable Optics Module"
PID: SFP-10G-SR      , VID: V03, SN: FNS17021Q1B

NAME: "SAT201/PORT-58", DESCR: "Cisco 10GE SFP+ SR Pluggable Optics Module"
PID: SFP-10G-SR      , VID: V03, SN: JUR19060MEP

NAME: "SAT201/PORT-59", DESCR: "Cisco 10GE SFP+ SR Pluggable Optics Module"
PID: SFP-10G-SR      , VID: V03, SN: SPC17050EU4

NAME: "SAT201/PORT-60", DESCR: "Cisco 10GE SFP+ SR Pluggable Optics Module"
PID: SFP-10G-SR      , VID: V03, SN: SPC17050ETY

NAME: "SAT201", DESCR: "NCS5002 Satellite Chassis"
PID: NCS-5002        , VID: V00, SN: FOC1939R191

NAME: "SAT201/FT0", DESCR: "NCS 5002 Router Fan Back to Front AirFlow"
PID: NCS-5002-FN-BK  , VID: N/A, SN: N/A

NAME: "SAT201/FT1", DESCR: "NCS 5002 Router Fan Back to Front AirFlow"
PID: NCS-5002-FN-BK  , VID: N/A, SN: N/A

NAME: "SAT201/PT0", DESCR: "NCS5K Power Module AC 650W Back to Front AirFlow"
PID: NC5K-PAC-650W-BK , VID: V00, SN: LIT193813NX

NAME: "SAT201/PT1", DESCR: "NCS5K Power Module AC 650W Back to Front AirFlow"
PID: NC5K-PAC-650W-BK , VID: V00, SN: LIT193813MP

```

```

RP/0/RSP0/CPU0:router# show environment nv satellite all ?
current      Current information(cisco-support)
fan          Fan information(cisco-support)
power        Power Supply information(cisco-support)
temperature   Temperature information(cisco-support)
voltage      Voltage information(cisco-support)
|            Output Modifiers
<cr>

```

```

RP/0/RSP0/CPU0:sat-host# show environment nv satellite all fan

```

```

=====
Location      FRU Type          Fan speed (rpm)
              FAN_0    FAN_1
-----
SAT150/FT0    SAT-FANTRAY      5515    5888
SAT150/FT1    SAT-FANTRAY      5515    5863

```

```

RP/0/RSP0/CPU0:TARDIS#show environment nv satellite all voltage

```

```

.....
Location  VOLTAGE          Value      Crit      Minor      Minor      Crit
          Sensor          (mV)      (Lo)      (Lo)      (Hi)      (Hi)
-----
          V3P3VA          3348      3135      3168      3436      3465
          V2P5VA          2525      2375      2400      2600      2625
          V1P2VA          1204      1140      1152      1248      1260
          V5P0           4991      4750      4800      5200      5250
          V3P3           3366      3135      3168      3436      3465
          V2P5           2499      2375      2400      2600      2625
          V1P8           1799      1710      1728      1872      1890
          V1P8_ZR        1810      1710      1728      1872      1890

```


V1P5	1525	1425	1440	1560	1575
V1P5_PCH	1511	1425	1440	1560	1575
V1P05_VTT	1055	997	1008	1092	1103
V1P0	1002	950	960	1040	1050
V1P0_PCH	1001	950	960	1040	1050
V0P8_CPU	802	760	768	832	850
V3P3_BRCM	3326	3135	3168	3436	3465
CPUVCC	862	284	288	1580	1596
TDNT_VDD_V1P0	1005	903	912	1040	1050
TDNT_AVDD_V1P0	1005	950	960	1040	1050
V12P0_MZ	12114	11400	11520	12480	12600
PHY4G_V1P0_MZ	1007	950	960	1040	1050
PHY1_V1P0_MZ	1009	950	960	1040	1050
V1P2_MZ	1218	1140	1152	1248	1260
V2P5_MZ	2521	2375	2400	2600	2625
V3P3_MZ	3388	3135	3168	3436	3465

Monitoring the Satellite Environment

You can use the **show environment nv satellite all temperature** command in the exec mode to monitor the status of satellite environment.

```
RP/0/RSP0/CPU0:router# show environment nv satellite all temperature
=====
Location  TEMPERATURE          Value  Crit Major Minor Minor Major  Crit
          Sensor              (deg C) (Lo) (Lo) (Lo) (Hi) (Hi) (Hi)
-----
SAT150   Inlet                   20    -10  -   -   -   -   -
```

Reloading the Satellite and Monitoring DOM Parameters

In order to reload the satellite device, use the **hw-module satellite satellite id/all reload** command.

```
RP/0/RSP0/CPU0:router# hw-module satellite 101 reload

Reload operation completed successfully.
RP/0/RSP0/CPU0:May  3 20:26:51.883 : invmgr[254]: %PLATFORM-INV-6-OIROUT : OIR: Node 101
removed
```

In order to see the DOM parameters of the SFPs and XSPs or access ports and ICL ports of the satellite, use the **show controllers gigabitEthernet interface phy** command.

For access ports

```
RP/0/RSP0/CPU0:Saturn#show controllers gigabitEthernet 100/0/0/22 phy
Wed Apr  8 17:42:32.100 UTC

Port: 22
Xcvr Type: SFP
Vendor Name: CISCO-FINISAR
CLEI Code: IPUIALJRAA
Part Number: 10-2143-01V01
Product Id: SFP-GE-S

Thresholds:          Alarm High          Warning High          \
Warning Low          Alarm Low          \
Temperature:          109C          103C          \
-13C          -29C          \
Voltage:          3900uV          3700uV          \
```

Port Level Parameters Configured on a Satellite

```

                2900uV                2700uV
                Bias:                15mAmps                12mAmps \
                2mAmps                1mAmps
Transmit Power: 0.63100 mW (-1.99971 dBm) 0.63100 mW (-1.99971 dBm) \
0.07900 mW (-11.02373 dBm) 0.06600 mW (-11.80456 dBm)
Receive Power: 1.25800 mW (0.99681 dBm) 0.79400 mW (-1.00179 dBm) \
0.01500 mW (-18.23909 dBm) 0.01000 mW (-20.00000 dBm)
Temperature: 32 C
Voltage: 3327 uV
Bias: 5 mAmps
Tx Power: 0.28100 mW (-5.51294 dBm)
Rx Power: 0.000 mW (<-40.00 dBm)

```

For ICL port

```

RP/0/RSP0/CPU0:Saturn#show controllers nvFabric-TenGigE 100/0/0/46 phy
Wed Apr 8 17:46:57.045 UTC

```

```

Port: 46
Xcvr Type: SFP
Vendor Name: CISCO-FINISAR
CLEI Code: COUIA75CAA
Part Number: 10-2457-02V02
Product Id: SFP-10G-LR
Thresholds:                Alarm High                Warning High                \
Warning Low                Alarm Low
Temperature:                75C                70C                \
0C                -5C
Voltage:                3630uV                3465uV                \
3135uV                2970uV
Bias:                70mAmps                68mAmps                \
2mAmps                1mAmps
Transmit Power: 2.23800 mW (3.49860 dBm) 1.12200 mW (0.49993 dBm) \
0.15100 mW (-8.21023 dBm) 0.06000 mW (-12.21849 dBm)
Receive Power: 2.23800 mW (3.49860 dBm) 1.12200 mW (0.49993 dBm) \
0.03600 mW (-14.43697 dBm) 0.01400 mW (-18.53872 dBm)
Temperature: 30 C
Voltage: 3366 uV
Bias: 34 mAmps
Tx Power: 0.86300 mW (-0.63989 dBm)
Rx Power: 1.01000 mW (0.04321 dBm)

```

Port Level Parameters Configured on a Satellite

These are the port-level parameters that can be configured on a satellite nV system:

- Admin state (shut and no shut)
- Ethernet MTU



Note For Cisco ASR 9000v access ports, the maximum MTU is 9212 for a hub and spoke topology.

For Cisco NCS 5000 series satellite access ports, the maximum MTU is 9158.

- Ethernet MAC Address.
- Ethernet link auto-negotiation that includes,

- Half and full duplex
 - Link speed
 - Flow control
- Static configuration of auto-negotiation parameters such as speed, duplex, and flow control
 - Carrier-delay



Note The Cisco ASR 9000v satellite does not support asymmetric carrier-delay values for access ports. If the host side configuration includes asymmetric values, the satellite applies the minimum of the two for both the up and down carrier-delay values. If any of the carrier-delay values from the host side is less than 100ms, then the satellite applies the value of 100ms, as that is the smallest supported value for carrier-delay.

The Cisco NCS 5000 Series satellite supports both symmetric and asymmetric carrier-delay values for access ports.

- Layer-1 packet loopback which includes,
 - Line loopback
 - Internal loopback
- All satellite access port features on Cisco ASR 9000 Series Router.

Loopback Types on Satellite Ports

There are two types of loopback interfaces that can be configured on satellite ports. They are,


- Line Loopback
- Internal Loopback

These illustrations show how the loopback interface types function on a satellite.

Figure 10: Line Loopback



Figure 11: Internal Loopback



You can specify the type of loopback to be used, as specified in this example:

```
Interface GigabitEthernet 100/0/0/0
```

```
loopback line | internal
```

Configuration Examples for Satellite nV System

This section contains configuration examples for the Satellite nV system:

Satellite System Configuration: Example

This example shows a sample configuration to configure connectivity for a Satellite System:

Satellite Global Configuration

The satellite ID, type, serial number, description, and satellite IP address are configured in the satellite global configuration sub-mode:

```
nv
 satellite 100
  type asr9000v
  serial-number CAT1521B1BB
  description milpitas bldg20
  ipv4 address 10.0.0.100
 !
 !
```

ICL (satellite-fabric-link) Interface Configuration

On an interface connected to a Satellite (TenGigE or Bundle interface), the ports associated with the satellite-id must be specified. All fabric links connected to the same Satellite must use the same (Host) IPv4 address. This Host IPv4 addresses can be used for the same Host to connect to different Satellites.



Note Before you remove or change a configuration on an ICL interface, shut down the ICL port.

```
interface Loopback1000
 vrf <vrf_name>
  ipv4 address 10.0.0.1 255.0.0.0
 vrf <vrf_name>
 interface TenGigE0/2/1/0
  description To Sat5 1/46
  ipv4 point-to-point
  ipv4 unnumbered Loopback1000
 nv
 satellite-fabric-link satellite 200
  remote-ports GigabitEthernet 0/0/0-30
 !
 !
 !
```



Note To manage satellite traffic, use the IP addresses from the global VRF of the router (shown in the examples). As mentioned in [Satellite Discovery and Control Protocol IP Connectivity](#) section, you can use a private VRF to prevent IP address conflict with global VRF. In such a case, the loopback interface and ICL interface (in the examples) must be assigned to the private VRF dedicated for satellite management traffic.

Satellite Interface Configuration

A Satellite interface can be used like any other regular Gigabit Ethernet interfaces:

```
interface GigabitEthernet200/0/0/0
l2transport
!
!

interface GigabitEthernet200/0/0/0
ip address 99.0.0.1 255.255.255.0
!
!

interface GigabitEthernet200/0/0/2
bundle id 100 mode active
!
!
```

This is a sample satellite interface configuration in the case of a dual home mode on the active and standby hosts:

Active Host

```
interface GigabitEthernet100/0/0/32
ipv4 address 1.1.1.1 255.255.255.0
!
```

Standby Host

```
interface GigabitEthernet100/0/0/32
ipv4 address 1.1.1.1 255.255.255.0
!
```

For an L3 interface, the IPv4 protocol states in the output of show ipv4 interface brief command show as up; up on the active host and up; down on the standby host.

Active host:

GigabitEthernet100/0/0/32	1.1.1.1	Up	Up
Standby host:			
GigabitEthernet100/0/0/32	1.1.1.1	Up	Down

For an L2 interface, the ports show as up on both the hosts.

Active host:

```
GigabitEthernet100/0/0/33      unassigned      Up              Up
Standby host:

GigabitEthernet100/0/0/33      unassigned      Up              Up
```



Note You cannot add the satellite interface to the same bundle as the physical ICL link.

Satellite Management Using Private VRF

You can use a special private VRF instead of the global default routing table, to configure the loopback interface and ICLs used for satellite management traffic. IP addresses in this VRF will not conflict with any other addresses used on the router.

```
router(config)# vrf NV_MGMT_VRF
router(config)# address ipv4 unicast

router(config)# interface Loopback 1000
router(config)# vrf NV_MGMT_VRF
router(config)# ipv4 address 10.0.0.1 / 24

router(config)# interface TenGige 0/1/0/3
router(config)# vrf NV_MGMT_VRF
router(config)# ipv4 point-to-point
router(config)# ipv4 unnumbered Loopback 1000
router(config)# nv
router(config-nv)# satellite-fabric-link satellite 500
router(config-nv)# remote-ports GigabitEthernet 0/0/28-39
router(config)# nv satellite 500
router(config)# ipv4 address 10.0.0.2 / 24
```

Configuration of Satellite using Auto-IP

```
show run nv satellite 1200
nv
satellite 1200
  type asr9000v
!
interface GigabitEthernet0/1/0/5
transceiver permit pid all
nv
  satellite-fabric-link satellite 1200
  remote-ports GigabitEthernet 0/0/0-7
!
!
!

show run nv satellite 100
nv
satellite 100
type ncs5002
!
Interface HundredGigE0/1/0/0
nv
satellite-fabric-link satellite 100
remote-ports TenGigE 0/0/0-79
```

```
interface TenGigE100/0/0/4
  ipv4 address 192.100.1.1 255.255.255.0
  ipv6 address 2000:100:1::1/64
```

Satellite Configuration with Dual-Homed Hosts

You can configure satellite with dual-homed hosts as shown in this example.

```
redundancy
  iccp
    group <group-id>
      member
        neighbor <ip-address>
      !
      backbone
        interface <interface-id>
      !
      isolation recovery-delay <value>
      nv satellite
        system-mac <macaddr>
      !
    !
  !
interface TenGigE0/1/0/0
  nv
    satellite-fabric-link {network | satellite <id>}
    redundancy
      iccp-group <group-id>
    !
    remote-ports <interface-id>
  !
!
nv
  satellite <id>
  type <type>
  device-name <name>
  redundancy
    host-priority <0-255>
  !
  serial-number <serial-number>
!
```

Dual-Home for Multiple Satellites with Single Physical ICLs on Both Hosts and Satellites

Dual Home Configuration for SAT1



Note When you use either a manual IP address or an IPv4 unnumbered loopback address for the ICL, the IP address must be different on both the hosts.

Host 1:

```

interface TenGigE0/1/1/23.100
  ipv4 point-to-point
  ipv4 address 100.100.1.101 255.255.255.0
  encapsulation dot1q 100
  nv
  satellite-fabric-link satellite 100
  redundancy
  iccp-group 1
  !
  remote-ports GigabitEthernet 0/0/0-35
  !
  !
  !

```

Host 2:

```

interface TenGigE0/1/1/2.100
  ipv4 point-to-point
  ipv4 address 100.100.1.102 255.255.255.0
  encapsulation dot1q 120
  nv
  satellite-fabric-link satellite 100
  redundancy
  iccp-group 1
  !
  remote-ports GigabitEthernet 0/0/0-35

```

Dual Home Configuration for SAT2**Host 1:**

Host1:

```

interface TenGigE0/1/1/23.200
  ipv4 point-to-point
  ipv4 address 100.100.1.101 255.255.255.0
  encapsulation dot1ad 100
  nv
  satellite-fabric-link satellite 200
  redundancy
  iccp-group 1
  !
  remote-ports GigabitEthernet 0/0/0-35
  !
  !
  !

```

Host 2:

```

interface TenGigE0/1/1/2.200
  ipv4 point-to-point
  ipv4 address 100.100.1.102 255.255.255.0
  encapsulation dot1ad 120
  nv
  satellite-fabric-link satellite 200
  redundancy
  iccp-group 1
  !
  remote-ports GigabitEthernet 0/0/0-35
  !

```


Configuring nV Satellite Auto Image Upgrade: Example

The following example shows how to configure nV satellite auto image upgrade:

```
config
  nv satellite 100 upgrade on-connect
!
```

Additional References

These sections provide references to related documents.

Image Upgrade for Cisco NCS 500x Satellite from Cisco IOS XR Software Release 6.0.0

Prerequisites:

- Cisco NCS 500x satellite has 32GB internal memory and it is upgraded with the latest BIOS.
- Cisco ASR 9000 Host has a RSP3/RSP4 Route processor with at least 2 GB free memory on the disk.

Follow these steps to upgrade the image for Cisco NCS 500x Satellite:

1. Copy the new NCS 500x image (*ncs5k-mini-x-v2.iso*) into the host disk0:



Note *ncs5k-mini-x-v2.iso* is a sample name for illustration purpose. Use the actual Cisco NCS 5000 image from Cisco webpage for the updated software.

2. Once the satellite is connected and the satellite IP address can be pinged successfully, telnet into the satellite using the username/password combination of root/root.
3. On the satellite telnet console, execute **install add source tftp://<HOST ICL IP address>/../disk0 ncs5k-mini-x-v2.iso**. The **show install log** command on the satellite Telnet console would indicate the install operation progress.
4. Once the install add operation completes and the console message indicates a success, the **show install inactive** command displays the package that just got added. Now, execute **install activate <package name>** to activate the new image.
5. Finally, the device reboots automatically and comes up with the new image. Once the image boots up completely, the satellite comes back in satellite mode and gets rediscovered.

Related Documents

Related Topic	Document Title
Satellite System software upgrade and downgrade on Cisco IOS XR Software	<i>Cisco ASR 9000 Series Aggregation Services Router Getting Started Guide</i>
Cisco IOS XR interface configuration commands	<i>Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Command Reference</i>
Satellite QoS configuration information for the Cisco IOS XR software	<i>Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide</i>
Bidirectional Forwarding Detection features on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide</i>
Layer-2 and L2VPN features on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Configuration Guide</i>
Layer-3 and L3VPN features on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router MPLS Layer 3 VPN Configuration Guide</i>
Multicast features on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router Multicast Configuration Guide</i>
Broadband Network Gateway features on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router Broadband Network Gateway Configuration Guide</i>
AAA related information and configuration on the satellite system	<i>Cisco ASR 9000 Series Aggregation Services Router System Security Configuration Guide</i>
Information about user groups and task IDs	<i>Configuring AAA Services on Cisco IOS XR Software module of Cisco IOS XR System Security Configuration Guide</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
There are no applicable MIBs for this module.	To locate and download MIBs for Cisco IOS XR software, use the MIB Locator tool.

RFCs

RFCs	Title
None	N.A.

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/support



CHAPTER 2

Configuring Multicast Offload on the Satellite nV System

This chapter describes the configuration of the Satellite Network Virtualization (Satellite nV) multicast offload on the Cisco ASR 9000 Series Aggregation Services Routers.

- [Need for Multicast Offload, on page 77](#)
- [Scope and Prerequisites, on page 78](#)
- [Multicast Offload Terminology, on page 80](#)
- [Overview of Multicast Offload, on page 80](#)
- [Restrictions for Multicast Offload, on page 85](#)
- [Configuring Satellite nV Multicast Offload, on page 87](#)


Need for Multicast Offload

The Satellite nV System architecture currently requires the Cisco ASR 9000 Series Router host to process all replications for supported multicast traffic and topology profiles. This is in line with the envisioned port extender functionality of the satellite devices where all protocol processing and forwarding decisions happen on the host to fully utilize the IOS-XR functionalities. With the introduction of support for a wider variety of satellite topologies, the satellites are no longer restricted by a need for direct connection to the hosts. The satellites can be connected to a Dual Host system to form a Ring topology or through a hub and spoke topology. The satellites can also be reached over an Layer-2 fabric connection through transport provider EVCs.

This use of satellite fabric connections over EVCs or sharing of the same fabric link by a chain of satellite devices has introduced new bandwidth conservation requirements. One such specific case is the forwarding of multicast traffic over satellite ports where the Host currently does all the replication even if all the copies are destined to the ports on the same satellite device. The satellite device is unaware of any multicast group or membership and if multiple receivers are present on the ports of the same satellite, multiple copies of the same packet is sent to the satellite for each of the receivers which can eventually oversubscribe the fabric bandwidth.

With the increasing scale of satellites over the new topologies, it is evident that the model of Host side replication is not very efficient or scalable. The nV Satellite multicast offload feature is introduced to solve this problem. This feature allows the Host to forward just the pre-replicated multicast streams per multicast route (S,G) to the satellites and offload the per satellite access port replication to the satellite device itself. The protocols still run on the Cisco IOS-XR Software modules of the Host but the final replication happens locally on the satellite device based on selective download of routes to the satellites.

Figure 12: Satellite nV System Multicast Offload Illustration



Without offload, Host 1 (assuming that it is the active for both Satellite 1 and Satellite 2) sends 4 copies to receivers 1,2,3 and 4 even if they join the same multicast group, (S1,G1). With offload, Host 1 sends 1 copy and Satellite 1 replicates it twice for receivers 1 and 2 and Satellite 2 replicates it twice for receivers 3 and 4.

The multicast offload feature provides a way for a multicast route which has one or more satellite extended interfaces as its route members known as outgoing interface element lists (OLEs) on the host data plane to be represented locally on the satellite, so that the satellite receives the main copy of a multicast packet from the host and also is capable of replicating that packet into the extended interfaces, that is, the route members (OLEs) of a multicast route.

Whenever a multicast route is offloaded to the satellite, the multicast data received by the satellite on the IC link is tagged with an identifier identifying the multicast route the data belongs to, the satellite hardware reads this tag and through a table lookup finds out the extended interfaces to which the data has to be replicated for a particular route and sends the multicast data on the route-members (OLEs) of that multicast route. The multicast feature channel allows for the host to specify to the satellite which extended interfaces are mapped to which multicast route (tag or multicast isid).

Scope and Prerequisites

These are the supported requirements and prerequisites:

- **Protocol:** IGMP Snooping protocol.



Note nV Satellite does not support offload for Layer 3 multicast, IPv6 multicast traffic or MLD snooping protocol.

- **Satellite topologies:** Hub and spoke topology with a single nV host.



-
- Note** nV satellite supports only these topology variants:
- L2 access and L2 core with Satellite Host providing only Layer 2 bridging functionality.
 - L2 access and L3 core connected over BVI on the Satellite Host.
 - ICL bundles and access bundles on satellites connected in a Hub and Spoke topology.
-

• **Transport types:**

- MVPN-MLDP Inband
- MVPN-GRE
- MVPN-P2MP-TE
- L2 bridging, VPLS
- VPLS LSM
- Native multicast under BVI for PIM, MLDP, and P2MP-TE
- MVPN IRB with PIM, MLDP, and P2MP-TE core tree type



-
- Note** The nV Satellite interfaces are recommended to be deployed as access or edge interfaces and hence do not support functionalities of core interfaces on multicast topologies.
-

Prerequisites:

All existing hardware and software prerequisites for multicast support on the Host are applicable. Similarly, all existing hardware and software prerequisites for the dual host topology is applicable.

In addition, the offload feature is not backward image or device compatible. Hence, before enabling offload, all devices in the system need to be upgraded to the following software:

- Cisco IOS XR Software Release 5.2.2 for offload on Cisco ASR9000v.
- Cisco IOS XR Software Release 6.1.2 for offload on Cisco ASR9000v and Cisco ASR9000v-V2 satellites in Hub and Spoke topology.
- Cisco IOS XR Software Release 6.2.1 for offload on Cisco NCS5000 series routers as satellites in Hub and Spoke topology.

For better convergence and redundancy for connectivity failures to the Satellite Hosts or the core, the IGMP snoop sync feature needs to be mandatorily turned on for the redundancy group of the Dual Host System. Otherwise, there can be traffic drops when the Designated multicast forwarder picked by satellite does not align to the Unicast Active Host.



Note For information on bidirectional sync on dual-homed satellites using IGMP Snooping, see the *Bidirectional Internet Group Management Protocol Snoop Synchronization for Satellite Dual-Homed System* section in *Multicast Configuration Guide for Cisco ASR 9000 Series Routers*

Multicast Offload Terminology

These are the different terms associated with the Multicast offload solution.

- **Unicast ISID:** It is the ISID value set in 802.1ah header (802.1ad vlan headers for Cisco NCS 5000 series satellites) for unicast data packet to satellite and for all data packets from satellite to represent satellite access port. Unicast ISID is used to identify the Slot/Subslot/Port of a satellite for unicast packets. It is scoped over a satellite.
- **Multicast ISID:** It is the ISID value set in 802.1ah header (802.1ad vlan headers for Cisco NCS 5000 series satellites) for multicast data packet from host to satellite to represent offloaded Multicast Route (S, G) per Ring.
- **Designated Multicast Forwarder:** Each satellite selects a host as the designated multicast forwarder and replicates multicast packets for all receivers (satellite access ports) for route (S, G) from the designated multicast forwarder.
- **Backup Multicast Forwarder:** Multicast packets received from Backup Multicast forwarder are not replicated on the satellite for access ports.
- **Active Unicast Host:** The unicast customer data traffic is switched through Active Unicast Host while traversing through the satellite.
- **Standby Unicast Host:** In case of a lost connection to Active Unicast Host due to failures such as cut cables or connection interface failure, the Standby unicast host shall become the new Active Unicast Host.
- **Primary Host:** Specifies the Host with the lowest chassis MAC in a Dual head topology.
- **Secondary Host:** Specifies the Host with a higher chassis MAC than the Primary Host in Dual head topology.
- **Primary OLE:** Specifies an OLE selected amongst the active OLEs of an offloaded route to represent the receiver of the single instance/ replication of offloaded multicast traffic sent from the host to the satellite. For ring topology, one primary OLE is selected per offloaded route (S,G), for each BD and each ring. For Hub and Spoke topology, the primary OLE is scoped for each offloaded route, BD and satellite.

Overview of Multicast Offload

As in the case of existing Satellite nV System architecture, the IGMP snooping protocol runs on the satellite hosts with full Cisco IOS XR Software feature parity even with offload. The joins received on the satellite access ports reach the active host and gets processed by the IGMP snoop module.

- The IGMP snoop sync feature running over ICCP protocol synchronizes the joined multicast group membership information to the other host. This keeps the IGMP protocol state for receivers joining over satellite access ports in sync across both hosts in the case of dual host topology.
- In the case of dual host topology, the hosts independently offload eligible routes to the relevant satellite devices through SDAC. This includes the list of local ports that have expressed interest in this route.
- The joins also get propagated to the core by each of the hosts (dual host topology), if they have an active link to the core. This allows better convergence in case of a redundancy switchover on the satellites.
- When traffic from the source is received on the host, a special indication in the satellite 802.1ah encapsulation header (802.1ad vlan headers for Cisco NCS 5000 series satellites) signals the satellite devices about the offloaded route to which this packet belongs to. The host then transmits the packet out of a dynamically elected primary OLE for each route (S,G) on a specific bridge domain. Then, the satellite devices locally replicate the traffic to all the intended local receivers on that device.
- As the same traffic stream comes in from both the hosts (in the case of dual host topology), the satellites pick a designated multicast forwarder at discovery and continue to replicate from that host until that host goes away or loses its core connectivity. Thus the offload topology is a full active / active system.

Figure 13: IGMP Join Programming (Access L2 and Core L2/BVI)


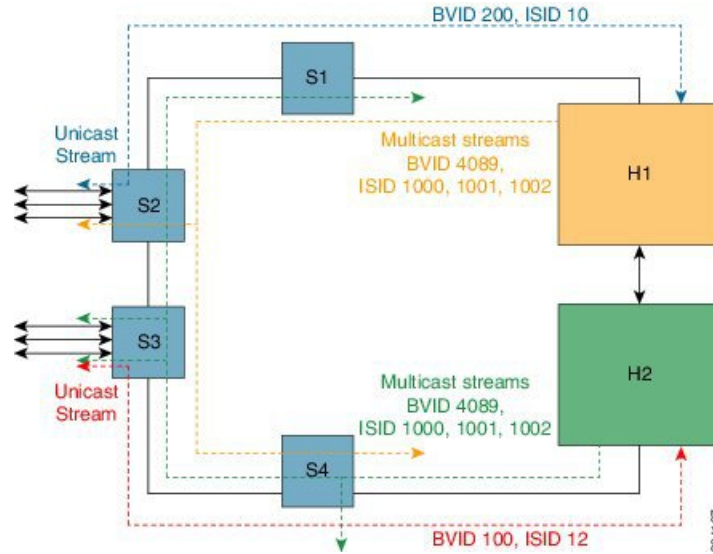


Figure 14: Multicast Offload Data Flow



These illustrations show the multicast offload functionality on Satellite nV System that use hub and spoke topologies.

Figure 15: Multicast Offload Data Flow with ICL bundles to different satellites in Hub and Spoke topology





Figure 16: Multicast Offload with ICL bundle from different line cards to a satellite



Behavior of Multicast Offload During Satellite Topology Events

These are the behavior of the Multicast offload during various satellite topology events:



Note These are applicable in the case of dual host ring topology setup.

- **Split brain:** During a split brain event, the hosts function independently and the system behaves as two separate chains with IGMP snoop sync broken between satellites and each host offloading to its own set of satellites. On recovery, there is reconciliation and some reprogramming that can result in traffic loss as IGMP snoop protocol states sync up.
- **RSP failover:** This is expected to be hitless for traffic going over routes that have been offloaded already. New joins during the failover may experience the same set of events and traffic impact as the split brain condition.
- **Ring break and Recovery:** On ring breakage, any satellites aligned to other host as designated multicast forwarder has to realign to the remaining host on the ring. This can have a minimal (sub-second for L2 core) traffic impact on a stable system where the routes are already synchronized. On recovery, there is no revertive switchover and so there must be no traffic impact.



Note For more information, refer *Features of Dual Home Network Architecture* section in the *Configuring the Satellite Network Virtualization (nV) System* chapter.

Difference between Layer 2 and Layer 3 (over BVI) core support

For the Dual host topology systems connecting to an Layer 3 core through BVI, the current recommendation is to align all satellites to a single host as Unicast active. This is required as BVI protocol states are kept down on the standby host to avoid ECMP being triggered from the core for an active/standby Satellite nV system that could lead to traffic drop of half the packets on the standby host.

Currently, multicast offload which is implemented as an inherently active/active system needs to track the BVI redundancy states as well to avoid picking the Designated multicast forwarder on the host having the standby BVI. While this ensures a common feature model with the Layer 2 core, the benefits of an active/active system, especially, the sub-second switchover convergence is not valid anymore. This can be improved by at least allowing IGMP snoop sync to continue by configuring an internal querier with system IP address lower than the BVI IP address and a query max timeout of 1s.

Uplink/Core Isolation Tracker

Under redundancy group configuration, there is an option to configure a backbone interface. Multiple backbone interfaces can be configured depending on the number of links from the Host to the core. ICCP protocol keeps a watch over the link states of these interfaces and if all of these backbone interfaces go down, then a core isolation event is notified to the client.

As part of the multicast offload feature, IGMP Snoop registers as a client to ICCP protocol and propagates these notifications to the satellite devices. Any satellite which still expect this host to be the designated multicast forwarder can then switchover to the other host as core connectivity is lost. In order to avoid traffic downtimes during flaps, this event is triggered only for core link up to down condition and is non-revertive. A satellite device stays with a host as a designated multicast forwarder until it goes down and does not switch back if the previous host that triggered a switchover comes up.

For more information, see *Configuring Interchassis Communication Protocol* section in the *Interface and Hardware Component Configuration Guide for Cisco ASR 9000 Series Routers*

Multicast Offload Scalability

These are the maximum scale capabilities of Multicast Offload Solution:

- The maximum number of rings over which the same offloaded route can be downloaded is 32.
- The maximum number of offloaded routes for each satellite or ring is 2000. The maximum number of multicast routes for each satellite in the hub and spoke topology is 1000.
- The maximum number of offloaded OLEs for each satellite is 7800 for Cisco ASR9000v/v2 satellites and up to 3500 on Cisco NCS 5000 satellites.
- The maximum number of OLEs for each route is 44 on Cisco ASR9000v/v2 satellites, 40 on Cisco NCS 5001 satellite and 80 on Cisco NCS 5002 satellite.
- The maximum number of offloaded routes for each system is equal to the maximum number of rings multiplied by 2000 up to the Host limit (128k).
- The maximum bandwidth of ICL to ICL traffic with offload enabled for each satellite is 6Gbps in each direction.

Multicast Offload Use Case Scenario

The Satellite nV multicast offload feature has been designed and optimized specifically for an IPTV scenario where the Cisco ASR 9000 Series Satellite nV System replaces the traditional edge or aggregate router and switch solution. The Satellites feed multicast traffic from the Layer 2 or Layer 3 (through BVI) core to the Layer 2 access trunk links connecting DSLAMs which in turn terminate residential gateways and take care of subscriber aware processing, if any, as shown below.

Residential Gateway ---- DSLAM ---- Cisco ASR 9000 Satellite nV System---- Core ---- IPTV source

Therefore, the Satellite nV multicast offload solution does not need to support VLAN rewrite operations on individual offloaded replications as they all go over the same trunk video VLAN. This reduces the offload processing overhead on the satellite devices to achieve line rate replication. Similarly, the solution is optimized to an use case that requires minimal egress feature processing. QoS, ACL and other egress features act on the pre-replicated multicast stream and the configuration needs to be replicated to all participating OLEs in the offload, if at all feature processing is specifically required.

Network design needs to ensure that there is no congestion on the satellites post replication. This is critical and different from a non-offload solution as the Host takes care of correct priority aware traffic shaping through Auto QoS or a user specified MQC policy for the non-offload case. For offloaded traffic, as the replication happens locally on the satellite devices, the Host QoS is unaware of the total traffic volume post replication, and therefore cannot include it in its bandwidth computations without statically reserving bandwidth on all offload participating OLEs.

Such a permanent reservation might be sub-optimal in most cases and a rigid reservation may not cater to user needs. However, a solution based on intelligent network design is generally straightforward for the IPTV roll outs, as they have well defined bandwidth planning, given the load generally remains constant over time. For residential triple and quadruple play cases, with possibility of other priority and internet traffic causing oversubscriptions, a simple QoS port-shaper policy on the egress video traffic DSCP/CoS markers or VLAN can be used to classify and allocate a reserved bandwidth for multicast offloaded video streams by shaping the remaining traffic to 60-70% of the access port interface bandwidth as required.

Due to various scale bottlenecks and for practical reasons, the need for a simple service QoS model as above the Satellite nV multicast offload solution is recommended to be deployed on a -SE version of the ethernet line cards.

Restrictions for Multicast Offload

These are the restrictions and considerations of the Multicast offload feature:

- Only Layer-2 multicast offload through IGMP snooping for IPv4 multicast traffic is supported. IPv6 multicast traffic / MLD snoop and Layer-3 multicast traffic replication offload is not supported.
- Offloaded OLE interface statistics (except for the primary OLE) are only incremented on the main interface, even if the OLE is actually on a sub-interface. L2VPN statistics reflect updates against the correct OLE and multicast route.
- Endian type mismatching RSPs across the two hosts of the Dual Host system is not supported. IGMP snoop sync, a pre-requisite for this feature does not function on such hardware.
- Multicast offload does not support broadcast, multicast router port, or TCN flood event offloading. All such traffic is replicated on the host.

- Only normal Satellite EFPs are considered as OLEs eligible for offload. An OLE consisting of a pseudo wires and BNG subscribers over satellite ports are not supported as eligible OLEs for offload. These might fall back to legacy host side replication.
- VLAN rewrite support for offloaded replications is not supported. All offloaded replications carry the same VLAN tag as the original copy sent by the Host over the dynamically picked primary OLE. The same VLAN rewrite configuration, if required, needs to be added on all the EFPs of a route as any OLE can be dynamically selected as the primary OLE on host. The same applies for any other egress feature configuration on the OLEs.
- The satellite hardware can only replicate once per physical access port per multicast route (S,G). Therefore, multiple sub-interfaces connected to receivers for same route in the same bridge domain cannot be offloaded. This is still acceptable for an IPTV scenario, where the whole port is a trunk to DSLAM devices downstream.
- Split horizon does not work with offloaded satellite OLEs on a bridge domain. So, multicast source or multicast routers (mrouter port) must not be over satellite EFPs within the same bridge domain where receivers through offloaded satellite EFPs are also present.
- Offload functionality is disabled by default. If it is enabled, there will be a transient downtime including impact to traffic for all IGMP snoop routes on the system as the offload mode switches. Similarly, all routes (including non-offloaded routes) in the system are impacted if offload is disabled on any of the bridge domains as IGMP snoop moves to non-offload mode.



Note Local replication can co-exist with offloaded replication for same route/bridge domain/satellite combinations.

- While sub-second convergence and fast reroutes are possible for an Layer-2 core failover scenarios (ring break, split brain, core link down and so on), L3 core (over BVI) requires the BVI redundancy state toggles post failover before protocol states can be synchronized between Hosts. This cold failover mode for L3 core over BVI is slower and in the order of seconds compared to the Layer-2 core case.
- VLAN ranges cannot be used for IGMP ports involved in the sync between the Dual Hosts of a satellite system. Unambiguous VLAN id configuration is required.
- Cisco ASR 9000v and Cisco ASR9000v-V2 only support 12 Gbps of ICL to ICL bidirectional multicast replication and a maximum of 6Gbps for each direction while processing multicast offloaded traffic even if it does not actually participate in local replications. With respect to Hub and Spoke topology, support for multicast offload traffic processing on the satellite is restricted to 12 Gbps or the ICL bandwidth, whichever is lower, in the case of single host topology and 6 Gbps for each host in the case of dual host topology.
- The offload solution has been optimized and specifically characterized for the scale numbers mentioned. There might be a fallback to legacy Host side replication if any of those numbers are exceeded or if any of the limitations stated above make an OLE ineligible for offload. However, these fallbacks may be delayed or may not work gracefully, especially in scale exceed cases. Hence, network design must ensure that the scale numbers are complied for optimal performance.
- On reload of one of the hosts in the dual host system, the hosts might go out of sync because offload failures are not synchronized between hosts. A restart of the IGMP snoop process or a receiver leave/join can resolve this. The impact is only for scale exceed cases, which is not recommended in general.

- In the case of Layer 2 Multicast offload, the satellite does not support any egress features on the replicated packet on satellite. Egress features such as ACL, QoS shaping/ policing, VLAN rewrite, SPAN, netflow will work on the host on the primary OLE copy sent to the satellite. If any egress feature need to applied to offloaded multicast streams, it needs to be configured on all the EFP as any OLE can be selected as primary OLE on host.
- ICLs (physical interfaces or bundle interfaces) must be main interfaces. This means that sub-interfaces on ICLs are not supported.
- You cannot configure satellite rings and satellite hub and spoke topologies on the same host if any of the satellites are participating in Layer 2 multicast offload. The satellites can be configured in either ring or hub and spoke topologies on a per-router basis.
- Restrictions for Offload on Cisco NCS5000 Series satellites:
 - The customer vlan ids from 3967 to 4095 are reserved vlan ids for the satellite ports participating in the offload.
 - There is no support for partitioned ICLs for the satellites participating in the offload.
 - There is no support for L2 multicast offload in a dual host case for hub and spoke topology.
 - Only MC-LAG cold standby mode is supported on access ports participating in L2 multicast offload.
 - Access bundles are supported as offloaded OLEs only in Hub and Spoke topology. Further, there is no support for access bundles across satellites in any dual host nV satellite topology.

Configuring Satellite nV Multicast Offload

Prerequisites

All the existing configuration for implementing multicast and Satellite nV System have to be configured and the multicast offload is only an incremental function that allows for the final routes to be offloaded to the satellite instead of the local Cisco ASR 9000 Series Router line card hardware.

By default, multicast offload is disabled on the host.



Note Refer the *Multicast Configuration Guide for Cisco ASR 9000 Series Routers* for information on Multicast routing and IGMP snooping. All the configuration has to be manually synchronized between the two hosts for the functionality to perform correctly.

Enabling Multicast Offload on a Bridge Domain

To enable multicast offload use the **nv satellite offload ipv4 multicast enable** command:

```
RP/0/0/CPU0: (config-l2vpn-bg-bd-ac)#show run l2vpn
l2vpn
 bridge group bg1
  bridge-domain bd1
  nv
  offload ipv4 multicast enable
```

By default, multicast offloading is disabled on all bridge domains. To enable IPv4 Multicast offloading (including IGMP), run this command:

```
(config)#l2vpn
(config-l2vpn)#bridge group <bg>
(config-l2vpn-bg)#bridge-domain <bd>
(config-l2vpn-bg-bd)#nv
(config-l2vpn-bg-bd-nv)#offload ipv4 multicast enable
```



Note This configuration needs to be enabled on both hosts participating in the Dual Host system of the Ring topology.

Configuration examples for Satellite nV multicast offload

These show command outputs provide the necessary details about multicast offload:

```
RP/0/RSP0/CPU0:router#show igmp snooping group summary debug
```

```
Bridge Domain bg:bd
```

#Mem	Master	OLE		#Mem	#Inc	#Exc	Annot		Master				
Group	Annot	Source	InterfaceName	Ver	GM	Ports	Ports	Ports	Key	ISID	Ver	OLE	XID
225.0.0.1		{}		V3	EX	2	-	-	0x2-c8	0x3e8	0x4	0xa0000005	
2		Bundle-Ether50		V3	EX	2	-	-	0x2-64	0x3e8	0x4	0xa0000005	
225.0.0.2		{}							0x2-c8	0x3ed	0x6	0xa0000004	
1		Bundle-Ether40											
225.0.0.3		{}		V3	EX	2	-	-	0x2-c8	0x3e9	0x4	0xa0000005	
2		Bundle-Ether50											
225.0.0.4		{}		V3	EX	2	-	-	0x2-c8	0x3ea	0x4	0xa0000005	
2		Bundle-Ether50											
225.0.0.5		{}		V3	EX	2	-	-	0x2-64	0x3e9	0x4	0xa0000005	
2		Bundle-Ether50											
225.0.0.6		{}		V3	EX	2	-	-	0x2-c8	0x3eb	0x4	0xa0000005	
2		Bundle-Ether50											
225.0.0.7		{}		V3	EX	2	-	-	0x2-64	0x3ea	0x4	0xa0000005	
1		Bundle-Ether50							0x2-c8	0x3ee	0x6	0xa0000004	
1		Bundle-Ether40											
225.0.0.8		{}		V3	EX	2	-	-	0x2-64	0x3eb	0x4	0xa0000005	
1		Bundle-Ether50							0x2-c8	0x3ef	0x5	0xa0000004	
1		Bundle-Ether40											
225.0.0.9		{}		V3	EX	2	-	-	0x2-c8	0x3ec	0x4	0xa0000005	
2		Bundle-Ether50											
225.0.0.10		{}		V3	EX	2	-	-	0x2-64	0x3ec	0x4	0xa0000005	
2		Bundle-Ether50											

```
RP/0/RSP0/CPU0:router#show igmp snooping port group debug
```

```
Key: GM=Group Filter Mode, PM=Port Filter Mode
```

```
Flags Key: S=Static, D=Dynamic, E=Explicit Tracking, R=Replicated, I=Inherited OLE
```

```
OLE Offl Key: F=False, T=True, d=Offload disabled, H=Offload disabled, not in SHG-0,
```

```
M=Offload disabled, mrouter port, t=Offload disabled, topo not supported
```


Bridge Domain bg:bd

Port	InterfaceName	PM Group	Ver	GM	Source	Exp	Flgs	OLE Offl
Ul_ifh								
----	-----	--	----	---	--	-----	---	-----
BE40		- 225.0.0.1	V2	-	*	92	D	T
0x000036e0	GigabitEthernet200_0_0_4							
BE40		- 225.0.0.2	V2	-	*	86	D	T
0x00003720	GigabitEthernet200_0_0_3							
BE40		- 225.0.0.3	V2	-	*	87	D	T
0x00003760	GigabitEthernet200_0_0_2							
BE40		- 225.0.0.4	V2	-	*	85	D	T
0x000036e0	GigabitEthernet200_0_0_4							
BE40		- 225.0.0.5	V2	-	*	90	D	T
0x000037e0	GigabitEthernet100_0_0_2							
BE40		- 225.0.0.6	V2	-	*	87	D	T
0x00003760	GigabitEthernet200_0_0_2							
BE40		- 225.0.0.7	V2	-	*	88	D	T
0x00003720	GigabitEthernet200_0_0_3							
BE40		- 225.0.0.8	V2	-	*	89	D	T
0x00003720	GigabitEthernet200_0_0_3							
BE40		- 225.0.0.9	V2	-	*	86	D	T
0x00003760	GigabitEthernet200_0_0_2							
BE40		- 225.0.0.10	V2	-	*	88	D	T
0x000037e0	GigabitEthernet100_0_0_2							
BE50		- 225.0.0.1	V2	-	*	85	D	T
0x000037a0	GigabitEthernet200_0_0_1							
BE50		- 225.0.0.2	V2	-	*	91	D	T
0x000036a0	GigabitEthernet100_0_0_1							
BE50		- 225.0.0.3	V2	-	*	90	D	T
0x000037a0	GigabitEthernet200_0_0_1							
BE50		- 225.0.0.4	V2	-	*	85	D	T
0x000037a0	GigabitEthernet200_0_0_1							
BE50		- 225.0.0.5	V2	-	*	84	D	T
0x000036a0	GigabitEthernet100_0_0_1							
BE50		- 225.0.0.6	V2	-	*	85	D	T
0x000037a0	GigabitEthernet200_0_0_1							
BE50		- 225.0.0.7	V2	-	*	90	D	T
0x000036a0	GigabitEthernet100_0_0_1							
BE50		- 225.0.0.8	V2	-	*	90	D	T
0x000036a0	GigabitEthernet100_0_0_1							
BE50		- 225.0.0.9	V2	-	*	89	D	T
0x000037a0	GigabitEthernet200_0_0_1							
BE50		- 225.0.0.10	V2	-	*	85	D	T
0x000036a0	GigabitEthernet100_0_0_1							

RP/0/RSP0/CPU0:router# show nv satellite multicast satellite 210

```

-----
Satellite 210
-----
Multicast channel state: Open
Host state: Designated Forwarder (DF)
RP/0/RSP0/CPU0:SAT-SCALE-HOST1#
    
```



Note For more information on IGMP Snooping commands, see *Multicast Command Reference for Cisco ASR 9000 Series Routers*. For more information on L2VPN forwarding commands, see *VPN and Ethernet Services Command Reference for Cisco ASR 9000 Series Routers*



INDEX

A

active/standby [9](#)
ASR9001 [20](#)
auto-ip [33](#)

B

BNG [20](#)
 interoperability [20](#)

C

CGN [20](#)
chassis management [21](#)

H

host ip address, configuring [34](#)

I

inter-chassis link [13](#)

N

nV [20](#)

P

port extender model [6](#)
private vrf [70](#)

Q

quality of service [21](#)

S

Satellite Discovery And Control (SDAC) Protocol [14](#)
satellite ethernet [20](#)
 l2vpn [20](#)
 layer 2 information [20](#)
satellite global configuration [68](#)
satellite interface configuration [69](#)
satellite nV system [4, 24, 31](#)
 benefits [4](#)
 defining [31](#)
 restrictions [24](#)
sdac ip connectivity [14](#)
split brain handling [9](#)
sync-e [43–44, 46](#)
 configuration example [46](#)
 restrictions [44](#)

T

time of day [21](#)

