Cisco ASR 9000 Series Aggregation Services Router nV System Configuration Guide, Release 6.1.x

First Published: 2016-11-01

Americas Headquarters

CONTENTS

Preface

The Preface contains the following sections:

Changes to This Document

This table lists the technical changes made to this document since it was first released.

Table 1: Changes to This Document

Date	Change Summary
November 2016	Initial release of this document.

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.

New and Changed nV System Features

This chapter lists all the features that have been added or modified in this guide. The table also contains references to these feature documentation sections.

• New and Changed Satellite nV Feature Information, on page 1

New and Changed Satellite nV Feature Information

Table 2: New and Changed Features

Feature	Description	Changed in Release	Where Documented
Native Cisco NCS 5000 SMU push capability through the Cisco ASR 9000 Series Router host.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter • Native Image and SMU Push Support for Cisco NCS 5000 Series Satellites, on page 36
Capability to include upto 10 ICLs that can encompass 10 members in a bundle for each Cisco ASR 9000 Series Router host.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter • Support of 10x10G ICLs on Cisco NCS 5000 Series Satellites, on page 36
Support for Dual Home network topology on Cisco NCS 500x Series satellite.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter

Feature	Description	Changed in Release	Where Documented
Support for fabric CFM over bundle ICLs in Layer 1 Hub and Spoke topology for Cisco ASR 9000v satellites.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter
Fabric CFM for both physical and bundle ICLs for Cisco NCS 500x Series satellite.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter
QoS offload on Cisco NCS 500x Series satellite.	This feature was introduced.	Release 6.1.2	*Configuring the Satellite Network Virtualization System* chapter **Note** For detailed information on QoS offload, see *Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide.*
Layer 2 multicast offload for Hub and Spoke network topology in Satellite nV System with Cisco ASR 9000v satellites	This feature was introduced.	Release 6.1.2	Overview of Multicast Offload, on page 94

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 3: Feature History for Configuring Satellite System

Release	Modification
Release 4.2.1	• Support for Satellite Network Virtualization (Satellite nV) Service was included on the Cisco ASR 9000 Series Router for Cisco ASR 9000v Satellite .
Release 4.2.3	• Support for 36-Port 10-Gigabit Ethernet Line Card was included.
Release 4.3.0	• Support for Cisco ASR 9001 and Cisco ASR 9922 Series Routers as hosts was included. • Support for Cisco ASR 901, and Cisco ASR 903 as Satellite devices was included. • Support for Cisco ASR 901-1G series router. • Username and password (limited AAA) support for satellites.
Release 4.3.1	• Support for Auto-IP feature was included. • Support for Link Layer Discovery Protocol (LLDP) over Satellite access interface over bundle ICL was included. • Caveat for multiple members from the same NP belonging to same ICL bundle was removed . • Procedure to convert a Cisco ASR 901 or Cisco ASR 903 Router to a satellite was added.

Release 5.1.1	These features are included on Cisco ASR 9000v and Cisco ASR 901 satellites:
	• Support for Simple Ring Satellite nV topology was included.
	• Support for dual-homed Satellite nV network architecture was included.
	• Support for Layer 2 Fabric network architecture was included.
	• Support for Fabric Ethernet Connectivity Fault Management (Ethernet CFM) was included.
	• Support for 1G ICL on ports 1/45 and 1/46 on Cisco ASR 9000v satellite was included for all the Satellite nV topologies by using 1G SFPs.
	• Support for these new satellites was included:
	• Cisco ASR 901 Series Aggregation Services Router Chassis, Ethernet-only interfaces, 10 GE, DC power, USB **(A901-6CZ-F-D)**
	• Cisco ASR 901 Series Aggregation Services Router Chassis, Ethernet and TDM interfaces, 10 GE, DC power, USB **(A901-6CZ-FT-D)**
	• Cisco ASR 901 Series Aggregation Services Router Chassis, Ethernet-only interfaces, 10 GE, AC power, USB **(A901-6CZ-F-A)**
	• Cisco ASR 901 Series Aggregation Services Router Chassis, Ethernet and TDM interfaces, 10 GE, AC power, USB **(A901-6CZ-FT-A)**
	• Support for QoS Offload over Satellite was supported. See *Cisco ASR 9000 Series Aggregation Services Router QoS Configuration Guide* for more details.
Release 5.2.0	• Support for SyncE offload on the Cisco ASR 9000v satellite with the host running Cisco IOS XR 32-bit operating system.
	• A9K-MPA- 4X10GE
Release 5.3.0	• Support for Cisco ASR 903 Router as satellite is removed.

Release 5.3.1	• Support for Unidirectional Link Detection (UDLD) protocol over Satellite access interface over bundle ICL was included.
	• The Performance Monitoring (PM) for Connectivity Fault Management (CFM) on the nV Satellite system has been improved by processing the CFM PM frames on the satellite instead of the ASR9K host.
	• Support for ICL Fabric Port Monitoring was included.
	• Support for Soft Minimum Active Links for Dual Home Topology was included.
Release 5.3.2	• Support for Multiple ICCP groups for Dual Head topology was included.
	• Support for Dynamic ICL for Cisco ASR 9000v satellite was included.
	• Support for access bundle with fabric redundancy feature was included.
	• Support for 802.3ah Loopback on bundled and non-bundled ICLs was included.
	• Support for ARP redundancy in Dual Head topology was included.
	• Satellite nV Usability Enhancements were introduced.
Release 5.3.3	Support for A9K-MOD400 line card was added with following Modular Port Adaptors (MPAs): • A9K-MPA-2X10GE • A9K-MPA-4X10GE • A9K-MPA-8X10GE • A9K-MPA-20X10GE Support for Cisco ASR 901 Router as satellite type is removed.
Release 5.3.4	Support tunable DWDM SFP + configuration on satellite console.

Release 6.0.1	• Support for Multiple ICL for Layer 2 Fabric Network Topology was included.
	• Support for DHCP on Satellite nV System Bundle Over Bundle was included.
	• Support for L1 Physical ICL while configuring CFM on a Satellite nV Fabric link interface.
	• Support for Cisco NCS 5000 Series satellite to Cisco ASR 9000 Series Host was included. This includes the support for:
	• 10/100 GigE ICL ports and 1/10GigE access ports.
	• Inband image upgrade through the nV pie .
	• Improved manageability options such as nV fabric port management, satellite login/password configuration, and satellite host name remote configuration.
	• Dynamic ICL through configurable fabric port feature.
	• Bundle ICL support including Bundle over Bundle.
	• A9K-MOD400G-SE/TR with MPA-20x10G as satellite hosting line card on Cisco ASR 9000 series host.
Release 6.0.2	Support for Delayed Switchback for nV Dual Hosts was added.
Release 6.1.2	Support for these features is included on the Cisco NCS 5000 Series Satellite:
	• Native Cisco NCS 5000 Series Satellite SMU push capability through the Cisco ASR 9000 Series Router host.
	• Capability to include upto 10 ICLs that can encompass 10 members in a bundle for each Cisco ASR 9000 Series Router host.
	• Fabric CFM for both physical and bundle ICLs.
	• Support for Dual Home network topology.
	• QoS offload on Cisco NCS 5000 Series Satellite.
	Support for fabric CFM over bundle ICLs in Layer 1 Hub and Spoke topology for Cisco ASR 9000v satellites.
	Layer 2 multicast offload for Hub and Spoke network topology with Cisco ASR 9000v satellites.

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 with Cisco ASR 9000 Enhanced Ethernet line cards or the Cisco ASR 9000 High Density 100GE Ethernet 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-V1, 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)

> **Note**
>
> RSP2 is not qualified with Cisco NCS 5000 Series satellites. Also, the RSP needs to have at least 2 GB of free space in the memory for nV pies/ SMUs. Hence, RSP-440 TR cards are not recommended with Cisco NCS 5000 Series satellites. For more information, see Hardware Installation Guide for Cisco NCS 5000 Series Routers.

- Software — Cisco IOS XR Software Release 4.2.1 or later. To use Cisco ASR9000v-V2 as satellite, Cisco IOS XR Software Release 5.2.2 or later must be installed on the host. To use Cisco NCS 5002 Series Router as satellite, Cisco IOS XR Software Release 6.0.x or later must be installed on the host. To use Cisco NCS 5001 Series Router as a satellite, Cisco IOS XR Software Release 6.0.1 or later must be installed on the host.

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 9001, Cisco ASR 9904, Cisco ASR 9006, Cisco ASR 9010, Cisco ASR 9906, Cisco ASR 9910, Cisco ASR 9912 and 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. When the Satellite nV service was introduced in Cisco IOS XR Release 4.2.x, Cisco ASR 9000v was used as the satellite device. It supports one Gigabit Interchassis Links (ICL) in two of the ports (1/45 and 1/46).

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 and 100-Gigabit Ethernet interfaces can be used as ICL ports. The 1GE and 10GE ICL ports are supported on the Cisco ASR 9000 Enhanced Ethernet line card and the100 GE ICL ports on the Cisco ASR 9000 High Density 100GE Ethernet line cards..

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.

Figure 1: Cisco ASR 9000 Series Router nV Switching System
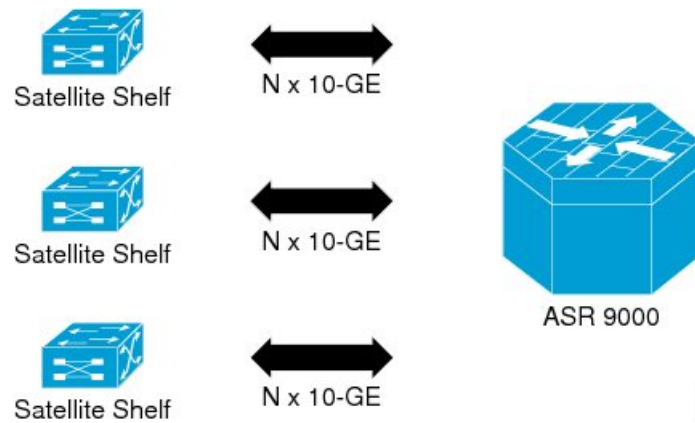
Note   Though the above figure shows Nx10GigE links, Satellite nV System also supports Nx1GigE and Nx100GigE links.

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   Similar port scalability can be achieved for 10GigE access ports and 100GigE uplinks using the Cisco NCS 5000 Series satellites from Cisco IOS XR Software Release 6.0.x.

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 Satellite Network Virtualization (nV) hardware compatibility matrix for the Cisco ASR 9000 Series Routers.

Table 4: Cisco ASR 9000 Series Router Satellite nV Hardware Compatibility Matrix

Line Cards	Cisco ASR 9000v Satellite Supported Version	Cisco NCS 5000 Series Satellite Supported Version
A9K-MPA-20X1GE on MOD80 and MOD160	4.2.1	–
A9K-MPA-2X10GE on MOD80 and MOD160	4.2.1	6.0.1
A9K-MPA-4X10GE on MOD80 and MOD160	4.2.1	6.0.1
A9K-MPA-8X10GE on MOD160	4.3.1	6.0.1
A9K-24X10GE-SE	4.2.1	6.0.1
A9K-24X10GE-TR	4.2.1	6.0.1
A9K-36X10GE-SE	4.2.3	6.0.1

Line Cards	Cisco ASR 9000v Satellite Supported Version	Cisco NCS 5000 Series Satellite Supported Version
A9K-36X10GE-TR	4.2.3	6.0.1
A9K-40GE-TR	5.2.2	–
A9K-40GE-SE	5.2.2	–
A9K-4X100G-SE	5.3.1	6.0.1
A9K-4X100G-TR	5.3.1	6.0.1
A9K-8X100G-SE	5.3.1	6.0.1
A9K-8X100G-TR	5.3.1	6.0.1
A9K-4T16GE-TR	5.3.2	6.0.1
A9K-4T16GE-SE	5.3.2	6.0.1
A9K-MPA-2X10GE on MOD200	6.0.1	6.0.1
A9K-MPA-2X10GE on MOD400	5.3.3	6.0.1
A9K-MPA-4X10GE on MOD200	6.0.1	6.0.1
A9K-MPA-4X10GE on MOD400	5.3.3	6.0.1
A9K-MPA-8X10GE on MOD200	6.0.1	6.0.1
A9K-MPA-8X10GE on MOD400	5.3.3	6.0.1
A9K-MPA-20x10GE on MOD200	6.0.1	6.0.1
A9K-MPA-20x10GE on MOD400	5.3.3	6.0.1
A9K-48X10GE-1G-SE	6.2.2	6.2.2
A9K-48X10GE-1G-TR	6.2.2	6.2.2
A9K-24X10GE-1G-SE	6.2.2	6.2.2

IOS XR 64 Bit Satellite nV Hardware Compatibility Matrix

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

Table 5: IOS XR 64 Bit Satellite nV Hardware Compatibility Matrix

Line 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

Line Cards	Cisco ASR 9000v Satellite Supported Version	Cisco NCS 5000 Series Satellite Supported Version
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
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

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.

Note In releases later than Cisco IOS XR Software Release 4.2.1, attachment models beyond the port extender model are also supported.

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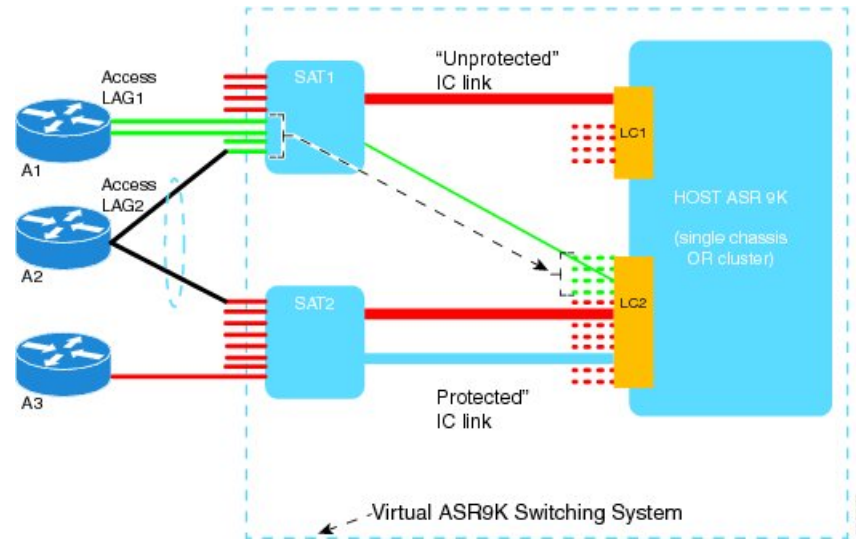
In this model, a single Cisco ASR 9000 Series Router hosts two satellite switches, SAT1 and SAT2, to form an overall virtual Cisco ASR 9000 switching system; represented by the dotted line surrounding the Cisco ASR 9000 Series Router, SAT1, and SAT2 in the Figure.

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 4 different network topologies for the ICLs between satellite switches and the host:

- Hub and Spoke network topology

- Dual Home network topology

- Simple Ring topology

- Layer 2 Fabric network topology

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 12. Dual Home, Simple Ring and Layer 2 Fabric are the Advanced Satellite nV network topologies and are covered in the section Advanced Satellite nV System Network Topologies, on page 14.

All these topologies 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).

These topologies also allow a satellite switch to be geographically at a separate 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*.

This table summarizes the network encapsulation techniques used by different Satellite nV System topologies.

Table 6: Supported Satellite Network Encapsulation

Topology Type	SDAC Discovery Protocol Packets	SDAC Control Protocol Packets	Data Packets
Hub and Spoke / Dual Home	Untagged LLC SNAP	Untagged TCP	802.1ad + customer payload
Layer 2 Fabric	Single tagged LLC SNAP	Single tagged TCP	802.1ah + customer payload **Note** Supports both 802.1ad and 802.1q as outer tag. Using 802.1q outer tag is a non-standard way of 802.1ah Encapsulation.
Simple Ring	Untagged LLC SNAP	Untagged TCP	802.1ah + customer payload **Note** Supports both 802.1ad and 802.1q as outer tag. Using 802.1q outer tag is a non-standard way of 802.1ah Encapsulation.

Advanced Satellite nV System Network Topologies

The host running Cisco IOS XR 32-bit Operating System supports all 3 network topologies namely, Dual home, Simple Ring, Layer 2 fabric network topologies, with the Cisco ASR 9000v Satellite from Release 5.1.1 onwards.

However with the Cisco NCS 5000 Series satellite, the host running Cisco IOS XR 32-bit Operating System supports only dual home network topology from Release 6.1.2 onwards.

Dual Home Network Topology

In the dual home network topology, 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 Network Architecture
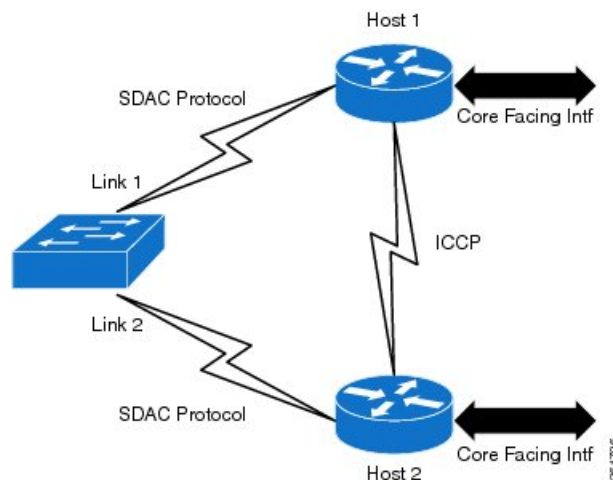
Features of the Dual Home Network Topology

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

- **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 here 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. If the host-priority is same and a simple ring topology is used, the lower hop count is used. If the hop count is same, the lower chassis MAC address is used for picking the active 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, the hop-count (path length) 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, on page 17.

- **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.

- **Hop Count** – This only affects simple rings, and provides a default traffic engineering algorithm based on number of intervening Satellite devices.

- **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.

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* .

Soft Minimum Active Links

The Soft minimum active links feature for dual home network topology 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.

Simple Ring Satellite nV Topology

These are the salient features of this topology:

- A satellite or ring of satellites can be dual-homed to two hosts. In the following figure, all the three satellites are connected to the redundant hosts Host A and Host B.

- The two hosts communicate using the ORBIT protocol over ICCP.

- In simple ring topology, the satellite chassis serial number is a mandatory configuration to identify the satellite.

- When the ring span is broken. the satellite and hosts detect the link failure using LOS mechanism and perform the necessary switching based Dual Home management.

- The link failure is detected by LOS (loss of signal) in the case of Ring and Hub and Spoke topologies.

Figure 4: Simple Ring Topology
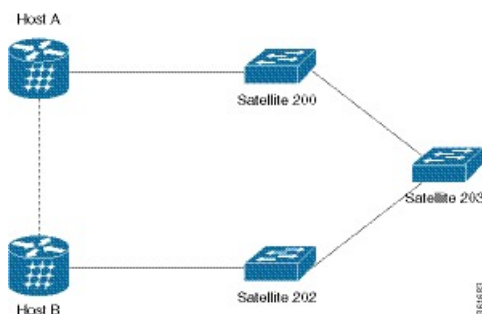
For configuration samples of Dual home architecture, see Satellite Configuration with Dual-homed Hosts. For a sample configuration of the simple ring topology, see the Configuration Examples for Satellite nV System section.

Simple Ring Topology Configuration

This is a sample ICL running configuration for a simple ring topology:

```
interface GigabitEthernet0/1/0/0
 ipv4 point-to-point
 ipv4 unnumbered Loopback10
 nv
  satellite-fabric-link network
   redundancy
    iccp-group 2
   !
   satellite 500
    remote-ports GigabitEthernet 0/0/0-9
   !
   satellite 600
    remote-ports GigabitEthernet 0/0/0-9
```

```
     !
   satellite 700
    remote-ports GigabitEthernet 0/0/0-9
 !
   satellite 800
    remote-ports GigabitEthernet 0/0/0-9
    !
```

Limitations of Simple Ring Topology

- If one of the satellite in a simple ring setup is removed from the ICL configuration, the subtending satellites remain in the connected state.

- When the configuration for a new satellite is applied, then the existing conflicting nV configuration must be removed.

- In a simple ring or a dual head topology, the local Layer-2 xconnects having satellite interfaces from active and standby satellites are not supported. All the satellites part of the local xconnect have to be in active state.

- Bundle ICL interfaces are not supported in the Simple Ring topology, but sat-ether port bundle is supported.

- When you activate a new image on the satellite in a simple ring topology based network, you need to initiate install transfer followed by an install activate. For more information, see Installing a Satellite.

- Satellite access bundles with bundles members spanning across different satellites are not supported. For example, if there are three satellites, namely, sat 100, sat200, and sat 300, then you cannot have an access bundle with members from sat 100, sat 200 and sat 300. This is because each of the satellite could be active/standby to different hosts and hence leads to unpredictability in their behavior.

- In a simple ring topology, the access ports might not achieve full line rate. All downstream traffic sent from host to satellite for an access port is subject to a port level shaper on the host. There is an overhead of 22 bytes of nv tag on each data packet. This shaper oversubscribes the port based flow to account for the overhead. But there is a hardware limitation on the traffic manager used on the host line card, which has a metering granularity of 4 bytes. Therefore the host interprets 22 bytes as 20 bytes. Hence, shaping is done assuming 20 bytes overhead for each packet. However, 22 bytes are stripped off from each packet received on ICL on the satellite before it is sent out on access port. This results in approximately 3% undersubscription of data packets on the access port. The percentage of undersubscription reduces as the packet size increases.

Layer 2 Fabric Network Architecture

In the Layer 2 Fabric network architecture, a satellite is connected to one or two hosts through one of two Ethernet Virtual Circuits (EVC) of Layer 2 Fabric network. An EVC can be identified by two transports VLAN IDs, such as TP-VID-S and TP-VID-H. TP-VID-S is the VLAN ID assigned by the satellite side transport and TP-VID-H is the VLAN ID assigned by the host. The CFM based Fast Fabric Link Failure Detection is supported only in the Layer 2 Fabric Network Architecture. These illustrations show different variants of Layer 2 Fabric network topology.

> **Note** CFM is mandatory in the case of Layer 2 Fabric Network Architecture to ensure that link failure detection is fast.

Note Fabric CFM is supported on physical ICLs with minimum CCM interval of 1 second for Cisco ASR 9000v 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.

Figure 5: Layer 2 Fabric Satellite Network Architecture with dual host
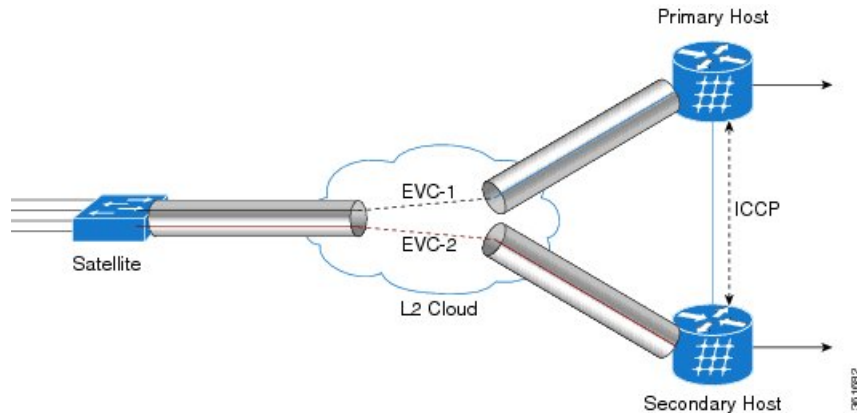
Figure 6: Layer 2 Fabric Single Home (SH) with Single Physical ICL on Host and Satellite

Figure 7: DH : Single physical ICLs on hosts and Multiple physical ICLs on satellite
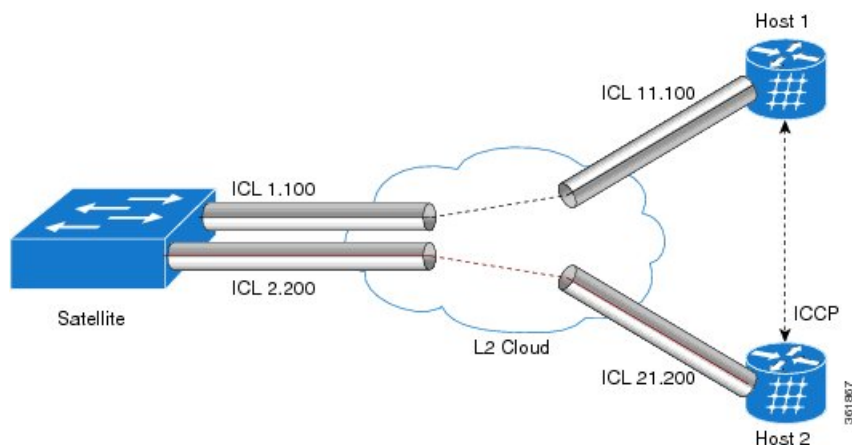
Limitations of Layer 2 Fabric Network Topology

1. Bundle interfaces are not supported in Layer 2 Fabric architecture. In the Layer 2 Fabric topology, bundle ICL is not supported but sat-ether port bundle is supported.

2. Point to Multi-point Layer 2 cloud is not supported.

3. A Satellite can support only one encapsulation on a given physical interface. So, the Layer 2 fabric connections from both hosts must be configured with the same encapsulation type.

4. A Satellite cannot support multiple Layer 2 fabric connections with the same VLAN on the same physical ICL interface.

5. When Satellite ethernet bundle interfaces are configured on the access ports, the bundle wait timer needs to be configured to zero to get better convergence.

6. Cisco ASR 9000v has these limitations on the Layer 2 Fabric network topology:

    • The usable VLAN range is from 2 to 4093.

    • Only four 10 Gigabit Ethernet ICLs can be used on the ports 1/45 to 1/48.

    • Only two 1 Gigabit Ethernet ICLs can be used on the ports 1/45 and 1/46.

    • Only 44 satellite ports are present (namely, Gig1/1 to Gig1/44).

    • On a satellite, a VLAN can be used on a single ICL only. Though there can be two L2FAB ICLs on a physical port, they must not use the same VLAN.

    • L2FAB ICL cannot be configured on a bundle interface.

    • When service policies are applied on L2FAB ICL, it will fetch policy statistics at physical port level due to hardware limitation on the platform.

7. When a single physical ICLs is connected on both hosts and multiple physical ICLs are connected on satellite, the following limitations are applied:

    • A satellite can support different encapsulations from both hosts till the layer 2 fabric connections from both the hosts terminate on different satellite ICL links.

    • A satellite can support multiple Layer 2 fabric connections with same VLAN on different physical ICL interfaces.

8. In the Dual-home scenario, for multiple satellites with a single physical ICL connected on hosts and Satellite, both hosts must use same encapsulation and VLANs must be different from Satellite point of view.

9. In the Dual-home scenario, for multiple satellites with multiple physical ICLs on one host, both hosts must use same encapsulation and VLANs must be different from Satellite point of view.

10. In the Dual-home scenario, if the L1 and L2 Fabric connections are on the same host, there is no limitation till same type of topology is used to connect both hosts on the Satellite side. There is no restriction on the host side.

11. In L2 Fabric ICL topology, the ICL port encapsulation does not change without shutting down the ICL port. To change the encapsulation, in the sub-interface, firstly do a port shut and commit the configuration. Secondly, with no encapsulation commit the configuration. And then change the encapsulation, do a `no shut` and commit the configuration.

12. In a Layer 2 Fabric network topology, the access ports might not achieve full line rate. All downstream traffic sent from host to satellite for an access port is subject to a port level shaper on the host. There is an overhead of 22 bytes of nv tag on each data packet. This shaper oversubscribes the port based flow to account for the overhead. But there is a hardware limitation on the traffic manager used on the host

line card, which has a metering granularity of 4 bytes. Therefore the host interprets 22 bytes as 20 bytes. Hence, shaping is done assuming 20 bytes overhead for each packet. However, 22 bytes are stripped off from each packet received on ICL on the satellite before it is sent out on access port. This results in approximately 3% undersubscription of data packets on the access port. The percentage of undersubscription reduces as the packet size increases.

General Limitations of Satellite nV System Network Topologies

1. 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.

2. With the Cisco NCS 5000 Series satellite, the host running Cisco IOS XR 32-bit Operating System supports only dual home network topology from Release 6.1.2 onwards.

3. 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.

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 Cisco NCS 5000 Series Satellite uses a similar hashing function to load-balance satellite to host traffic between the members of the ICL bundle.

The Cisco ASR 9000v Satellite uses random hashing 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   Cisco IOS XR Software supports the co-existence of redundant and non-redundant ICL modes on the same satellite shelf from Cisco IOS XR Software Release 4.3.x onwards in the case of Cisco ASR 9000v satellite and from Cisco IOS XR Software Release 6.0.x for Cisco NCS 5000 Series satellite.

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. Additionally, redundant ICL mode is not supported for Layer 2 fabric and simple ring network topologies.

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

Multiple ICL for the Layer 2 Fabric Network Topology

This feature allows you to configure 4 physical Layer 2 Fabric ICLs for each satellite. This indicates that in a dual home topology setup, each physical ICL port can have a pair of Layer 2 Fabric ICLs, one each from active and standby host. Hence, the Layer 2 Fabric network topology provides up to 8 Layer 2 Fabric ICLs. 1GE ports supporting Dynamic ICL feature can be used as an ICL on Cisco ASR9000v or Cisco ASR9000v2. These ICLs can be a combination of 1GigE and 10GigE links. Layer 2 Fabric ICLs to active and standby hosts can terminate on a single physical port on satellite or on separate ports. The encapsulation for an ICL can be based on IEEE 802.1ad or 802.1q, which is independent of encapsulation on other ICLs.

Note   This feature is supported only on Cisco ASR 9000v and ASR 9000v2 satellite boxes. Bundling of multiple ICLs is not supported. Failover to standby host does not happen unless there is total loss of connectivity to the active host.

Multiple ICL for the Layer 2 Fabric Network Topology Configuration

A time delay of approximately 40 seconds is required before and after a cross link in L2fab topology is added or deleted.

If a CFM on ICL is present, first remove the ICL interface configuration and then remove the nV satellite global configuration.

Note   If the CCM fails in any one of the ICL, CFM brings down the SDAC immediately. The remaining ICLs take 30 seconds time to go down SDAC discovery (after SDAC timeout). Hence, satellite hardware programming takes time to delete the MIM port. So, we recommend to follow the above steps to remove the nV satellite configuration when you have the CFM on fabric port.

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 45.

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.

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.

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.

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 network topology, 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.

Note In the case of Cisco IOS XR Software Release 6.0.1, this feature is only supported on Cisco NCS 5000 series satellites.

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. Multiple redundancy groups are supported on all the advanced satellite network topologies, such as:

- Dual Home network topology
- Hub & Spoke topology
- Simple Ring topology
- Layer 2 Fabric topology

In a multiple ICCP redundancy group setup, failovers can happen similar to the case of a basic Dual Home setup. See the Dual Home Network Topology, on page 15 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. Layer 2 Multicast Offload with Hub and Spoke Topology is supported.

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 8: Multiple Access Bundles over single ICL Bundle to Single Host
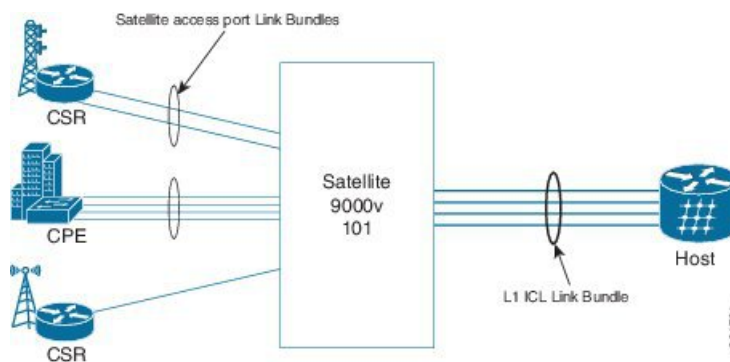
In this case, different members of the access bundle can be configured over different ICL bundle links.

Figure 9: Multiple Access Bundles over many ICL bundles to Single Host
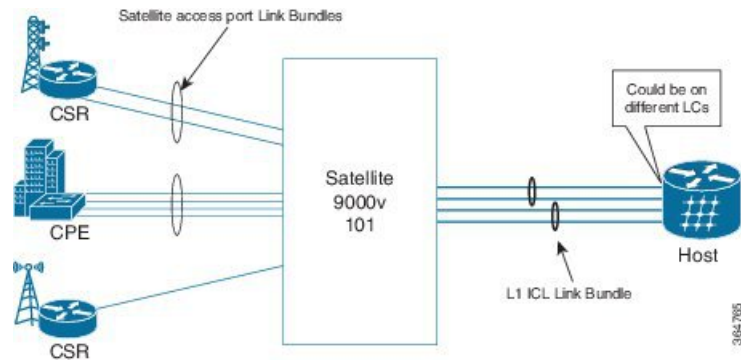
Figure 10: Multiple Access Bundles over many ICL bundles to Dual-Homed Hosts
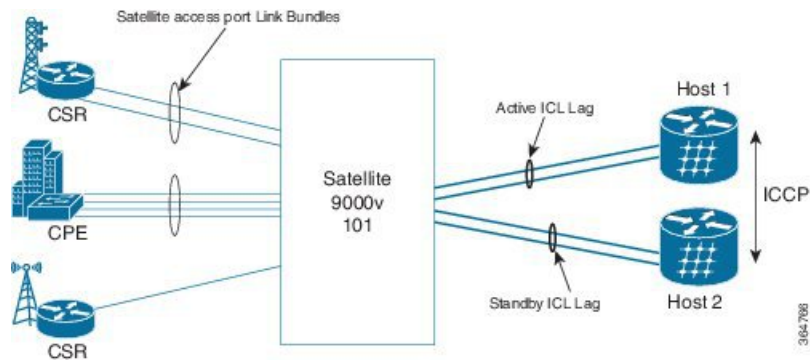
Note In the figures 8, 9, and 10, the satellite can also be Cisco NCS 5000 Series Router in place of Cisco ASR 9000v.

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 11: Access Bundle over ICL Bundles to two different hosts in CPE dual homing
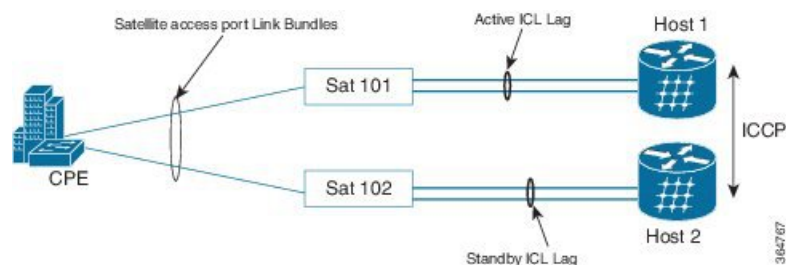
Port Partitioning

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.

Note Port partitioning is not supported for simple ring and Layer 2 fabric network topologies.

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. The BFD over Satellite is supported only on ASR 9000 Enhanced Ethernet Line Card and is supported in asynchronous mode. BFD over satellite is not supported in echo mode.

Note
- The **bfd multipath include location** *node-id* command is not supported on ASR 9000 Ethernet Line Card. Hence, BFD over Satellite Interfaces feature does not work on the ASR 9000 Ethernet Line Card.

- The nV Satellite access port bundles does not support BFD over bundles (BoB) when the Inter-Chassis Link (ICL) is also a bundle (Bundle over Bundle topology).

BNG Interoperability

The BNG interoperability allows BNG to exchange and use information with other larger heterogeneous networks. These are the key features:

- BNG Coexists with ASR9001:

  ASR9001 is a standalone high processing capability router that comprises of a route switch processor (RSP), linecards (LC), and ethernet plugs (EPs). All BNG features are fully supported on the ASR9001 chassis.

- BNG Supports nV Satellite:

  The only topology that is supported with BNG-nV Satellite is - bundled Ethernet ports on the CPE side of the Satellite node connected to the Cisco ASR 9000 through non-bundle configuration (static-pinning). That is,

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

  Although the following topology is supported on Satellite nV System (from Cisco IOS XR Software Release 5.3.2 onwards), it is not supported on BNG:

  - Bundled Ethernet ports on the CPE side of the satellite node, connected to the Cisco ASR 9000 through bundle Ethernet connection.

- BNG interoperates with Carrier Grade NAT (CGN):

  To address the impending threat from IPv4 address space depletion, it is recommended that the remaining or available IPv4 addresses be shared among larger numbers of customers. This is done by using CGN,

which primarily pulls the address allocation to a more centralized NAT in the service provider network. NAT44 is a technology that uses CGN and helps manage depletion issues of the IPv4 address space. BNG supports the ability to perform NAT44 translation on IPoE and PPPoE-based BNG subscriber sessions.

Note For BNG and CGN interoperability, configure the BNG interface and the application service virtual interface (SVI) on the same VRF instance.

Restrictions

- Only bundle access with non-bundle ICLs are supported for BNG interfaces over Satellite nV System access interfaces.

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

All Layer-2 and Layer-3 multicast features, including IGMP, IGMP snooping, PIM, mLDP, MVPNv4,MVPNv6, P2MP TE, are supported on Satellite Ethernet interfaces, as they are supported 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 Topology

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 home Satellite nV System topology. In the Dual Head

Satellite nV System topology, 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 node that is defined as slave. The port that is defined as master controls the loopback status on the remote node. You can conduct various performance and quality testing on the remote port that is in slave state using this feature. This loopback mode is disruptive because the traffic does not flow through the router in slave 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 node that is in master state 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 master node to enable and disable the loopback respectively on the slave 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 master
    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.

CFM Performance Monitoring

The Connectivity Fault Management (CFM) feature is supported on the nV Satellite system. To use this feature the Maintenance End Points (MEPs) must be configured on the satellite access ports so that they belong to the CFM service over which performance monitoring is being implemented. In an nV Satellite system, the MEPs that are set on the satellite access ports are configured on the Cisco ASR 9000 Series host. For additional details on CFM configuration, refer to the topic "Ethernet CFM" in the *Interface and Hardware Component Configuration Guide for Cisco ASR 9000 Series Routers*

CFM Performance Monitoring (PM) applies time-stamps or sequence numbers on CFM PM frames (PDUs), at the two MEPs (controller and responder) between which performance is being measured. The controller MEP initiates the PM process by sending frames to the responder. The responder MEP processes the frames and sends them back to the controller, for the latter to evaluate the frames and calculate the PM result.

Prior to Cisco IOS XR Release 5.3.1, MEPs configured on Satellite Access ports processed (time-stamped) performance monitoring frames at the host. To ensure accuracy in performance monitoring, the frames are now processed on the satellite instead of the ASR9K host.

The features supported in Release 5.3.1 are:

- Up MEPs process (timestamp) only those CFM PM frames on the satellite that meet the criteria (see processing locations table), and the remaining CFM PM frames are processed on the ASR9K host.

  Down MEPs process all CFM PM frames on the ASR9K host.

- Only the responder functionality for two-way delay measurement is available on the satellite. The controller functionality is managed by the ASR9K host. Also, the two CFM processes, the Loopback process and the Synthetic Loss Measurement process are managed by the ASR9K host.

Note   From Release 5.3.1 onwards, if the satellite supports it, the processing of relevant frames on the satellite happens by default. To turn off the satellite processing functionality, use the **nv satellite sla processing disable** command in the ethernet cfm configuration mode. If the frames are not processed on the satellite, they are processed on the ASR9K host by default. However, this is not supported on the IOS XR 64 Bit operating system.

nV Satellite CFM Performance Monitoring Processing Locations

This table lists various parameters and functionalities, and the corresponding locations where the CFM PM frames are processed. The location can either be the satellite or the ASR9K host.

Parameter	ASR901	ASR9000v
Up MEP responder	Satellite	Satellite
Down MEP responder	Host	Host
Delay Measurement (v0)	Satellite	Satellite
Delay Measurement (v1)	*Frames are not processed*	Satellite
Bundle ICLs in Hub & Spoke topologies	Host	Host
Main interfaces, default, native VLAN, Double Tagged EFPs	Host	Host
VLAN ID range	Host	Host
EFPs matching a single dot1q VLAN ID	Satellite	Satellite
EFPs matching a single dot1ad VLAN ID	Host	Host
EFPs configured with a dot1q Tunneling Ethertype 0x9100 or 0x9200	Host	Host
Untagged EFPs	Satellite	Satellite
Unprotected satellite access ports	Satellite	Satellite
Maximum number of MEPs per Ethernet Flow Point (EFP)	1 - Satellite More than 1 - Host	1 - Satellite More than 1 - Host

Parameter	ASR901	ASR9000v
Maximum number of MEPs per Satellite *Frames from MEPs that exceed the limit are not processed either on the satellite or host, and a warning is logged.*	Up to 40 - Satellite	Up to 44 - Satellite

Note   If an unsupported configuration is committed, the CFM PM processing takes place on the host instead of the satellite. A warning is issued and it can be viewed in the output of the **show ethernet cfm configuration-errors** command.

Related Commands   Refer to the *Interface and Hardware Component Command Reference for Cisco ASR 9000 Series Routers* guide for details of these related commands:

Command	Description
nv satellite sla processing disable	
show ethernet cfm configuration-errors	
show ethernet cfm interfaces statistics	
show ethernet cfm local meps	

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

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 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) 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 and

- Fabric CFM for both physical and bundle ICLs and dual head topologies.

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

Restrictions of the Satellite nV System

Software restrictions of the satellite system are:

- The nV Satellite access interfaces are expected to be deployed as the 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.

- When you convert from one satellite topology to another topology, such as hub and spoke to Layer 2 Fabric network topology, you must remove the existing ICL configurations from the interface in one commit followed by the new ICL configurations on the interface in a separate commit.

- Access bundles having both satellite and non-satellite interfaces is not supported.

- For all nV satellite topologies, a maximum of up to 200 satellite access ports can be supported per NPU on both the –TR (Packet Transport Optimized) and the –SE (Services Edge Optimized) line card variants.

  With the buffer, memory and 8 queues per port restrictions, the use of –TR line card variants has restrictions for advanced, large scale nV satellite deployments like simple ring and L2 fabric topologies. For such network mode nV satellite configuration, an –SE line card variant is strongly recommended for hosting the satellites on the Cisco ASR 9000 Series Router host

- On A9K-40GE-TR and A9K-40GE-SE line card variants, only the first 16 ports can be used as ICL or fabric ports to host nV satellites. A9K-40GE-SE supports a maximum of 170 nV satellite access ports per NPU and A9K-40GE-TR supports a maximum of 40 nV satellite access ports per NPU.

- The nV Satellite access port bundles does not support BFD over bundles (BoB) when the ICL is also a bundle (Bundle over Bundle topology)

- 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

- The feature Provider Backbone Bridging Ethernet Virtual Private Network (PBB EVPN) is not supported with nV Satellite Access Interface bundles over ICL bundles.

- Mixing of interface types (1G, 10G, 100G) is not supported for bundle ICLs.

- The maximum number of satellites per system is 16.

- The range in which the satellite ID can be configured is from 100 to 230.

- Dual-home nV Satellite network topology 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 dual-home nV Satellite network topology.

- 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**.

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.

This table provides the release-wise support for CDP, LLDP, UDLD, Ethernet OAM, E-LMI and CFM:

		Single-Homed		**Dual-Homed**	
		Physical ICL	Bundle ICL	Physical ICL	Bundle ICL
Access Physical	CDP	IOS XR 4.2.1	IOS XR 5.1.1	IOS XR 4.2.1	IOS XR 5.1.1
	LLDP	IOS XR 4.2.1	IOS XR 4.3.1	IOS XR 4.2.1	IOS XR 4.3.1
	UDLD	IOS XR 4.2.1	IOS XR 5.3.1	IOS XR 4.2.1	IOS XR 5.3.1
	OAM	IOS XR 4.2.1	IOS XR 5.1.1	IOS XR 4.2.1	IOS XR 5.1.1
	E-LMI	IOS XR 4.2.1	N.A.	IOS XR 4.2.1	N.A.
	CFM	IOS XR 4.2.1	IOS XR 5.1.1	IOS XR 4.2.1	IOS XR 5.1.1
	SLA (802.1ag / ITU-T Y.1731)	IOS XR 4.2.1	IOS XR 5.1.1	IOS XR 4.2.1	IOS XR 5.1.1
	EDPL	IOS XR 5.2.2	IOS XR 5.2.2	IOS XR 5.2.2	IOS XR 5.2.2

		Single-Homed		Dual-Homed	
Access LAG/Bundle	CDP	IOS XR 4.2.1	IOS XR 5.3.2	IOS XR 4.2.1	IOS XR 5.3.2
	LLDP	IOS XR 4.2.1	IOS XR 5.3.2	IOS XR 4.2.1	IOS XR 5.3.2
	UDLD	IOS XR 4.2.1	IOS XR 5.3.2	IOS XR 4.2.1	IOS XR 5.3.2
	OAM	IOS XR 4.2.1	N.A.	IOS XR 4.2.1	N.A.
	E-LMI	N.A.	N.A.	N.A.	N.A.
	CFM	IOS XR 4.2.1	IOS XR 5.3.2	IOS XR 4.2.1	IOS XR 5.3.2
	SLA (802.1ag / ITU-T Y.1731)	IOS XR 4.2.1	IOS XR 5.3.2	IOS XR 4.2.1	IOS XR 5.3.2
	EDPL	IOS XR 5.2.2	N.A.	IOS XR 5.2.2	N.A.

Note 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

- RSP2 is not qualified with Cisco NCS 5000 Series satellite.

- 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 Cisco Enhanced Ethernet line card, SIP-700 line cards can co-exist in the system but cannot be used to connect directly to the satellite fabric ports. They can forward traffic to and from the satellite ports.

- Cisco ASR 9001 chassis has low disk space. Hence, memboot is preferred. RSP3 TR card has a low memory. Hence, nV pies for Cisco NCS 5000 Series satellite cannot be loaded unless there is 2 GB of free space.

- 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. This requires further free space and the low disk size on Cisco ASR 9001 can impose severe restrictions on the use of this feature.

- 100G ICLs are only supported on the 8x100G or 4x100G Cisco ASR 9000 High Density 100GE Ethernet line cards.

- 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.

- 1GE ports are just 10GE ports operating at restricted speeds with no auto negotiation or support for 10/100 M speeds. The interfaces continue to be named as 10GE interfaces with 1GE 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.

- 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.

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 6.3.1 , the command for configuring nV satellite auto image upgrade feature is **upgrade [on-connect | on-first-connect]**

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

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. Therefore, auto upgrade of satellites in a chain or a ring is done in series. So, for long chains or rings, it may be more time efficient to do the operation manually.

> **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** On Cisco ASR 9000 Series Router host running IOS XR 64 Bit operating system, auto upgrade is not supported for the Cisco NCS 5000 Series satellite due to restrictions in supporting the satellite nV pie for the Cisco NCS 5000 Series satellite on the Cisco IOS XR 64 Bit operating system

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:**	Enters global configuration mode.

	Command or Action	Purpose
	`RP/0/RSP0/CPU0:router# configure`	
Step 2	**nv** **Example:** `RP/0/RSP0/CPU0:router(config)# nv`	Enters the nV configuration submode.
Step 3	**satellite** *Sat ID* **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. .
Step 4	**serial-number** *string* **Example:** `RP/0/RSP0/CPU0:router(config-satellite)# serial-number CAT1521B1BB`	Serial number is used for satellite authentication.
Step 5	**description** *string* **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** *type_name* **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** *address* **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** *password* **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** *interface-type slot/subslot/port-range* **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 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**	Saves configuration changes.

Command or Action	Purpose
Example: RP/0/RSP0/CPU0:router(config)# end or commit	• 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.
		
		• 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 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 43.

> ✎
>
> **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 81 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** *interface-name* **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, 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:** `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	**nV** **Example:** `RP/0/RSP0/CPU0:router(config-if)# nv`	Enters the Network Virtualization configuration mode.
Step 7	**satellite-fabric-link satellite** *<id>* **Example:**	Specifies that the interface is an ICPE inter-chassis link.

Command or Action	Purpose
`RP/0/0RSP0/CPU0:router(config-int-nv)# satellite-fabric-link satelite 200`	
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/0/0RSP0/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.
	
	- 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.
	
	For information on QoS configuration on ICLs , see *Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide*.

Configuring the Inter-Chassis Links in a Dual Home Network Topology

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

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.

	Command or Action	Purpose
Step 2	**interface** *interface-name* **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** *num*
	
	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)# redundancy`
	`RP/0/RSP0/CPU0:router(config-nV-red)# minimum preferred links 2`
		(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.
		
		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**.

Step 8	**member neighbor** 9.9.9.9 **Example:** `RP/0/RSP0/CPU0:router(config-interface)# member neighbor 9.9.9.9`	Configures the LDP neighbor.
Step 9	**backbone interface** *interface_type* **Example:** `RP/0/RSP0/CPU0:router(config-interface)# backbone interface TenGigE0/1/0/3`	(Optional) Configures the backbone interface for the PE isolation.

	Command or Action	Purpose
Step 10	**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 11	**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. - 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 Inter-Chassis Links for a Simple Ring Topology

These are the steps for configuring Inter-chassis links in the case of a simple ring topology.

Procedure

	Command or Action	Purpose
Step 1	**configure** **Example:** `RP/0/RSP0/CPU0:router# configure`	Enters global configuration mode.
Step 2	**redundancy iccp-group** **Example:** `RP/0/RSP0/CPU0:router(config)# redundancy iccp group 2`	Configures the redundancy ICCP group between the hosts.
Step 3	**member neighbor** 9.9.9.9 **Example:**	Configures the LDP neighbor.

	Command or Action	Purpose
	`RP/0/RSP0/CPU0:router(config-redundancy-iccp-group)# member neighbor 9.9.9.9`	
Step 4	**host-priority <0-255>** **Example:** `RP/0/RSP0/CPU0:router(config-redundancy-iccp-group)# host-priority 128`	(Optional) Specifies the priority for the satellite on each of the host. The host with the lower priority is preferred as the active host. The default priority is 128.
Step 5	**nv satellite system-mac** *<mac_address>* **Example:** `RP/0/RSP0/CPU0:router(config-redundancy-iccp-group)# nv satellite system-mac dcddc.dcdc.dcdc`	(Optional) Specifies the MAC address. Two hosts in the same redundancy group will sync up the system MAC address and satellite priority information. The System MAC must be the same. If it is different, then the Host with low chassis MAC gets priority. If the system MAC is not configured, then it uses low host chassis MAC as the system MAC.
Step 6	**interface** *interface-name* **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 7	**ipv4 point-to-point** **Example:** `RP/0/RSP0/CPU0:router(config-if)# ipv4 point-to-point`	(Optional) Configures the IPv4 point to point address.
Step 8	**ipv4 unnumbered loopback0** **Example:** `RP/0/RSP0/CPU0:router(config-if)# interface unnumbered loopback0`	(Optional) Configures the IPv4 loopback address on the interface.
Step 9	**nv** **Example:** `RP/0/RSP0/CPU0:router(config-if)# nV`	Enters the nV satellite mode.
Step 10	**satellite-fabric-link network** **Example:** `RP/0/RSP0/CPU0:router(config-if-nV)# satellite-fabric-link network`	Specifies the network type of ICPE inter-chassis link.
Step 11	**redundancy iccp-group** **Example:** `RP/0/RSP0/CPU0:router(config-sfl-network)# redundancy iccp-group 2`	Configures the ICCP redundancy group.
Step 12	**satellite** *<satellite id>* **Example:**	Specifies the satellite ID of the satellite.

	Command or Action	Purpose
	`RP/0/RSP0/CPU0:router(config-sfl-network)# satellite 500`	
Step 13	**remote-ports** *interface-type* **Example:** `RP/0/RSP0/CPU0:router(config-sfl-network)# remote-ports GigabitEthernet 0/0/9,5`	Configures the remote satellite ports for the satellite 500.
Step 14	**satellite** *<satellite id>* **Example:** `RP/0/RSP0/CPU0:router(config-sfl-network)# satellite 600`	Specifies the satellite ID of the connected satellite in the simple ring..
Step 15	**remote-ports** *interface-type* **Example:** `RP/0/RSP0/CPU0:router(config-sfl-network)# remote-ports GigabitEthernet 0/0/9,5`	Configures the remote satellite ports for the satellite 600.

Step 16	**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.
		
		- 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 topology, 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. SeeSatellite Interface Configuration.

Configuring the Fabric CFM for L1 Physical ICL and L2 Fabric Network

This procedure gives you the steps to configure CFM on a Satellite nV Fabric link interface. CFM can be configured with L1 Physical ICL or L2 Fabric Network architecture. In both these cases, it can also be configured on Dual Home network topologies.

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.

	Command or Action	**Purpose**

Step 2	**interface** *interface-name*
	
	 Example:
	
	 (for L1 Physical ICL)
	
	 `RP/0/RSP0/CPU0:router(config)# interface Gigabit 0/1/0/0`
	
	 Example:
	
	 (for L2 Fabric Network)
	
	 `RP/0/RSP0/CPU0:router(config)# interface Gigabit 0/1/0/0.1`
		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 satelite 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
	- 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.

   When 10G are supported, 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 9001, Cisco ASR 9904, Cisco ASR 9006, Cisco ASR 9010, 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 43, Configuring the Host IP Address, on page 46 and Configuring Inter-Chassis Links and IP Connectivity, on page 47. 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.

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.

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.

✎ **Note**   The connection between host and satellite can be established and routed over an L2 access network.

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
!
```

SyncE Offload Configuration:

In general, the syncE offload is configured in global mode for Host A, Host B. SyncE is also configured on the downstream device, which is a Cisco ASR 9000 Series Router in this example.

Configuring syncE offload in global mode on Host A:

By default, QL option 1 is supported with the syncE offload configuration in global mode. In addition, you can also explicitly configure QL option 1 or 2 as:

```
frequency synchronization
 quality itu-t option 2 generation 2
!
```

Configuring syncE offload in global mode on Host B:

```
frequency synchronization
 quality itu-t option 2 generation 2
!
```

Configuring syncE in global mode in downstream device (Cisco ASR 9000 Series Router):

```
frequency synchronization
 quality itu-t option 2 generation 2
clock-interface timing-mode system
!
```

Configuring Clock Interface on Host A:

```
clock-interface sync 0 loca 0/RSP0/CPU0
 port-parameters
  bits-input e1 crc-4 sa4 hdb3
 !
 frequency synchronization
  selection input
  priority 1
  wait-to-restore 0
  ssm disable
  quality receive exact itu-t option 2 generation 2 PRS
 !
!
```

Configuring Clock Interface on Host B:

```
clock-interface sync 0 loca 0/RSP0/CPU0
 port-parameters
  bits-input e1 crc-4 sa4 hdb3
 !
 frequency synchronization
  selection input
  priority 1
  wait-to-restore 0
  ssm disable
  quality receive exact itu-t option 2 generation 2 PRS
 !
!
```

Configuring Clock Interface in Downstream Device (Cisco ASR 9000 Series Router):

```
clock-interface sync 0 location 0/RSP0/CPU0
 port-parameters
```

```
 bits-output e1 crc-4 sa4 hdb3
 !
 frequency synchronization
  wait-to-restore 0
  ssm disable
 !
!
```

Configuring syncE offload on ICLs

SyncE offload is configured at the physical interface level for both physical and bundle ICLs.

Host A Configuration:

```
interface TenGigE0/0/0/0
 bundle id 100 mode on
 frequency synchronization
  wait-to-restore 0
  quality transmit exact itu-t option 2 generation 2 PRS
 !
!
```

Host B Configuration:

```
interface TenGigE0/0/0/0
 ipv4 point-to-point
 ipv4 unnumbered Loopback100
 frequency synchronization
  wait-to-restore 0
  quality transmit exact itu-t option 2 generation 2 STU
 !
 nv
  satellite-fabric-link satellite 100
   redundancy
    iccp-group 1
   !
   remote-ports GigabitEthernet 0/0/0-20
  !
 !
!
```

Configuring syncE offload on the interface connected to sat-ether ports in downstream device (Cisco ASR 9000 Series Router):

```
interface GigabitEthernet0/0/0/0
 frequency synchronization
  selection input
  wait-to-restore 0
 !
!
```

Configuration for enabling 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
```

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 41.

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 advanced Satellite system topologies such as dual head, simple ring, or Layer 2 Fabric network topologies.

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.

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 55 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 In the case of simple ring topology, the image must be transferred to all the satellites using **install nv satellite transfer** *<range of satellites>* command followed by **install nv satellite activate** *<range of satellites>* command. You cannot use only the install nv satellite activate command in the case of simple ring topology.

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 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).

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*.

```
RP/0/RSP0/CPU0: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


RP/0/RSP0/CPU0:TARDIS# 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 The software versions show up as N/A in Cisco IOS XR Software Release 6.0.0 because inband image upgrade and monitoring is supported only from Cisco IOS XR Software Release 6.0.1 for Cisco NCS 5000 series satellites.

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.

Check if Upgrade is Required

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

```
RP/0/RSP0/CPU0: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
```

Output of Status Check

This example shows the output of **show nv satellite status** command for a Satellite configured in dual home network topology.

```
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 90000v. 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 Advanced Network Topologies

The following examples show commands used for Dual Home Network Topology and Simple Ring Topology.

Dual Home Network Topology

```
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
```

Simple Ring Topology

```
RP/0/RSP1/CPU0:Router# show nv satellite topology interface tenGigE 0/3/0/6

TenGigE0/3/0/6
--------------
  Redundancy-Group: 10
  Discovery status: Running
  Satellites:
    Satellite 100 (BVID 2002)
    ------------------------
     Received Serial Number: CAT1547B30S
      MAC address: 4055.3957.5f50
      Satellite fabric links:
        TenGE/0/0/0 (Remote ID: 0x1):
          Host (TenGigE0/3/0/6)
        TenGE/0/0/3 (Remote ID: 0x4):
          Sat 200 (TenGE/0/0/3 (Remote ID: 0x4))

    Satellite 200 (BVID 2003)
    ------------------------
      Received Serial Number: CAT1708U0LV
      MAC address: 8478.ac01.d2d8
      Satellite fabric links:
        TenGE/0/0/3 (Remote ID: 0x4):
          Sat 100 (TenGE/0/0/3 (Remote ID: 0x4))
        TenGE/0/0/0 (Remote ID: 0x1):
          Remote port not yet identified
```

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 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 chassis**, **show inventory fans** commands in the admin configuration mode to monitor the status of satellite inventory.

Note Along with a physical entity for the ASR-9000v/v2 satellite, logical entities are also created for the satellite and the power module. Both these entities (physical and logical) are seen in the inventory details command output and in SNMP MIBs. The logical entities can be identified by the lack of serial number (SN) and version identifier (VID).

```
RP/0/RSP0/CPU0:router(admin)# show inventory chassis

NAME: "module 0/RSP0/CPU0", DESCR: "ASR9K Fabric, Controller, 4G memory"
PID: A9K-RSP-4G, VID: V02, SN: FOC143781GJ
...
NAME: "fantray SAT100/FT0/SP", DESCR: "ASR9000v"
PID: ASR-9000v-FTA, VID: V00 , SN: CAT1507B228

NAME: "module SAT100/0/CPU0", DESCR: "ASR-9000v GE-SFP Line Card"
PID: ASR-9000v, VID: N/A, SN: /*Logical Entity of the Satellite*/
NAME: "module mau GigabitEthernet100/0/CPU0/8", DESCR: "CISCO-AVAGO    "
PID: SFP-GE-S, VID: V01, SN: AGM1424P08N

NAME: "module mau TenGigE100/0/CPU0/3", DESCR: "CISCO-FINISAR   "
PID: SFP-10G-SR, VID: V02, SN: FNS144502Y3
```

```
NAME: "power-module SAT100/PM0/SP", DESCR: "ASR-9000v Power Module"
PID: ASR-9000v, VID: N/A, SN: /*Logical Entity of the Power Module*/
NAME: "Satellite Chassis ASR-9000v ID 100", DESCR: "ASR9000v"
PID: ASR-9000v-AC-A, VID: V00 , SN: CAT12345678/*Physical Entity of the Satellite*/



RP/0/RSP0/CPU0:router(admin)# show inventory fans

NAME: "fantray 0/FT0/SP", DESCR: "ASR-9006 Fan Tray"
PID: ASR-9006-FAN, VID: V02, SN: FOX1519XHU8

NAME: "fantray 0/FT1/SP", DESCR: "ASR-9006 Fan Tray"
PID: ASR-9006-FAN, VID: V02, SN: FOX1519XHTM

NAME: "fantray SAT100/FT0/SP", DESCR: "ASR9000v"
PID: ASR-9000v-FTA, VID: V01 , SN: CAT1531B4TC

NAME: "fantray SAT101/FT0/SP", DESCR: "ASR9000v"
PID: ASR-9000v-FTA, VID: V01 , SN: CAT1542B0LJ

NAME: "fantray SAT102/FT0/SP", DESCR: "ASR9000v"
PID: ASR-9000v-FTA, VID: V01 , SN: CAT1531B4T7


RP/0/RSP0/CPU0:sat-host(admin)# show inventory | b GigabitEthernet100/

NAME: "module mau GigabitEthernet100/0/CPU0/0", DESCR: "CISCO-FINISAR    "
PID: SFP-GE-S, VID:  , SN: FNS11350L5E

NAME: "module mau GigabitEthernet100/0/CPU0/1", DESCR: "CISCO-FINISAR    "
PID: SFP-GE-S, VID: V01, SN: FNS0934M290

NAME: "module mau GigabitEthernet100/0/CPU0/2", DESCR: "CISCO-FINISAR    "
PID: SFP-GE-S, VID:  , SN: FNS12280L59




RP/0/RSP0/CPU0:TARDIS(admin)#show inventory
...
NAME: "fantray SAT100/FT0/SP", DESCR: "Cisco NCS 5002 Series Router Fan Back"
PID: NCS-5002-FN-BK, VID: N/A, SN: N/A
NAME: "fantray SAT100/FT1/SP", DESCR: "Cisco NCS 5002 Series Router Fan Back"
PID: NCS-5002-FN-BK, VID: N/A, SN: N/A
NAME: "power-module SAT100/PM0/SP", DESCR: "Cisco NCS 5000 Series Router power AC 650W Back"
PID: NC5K-PAC-650W-BK=, VID: V01, SN: LIT1919198Z
NAME: "power-module SAT100/PM1/SP", DESCR: "Cisco NCS 5000 Series Router power AC 650W Back"
PID: NC5K-PAC-650W-BK=, VID: V01, SN: LIT1919199H
NAME: "Satellite Chassis NCS5002 ID 100", DESCR: "80-Port 10 GE + 4-Port 100GE NCS5002
Chassis"
PID: NCS-5002, VID: V00, SN: FOC1920R0V7


RP/0/RSP0/CPU0:TARDIS(admin)#show inventory  rack
  Rack            Chassis PID        S/N
  ----            -----------        ----------
  0               ASR-9904-AC        FOX1739G94Y
  100             NCS-5002           FOC1920R0V7
```

Monitoring the Satellite Environment

You can use the**show environment temperatures** and **show environment fans**commands in the admin configuration mode to monitor the status of satellite environment.

```
RP/0/RSP0/CPU0:router(admin)# show environment temperatures

R/S/I   Modules Sensor              (deg C)
0/RSP0/*
        host    Inlet0              33.1
        host    Hotspot0            46.9

0/RSP1/*
        host    Inlet0              32.1
        host    Hotspot0            45.9

0/0/*
        host    Inlet0              37.3
        host    Hotspot0            52.3

0/1/*
        spa0    InletTemp           34.0
        spa0    Hotspot             34.5

        spa1    LocalTemp           38.0
        spa1    Chan1Temp           36.0
        spa1    Chan2Temp           39.0
        spa1    Chan3Temp           39.0
        spa1    Chan4Temp           48.0

        host    Inlet0              36.1
        host    Hotspot0            64.0

0/2/*
        host    Inlet0              39.2
        host    Hotspot0            54.6

0/3/*
        host    Inlet0              41.3
        host    Hotspot0            48.5

0/FT0/*
        host    Inlet0              42.3
        host    Hotspot0            36.1

0/FT1/*
        host    Inlet0              40.4
        host    Hotspot0            35.8

SAT100/FT0/*
        host    Hotspot0            53.0


SAT101/FT0/*
        host    Hotspot0            56.0


SAT102/FT0/*
        host    Hotspot0            53.0

RP/0/RSP0/CPU0:router(admin)# show environment fans

Wed Apr  8 17:40:00.313 UTC
```

```
  Fan speed (rpm) and run time (in hours) :
        FAN0    FAN1    FAN2    FAN3    FAN4    FAN5    FAN6    FAN7    FAN8         \
FAN9    FAN10   FAN11   FAN12   FAN13

0/FT0/* (Speed)
        7980    7830    7920    8010    7920    7920    7920    7950    7920         \
7980    7890    8010    8010    7950
0/FT0/* (Run Time)
        N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A     N/A \
    N/A     N/A     N/A     N/A
SAT100/FT0/*
        15000   14117   14117   0
```

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            \
    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 and 9194 for a ring or L2FAB 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 12: Line Loopback

Figure 13: Internal Loopback
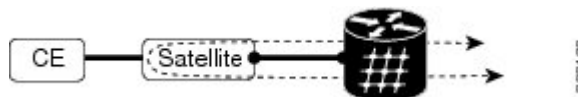
You can specify the type of loopback to be used, as specified in this example:

```
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 topology on
the active and standby hosts:
Active host:
interface GigabitEthernet100/0/0/32
 ipv4 address 1.1.1.1 255.255.255.0
!
```

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 a Satellite nV System in Simple Ring Topology

On Host1

```
redundancy
 iccp
  group 2
   member
```

```
    neighbor 9.9.9.9
   !
   nv satellite
    system-mac dcdc.dcdc.dcdc
   !
  !
 !
nv
 satellite 500
  type asr9000v
  ipv4 address 100.100.1.2
  description sat500
  redundancy
   host-priority 30
  !
  serial-number CAT1603U04Q
 !
 satellite 600
  type asr9000v
  ipv4 address 100.100.1.3
  description sat600
  redundancy
   host-priority 30
  !
  serial-number CAT1603U035
 !
 satellite 700
  type asr9000v
  ipv4 address 100.100.1.4
  description sat700
  redundancy
   host-priority 30
  !
  serial-number CAT1710U03C
 !
 satellite 800
  type asr9000v
  ipv4 address 100.100.1.5
  description sat800
  redundancy
   host-priority 30
  !
  serial-number CAT1651U09N
 !
!


RP/0/RSP0/CPU0:HOST1# show runn | b mpls ldp

mpls ldp
 router-id 8.8.8.8
 address-family ipv4
  neighbor 9.9.9.9 targeted
 !
 interface GigabitEthernet0/1/0/3
 !
!
End

RP/0/RSP0/CPU0:HOST1#show runn interface Gi0/1/0/0

interface GigabitEthernet0/1/0/0
 ipv4 point-to-point
 ipv4 unnumbered Loopback10
```

```
nv
 satellite-fabric-link network
  redundancy
   iccp-group 2
  !
  satellite 500
   remote-ports GigabitEthernet 0/0/0-9
  !
  satellite 600
   remote-ports GigabitEthernet 0/0/0-9
  !
  satellite 700
   remote-ports GigabitEthernet 0/0/0-9
  !
  satellite 800
   remote-ports GigabitEthernet 0/0/0-9
  !
 !
!
```

On Host2

```
nv
 satellite 500
  type asr9000v
  ipv4 address 100.100.1.2
  description sat500
  redundancy
   host-priority 101
  !
  serial-number CAT1603U04Q
 !
 satellite 600
  type asr9000v
  ipv4 address 100.100.1.3
  description sat600
  redundancy
   host-priority 101
  !
  serial-number CAT1603U035
 !
 satellite 700
  type asr9000v
  ipv4 address 100.100.1.4
  description sat700
  redundancy
   host-priority 101
  !
  serial-number CAT1710U03C
 !
 satellite 800
  type asr9000v
  ipv4 address 100.100.1.5
  description sat800
  redundancy
   host-priority 101
  !
  serial-number CAT1651U09N
 !
!
interface Bundle-Ether10
```

```
 bundle wait-while 0
 load-interval 30
 l2transport
 !
!

RP/0/RSP0/CPU0:HOST2# show runn | b mpls ldp

mpls ldp
 router-id 9.9.9.9
 address-family ipv4
  neighbor 8.8.8.8 targeted
 !
 interface GigabitEthernet0/0/0/2
 !
!
End


RP/0/RSP0/CPU0:HOST2# show runn | b redundancy

redundancy
 iccp
  group 2
   member
    neighbor 8.8.8.8
   !
   nv satellite
    system-mac dcdc.dcdc.dcdc
   !
  !



RP/0/RSP0/CPU0:HOST2# show runn interface Gi0/0/0/18

interface GigabitEthernet0/0/0/18
 ipv4 point-to-point
 ipv4 unnumbered Loopback10
 nv
  satellite-fabric-link network
   redundancy
    iccp-group 2
   !
   satellite 500
    remote-ports GigabitEthernet 0/0/0-9
   !
   satellite 600
    remote-ports GigabitEthernet 0/0/0-9
   !
   satellite 700
    remote-ports GigabitEthernet 0/0/0-9
   !
   satellite 800
    remote-ports GigabitEthernet 0/0/0-9
   !
  !
 !
!
```

Configuring a Satellite nV System in Layer 2 Fabric Network Topology

```
interface TenGigE0/0/0/0.102
ipv4 point-to-point
ipv4 unnumbered Loopback0
load-interval 30
encapsulation dot1q 102
nv
  satellite-fabric-link satellite 300
   ethernet cfm
    continuity-check interval 10ms
   !
   redundancy
    iccp-group 10
   !
   remote-ports GigabitEthernet 0/0/0-10
```

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

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 selected platforms using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL: *http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml*

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 **3**

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 91
- Scope and Prerequisites, on page 92
- Multicast Offload Terminology, on page 93
- Overview of Multicast Offload, on page 94
- Restrictions for Multicast Offload, on page 98
- Configuring Satellite nV Multicast Offload, on page 100

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.

Note Refer the *Configuring the Satellite Network Virtualization (nV) System* chapter for more information on Simple Ring and Layer 2 fabric topologies.

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.

Cisco ASR 9000 Series Aggregation Services Router nV System Configuration Guide, Release 6.1.x

91

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 14: Satellite nV System Multicast Offload Illustration
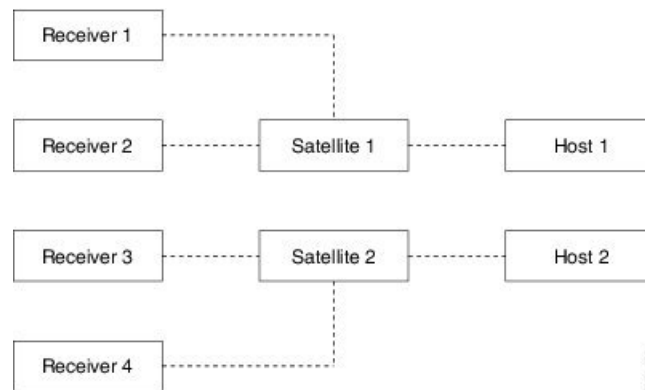
Without offload, Host 1 (assuming that it is the active for both Satellite 1 and Satellite 2) sends 4 copies to receivers 1,2,3 and 4 even if they join the same multicast group, (S1,G1). With offload, Host 1 sends 1 copy and Satellite 1 replicates it twice for receivers 1 and 2 and Satellite 2 replicates 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 a master 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**   There is no support for Satellite nV offload for Layer 3 multicast, IPv6 multicast traffic or MLD snooping protocol.

- **Satellite topologies**: Simple Ring topology with Dual Host, which is a variant of a simple chain with single host, and Hub and spoke topology with a single nV host.

Note   Only these topology variants are supported:

> • 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

Prerequisites:

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

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. This term applies to the simple ring Satellite nV System topology.

- **Backup Multicast Forwarder**: Multicast packets received from Backup Multicast forwarder are not replicated on the satellite for access ports. This term applies to the simple ring Satellite nV System topology.

- **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.

- **Master 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 master OLE is selected per offloaded route (S,G), for each BD and each ring. For Hub and Spoke topology, the master 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.

- 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 15: IGMP Join Programming (Access L2 and Core L2/BVI)
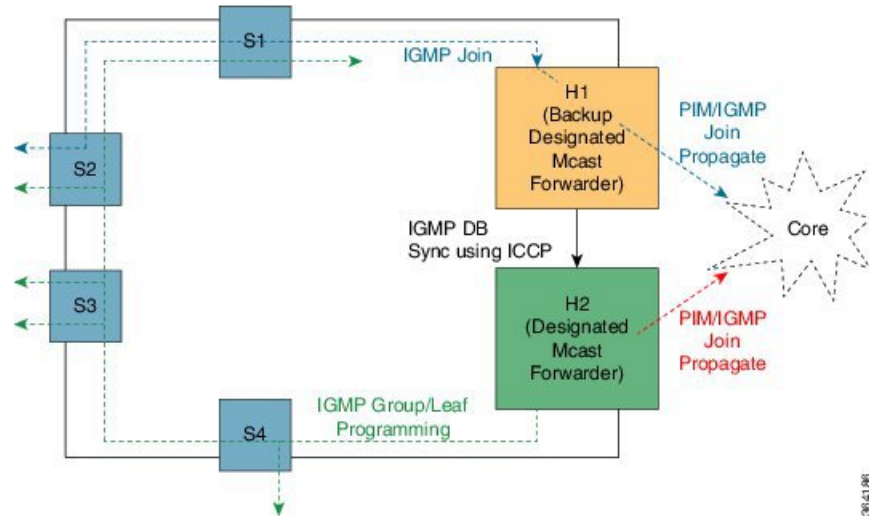
Figure 16: Multicast Offload Data Flow
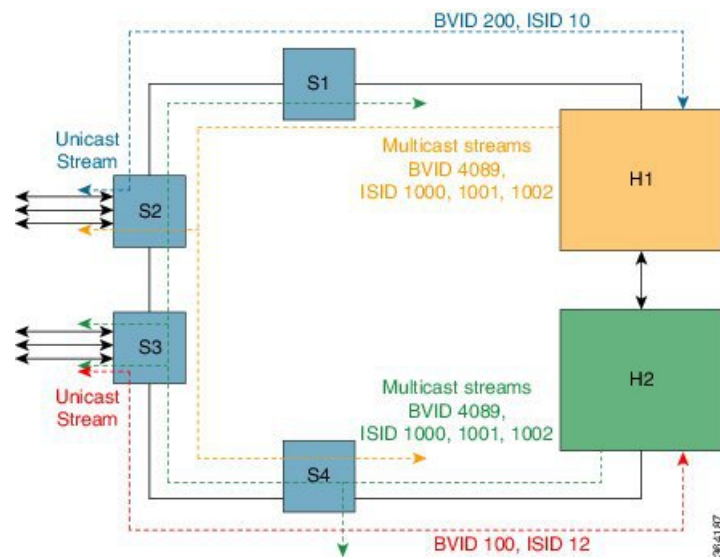
These illustrations show the multicast offload functionality on Satellite nV System that use hub and spoke topologies.

Figure 17: Multicast Offload Data Flow with ICL bundles to different satellites in Hub and Spoke topology
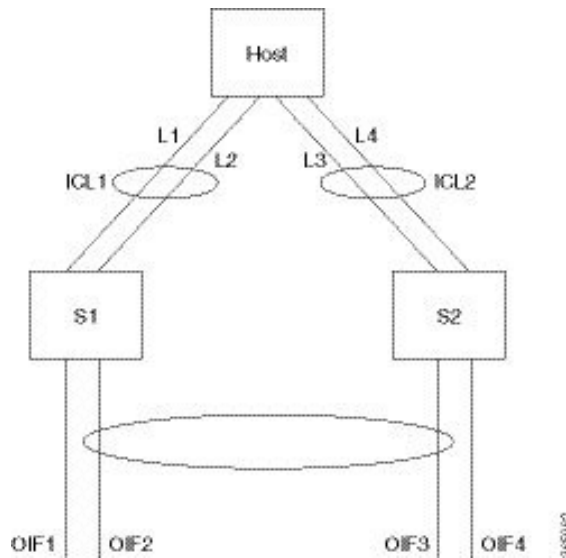
Figure 18: Multicast Offload with ICL bundle from different line cards to a satellite
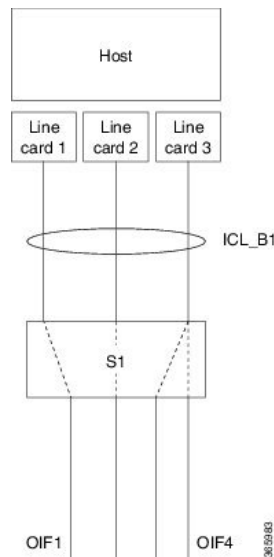
Behavior of Multicast Offload During Satellite Topology Events

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 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 master 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 master 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 master OLE on host. The same applies for any other egress feature configuration on the OLEs.

- 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 Master 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 Master 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.

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:routershow igmp snooping group summary debug

Bridge Domain bg:bd

                                      #Mem  #Inc  #Exc  Annot                        Master
      #Mem   Master OLE
Group            Source       Ver GM  Ports Ports Ports Key           ISID  Ver  OLE XID
      Annot InterfaceName
-----          ------         --- --  ----- ----- ----- ------------  ----- ---- ----------
      ----- -------------
225.0.0.1      {}            V3  EX  2     -     -     0x2-c8        0x3e8 0x4  0xa0000005
   2   Bundle-Ether50
225.0.0.2      {}            V3  EX  2     -     -     0x2-64        0x3e8 0x4  0xa0000005
   1   Bundle-Ether50
                                                       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:routershow 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


                                                                  OLE
Port                     PM Group           Ver GM Source    Exp  Flgs Offl
Ul_ifh     InterfaceName
----                     -- -----           --- -- ------    ---  ---- ----
---------- -------------
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*