Introduction

This module provides an overview of the Carrier Grade IPv6 (CGv6) on Cisco IOS XR software.

The following table lists changes made to the document.

Table 1: Feature History for Implementing CGv6 on ASR 9000 Router

Release	Modification
Release 4.2.0	Initial release of this document. CGv6 applications such as CGN or NAT44 are supported.
Release 4.2.1	These features were introduced: • DS-Lite. • Syslog and Bulk Port Allocation for NAT44 and DS-Lite.
Release 4.2.3	Support for multiple ISM line cards.
Release 4.3.0	These features were introduced: • Stateful NAT64 • Mapping of Address and Port-Translation Mode • High Availability • Destination-Based Logging
Release 4.3.1	These features were introduced: • IPv6 Rapid Deployment • Mapping of Address and Port-Encapsulation Mode • Point-to-Point Tunneling Protocol-Application Level Gateway on NAT44 • Real-Time Streaming Proocol-Application Level Gateway on Stateful NAT64
Release 5.1.1	Support for Virtualized Services Module (VSM) has been introduced in this release.

Release 5.2.0	These features were introduced: • NAT0 Mode • Static Destination NAT • Multiple NetFlow/Syslog Servers • Additional CGN Counters

CGv6 Overview and Benefits

To implement the CGv6, you should understand the following concepts.

CGv6 Overview

Internet Protocol version 4 (IPv4) has reached exhaustion at the international level (IANA). But service providers must maintain and continue to accelerate growth. Billions of new devices such as mobile phones, portable multimedia devices, sensors, and controllers are demanding Internet connectivity at an increasing rate. The Cisco Carrier Grade IPv6 Solution (CGv6) is designed to help address these challenges. With Cisco CGv6, you can:

• Preserve investments in IPv4 infrastructure, assets, and delivery models.

• Prepare for the smooth, incremental transition to IPv6 services that are interoperable with IPv4.

• Prosper through accelerated subscriber, device, and service growth that are enabled by the efficiencies that IPv6 can deliver.

Cisco CGv6 extends the already wide array of IPv6 platforms, solutions, and services. Cisco CGv6 helps you build a bridge to the future of the Internet with IPv6.

Cisco ASR 9000 Series Aggregation Services Router is part of the Cisco CGv6 solution portfolio and therefore different CGv6 solutions or applications are implemented on this platform (specifically on ISM service card). Carrier Grade Network Address Translation (CGN) is a large scale NAT that is capable of providing private IPv4 to public IPv4 address translation in the order of millions of translations to support a large number of subscribers, and at least 10 Gbps full-duplex bandwidth throughput.

Benefits of CGv6

CGv6 offers these benefits.

• Enables service providers to execute orderly transitions to IPv6 through mixed IPv4 and IPv6 networks.

• Provides address family translation but not limited to just translation within one address family.

• Delivers a comprehensive solution suite for IP address management and IPv6 transition.

IPv4 Address Shortage

A fixed-size resource such as the 32-bit public IPv4 address space will run out in a few years. Therefore, the IPv4 address shortage presents a significant and major challenge to all service providers who depend on large blocks of public or private IPv4 addresses for provisioning and managing their customers.

Service providers cannot easily allocate sufficient public IPv4 address space to support new customers that need to access the public IPv4 Internet.

Prerequisites for Implementing the CGv6

The following prerequisites are required to implement CGv6.

- You must be running Cisco IOS XR software Release 4.2.0 and above.

- You must have installed the CGv6 service package, **asr9k-services-p.pie** (to be used with RSP2) or **asr9k-services-px.pie** (to be used with RSP3).

- 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.

Note   All the error conditions result in a syslog message. On observation of Heartbeat failure messages, contact Cisco Technical Support with **show tech-support services cgn** information.

Note   If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Implementation of NAT

This section explains various implementations of NAT. The implementation of NAT over ISM and VSM are explained in the following chapters.

Implementing NAT with ICMP

This section explains how the Network Address Translation (NAT) devices work in conjunction with Internet Control Message Protocol (ICMP).

The implementations of NAT varies in terms of how they handle different traffic.

ICMP Query Session Timeout

RFC 5508 provides ICMP Query Session timeouts. A mapping timeout is maintained by NATs for ICMP queries that traverse them. The ICMP Query Session timeout is the period during which a mapping will stay

active without packets traversing the NATs. The timeouts can be set as either Maximum Round Trip Time (Maximum RTT) or Maximum Segment Lifetime (MSL). For the purpose of constraining the maximum RTT, the Maximum Segment Lifetime (MSL) is considered a guideline to set packet lifetime.

If the ICMP NAT session timeout is set to a very large duration (240 seconds) it can tie up precious NAT resources such as Query mappings and NAT Sessions for the whole duration. Also, if the timeout is set to very low it can result in premature freeing of NAT resources and applications failing to complete gracefully. The ICMP Query session timeout needs to be a balance between the two extremes. A 60-second timeout is a balance between the two extremes.

Implementing NAT with TCP

This section explains the various NAT behaviors that are applicable to TCP connection initiation. The detailed NAT with TCP functionality is defined in RFC 5382.

Address and Port Mapping Behavior

A NAT translates packets for each TCP connection using the mapping. A mapping is dynamically allocated for connections initiated from the internal side, and potentially reused for certain connections later.

Internally Initiated Connections

A TCP connection is initiated by internal endpoints through a NAT by sending SYN packet. All the external IP address and port used for translation for that connection are defined in the mapping.

Generally for the client-server applications where an internal client initiates the connection to an external server, to translate the outbound SYN, the resulting inbound SYN-ACK response mapping is used, the subsequent outbound ACK, and other packets for the connection.

The 3-way handshake corresponds to method of connection initiation.

Externally Initiated Connections

For the first connection that is initiated by an internal endpoint NAT allocates the mapping. For some situations, the NAT policy may allow reusing of this mapping for connection initiated from the external side to the internal endpoint.

Double NAT 444

The Double NAT 444 solution offers the fastest and simplest way to address the IPv4 depletion problem without requiring an upgrade to IPv6 anywhere in the network. Service providers can continue offering new IPv4 customers access to the public IPv4 Internet by using private IPv4 address blocks, if the service provider is large enough; However, they need to have an overlapping RFC 1918 address space, which forces the service provider to partition their network management systems and creates complexity with access control lists (ACL).

Double NAT 444 uses the edge NAT and CGN to hold the translation state for each session. For example, both NATs must hold 100 entries in their respective translation tables if all the hosts in the residence of a subscriber have 100 connections to hosts on the Internet). There is no easy way for a private IPv4 host to communicate with the CGN to learn its public IP address and port information or to configure a static incoming port forwarding.

Address Family Translation

The IPv6-only to IPv4-only protocol is referred to as address family translation (AFT). The AFT translates the IP address from one address family into another address family. For example, IPv6 to IPv4 translation is called NAT 64 or IPv4 to IPv6 translation is called NAT 46.

Jumbo Frame Support

Jumbo frames are frames that are larger than the standard Ethernet frame size, which is 1518 bytes. The definition of frame size is vendor-dependent, and are not part of the IEEE standard.

The Integrated Services Module (ISM) and Virtualized Services Module (VSM) both support Jumbo Frames.

To enable Jumbo Frame support, configure the Maximum Transmission Unit (MTU) value of both the ingress and egress interfaces. The default MTU value is 1512 bytes and the maximum value is 9216 bytes.

Jumbo Frame Support