



Configure IP-in-IP Tunnels

This chapter provides conceptual and configuration information for IP-in-IP tunnels.

Table 1: Feature History for Configure Tunnels

| | |
|----------------|--|
| Release 7.0.11 | This feature was introduced. |
| Release 7.0.14 | Support for the following feature was introduced in Configure Tunnels: <ul style="list-style-type: none">• Extended ACL must match on the outer header for IP-in-IP Decapsulation. |

- [IP-in-IP Decapsulation, on page 1](#)
- [ECMP Hashing Support for Load Balancing, on page 8](#)

IP-in-IP Decapsulation

Table 2: Feature History Table

| Feature Name | Release Information | Description |
|---|---------------------|--|
| IP header decapsulation for multistack packet in IP-in-IP tunnels | Release 7.3.5 | <p>Your routers are now enhanced to identify IP packets with multistack IP header encapsulations and perform multiple IP header decapsulations on these packets. With this ability, you can use IP packets to travel from the source endpoint to the destination endpoint and then back to the source in a IP-in-IP tunnel This feature facilitates identification of IP packets streamlined fault identification and isolation processes in your IP-in-IP network.</p> <p>This feature is enabled by default and doesn't require any configuration changes.</p> |

IP-in-IP encapsulation involves the insertion of an outer IP header over the existing IP header. The source and destination address in the outer IP header point to the endpoints of the IP-in-IP tunnel. The stack of IP headers is used to direct the packet over a predetermined path to the destination, provided the network administrator knows the loopback addresses of the routers transporting the packet. This tunneling mechanism can be used for determining availability and latency for most network architectures. It is to be noted that the entire path from source to the destination does not have to be included in the headers, but a segment of the network can be chosen for directing the packets.

In IP-in-IP encapsulation and decapsulation has two types of packets. The original IP packets that are encapsulated are called Inner packets and the IP header stack added while encapsulation are called the Outer packets.




Note The router only supports decapsulation and no encapsulation. Encapsulation is done by remote routers.

Starting with Cisco IOS XR Software Release 7.3.5, the routers are enhanced to handle multi-stack for-us packets within IP-in-IP tunnels efficiently. The for-us packets refers to the IP packets with the the source and destination terminating on the same router and they are multi-stack when they undergo multiple layers of IP header encapsulation at ingress, loopback and egress interfaces. With this improvement, your routers can identify IP packets with multi-stack tunnelling IP header encapsulations and perform multiple IP header terminations on these packets. This capability enables you to configure IP packets to travel from the source endpoint to the destination endpoint and then back to the source, facilitating streamlined fault identification and isolation processes.


The following topology describes a use case where IP-in-IP encapsulation and decapsulation are used for different segments of the network from source to destination. The IP-in-IP tunnel consists of multiple routers that are used to decapsulate and direct the packet through the data center fabric network.

Figure 1: IP-in-IP Decapsulation Through a Data Center Network



The following illustration shows how the stacked IPv4 headers are decapsulated as they traverse through the decapsulating routers.

Figure 2: IP Header Decapsulation



Stacked IP Header in an Encapsulated Packet

The encapsulated packet has an outer IPv4 header that is stacked over the original IPv4 header, as shown in the following illustration.

Figure 3: Encapsulated Packet

| | |
|-----------------------|--------------------|
| [-] Frame | |
| [-] EthernetII | |
| Preamble (hex) | fb55555555555d5 |
| Destination MAC | 62:19:88:64:E2:68 |
| Source MAC | 00:10:94:00:00:02 |
| EtherType (hex) | <auto> Internet IP |
| [-] IPv4 Header | |
| Version (int) | <auto> 4 |
| Header length (int) | <auto> 5 |
| ToS/DiffServ | tos (0x00) |
| Total length (int) | <auto> calculated |
| Identification (int) | 0 |
| [-] Control Flags | |
| Reserved (bit) | 0 |
| DF Bit (bit) | 0 |
| MF Bit (bit) | 0 |
| Fragment Offset (int) | 0 |
| Time to live (int) | 255 |
| Protocol (int) | <auto> IP |
| Checksum (int) | <auto> 33492 |
| Source | 192.xx.xx.xx |
| Destination | 127.0.0.1 |
| Header Options | |
| Gateway | 192.0.2.10 |
| [-] IPv4 Header | |
| Version (int) | <auto> 4 |
| Header length (int) | <auto> 5 |
| ToS/DiffServ | tos (0x00) |
| Total length (int) | <auto> calculated |
| Identification (int) | 0 |
| [-] Control Flags | |
| Reserved (bit) | 0 |

385413

Configuration

You can use the following sample configuration in the routers to decapsulate the packet as it traverses the IP-in-IP tunnel:

```

Router(config)# interface loopback 0
Router(config-if)# ipv4 address 127.0.0.1/32
Router(config-if)# no shutdown
Router(config-if)# interface tunnel-ip 10
Router(config-if)# ipv4 unnumbered loopback 1
Router(config-if)# tunnel mode ipv4 decap
Router(config-if)# tunnel source loopback 0

```

- **tunnel-ip**: configures an IP-in-IP tunnel interface.
- **ipv4 unnumbered loopback address**: enables ipv4 packet processing without an explicit address, except for loopback address.
- **tunnel mode ipv4 decap**: enables IP-in-IP decapsulation.
- **tunnel source**: indicates the source address for the IP-in-IP decap tunnel with respect to the router interface.



Note You can configure the tunnel destination only if you want to decapsulate packets from a particular destination. If no tunnel destination is configured, then all the ip-in-ip ingress packets on the configured interface are decapsulated.

Running Configuration

```

Router# show running-config interface tunnel-ip 10
...
interface tunnel-ip 10
ipv4 unnumbered loopback 1
tunnel mode ipv4 decap

```

Extended ACL to Match the Outer Header for IP-in-IP Decapsulation

Starting with Cisco IOS XR Software Release 7.0.14, extended ACL has to match on the outer header for IP-in-IP Decapsulation. Extended ACL support reduces mirrored traffic throughput. This match is based only on the IPv4 protocol, and extended ACL is applied to the received outermost IP header, even if the outer header is locally terminated.

Sample configuration:

```

Router#show running-config interface bundle-Ether 50.5
Tue May 26 12:11:49.017 UTC
interface Bundle-Ether50.5
ipv4 address 101.1.5.1 255.255.255.0
encapsulation dot1q 5
ipv4 access-group ExtACL_IPinIP ingress
ipv4 access-group any_dscpegg egress
!

Router#show access-lists ipv4 ExtACL_IPinIP hardware ingress location$
Tue May 26 12:11:55.940 UTC
ipv4 access-list ExtACL_IPinIP
10 permit ipv4 192.168.0.0 0.0.255.255 any ttl gt 150
11 deny ipv4 172.16.0.0 0.0.255.255 any fragments
12 permit ipv4 any any

```

Configure Tunnel Destination with an Object Group

Table 3: Feature History Table

| Feature Name | Release Information | Description |
|---|---------------------|--|
| Configure Tunnel Destination with an Object Group | Release 7.5.4 | <p>You can now assign an object group as the destination for an IP-in-IP decapsulation tunnel. With this functionality, you could configure an IPv4 or IPv6 object group consisting of multiple IPv4 or IPv6 addresses as the destination for the tunnel instead of a single IPv4 or IPv6 address. Using an object group instead of a singular IP address. This helps reduce the configuration complexity in the router by replacing the multiple tunnels with one destination with a single decapsulation tunnel that supports a diverse range of destinations</p> <p>The feature introduces these changes:</p> <ul style="list-style-type: none">• CLI: New tunnel destination command.• YANG Data Model: New object-group option supported in Cisco-IOS-XR-um-iftunnel-cfg.yang Cisco native model (see GitHub). |

In IP-in-IP Decapsulation, the router accepts a packet on a tunneled interface only when the tunnel IP address matches the source IP address of the incoming packets. With this implementation, the user needs to configure separate interface tunnels for each IP address that the router needs to receive the traffic packets. This limitation often leads to configuration overload on the router.

You can eliminate the configuration overload on the router by assigning an object group as the tunnel destination for IPv4 and IPv6 traffic types. That is, the router matches the source IP address of the incoming packet against the object group available as the tunnel destination. The decapsulation tunnel accepts the incoming traffic packets when there's a match between the packet source and the object group. Otherwise, the router drops the packets.

Restrictions

The following restrictions are applicable to the tunnel destination with an object group feature:

- GRE tunnels don't support configuring object groups as the tunnel destination.

- The router supports configuring tunnel destination with an object group only when the tunnel source is tunnel source direct.
- You can configure the object group as tunnel destination only on default VRF.
- Configuring object groups as the tunnel destination isn't applicable to tunnel encapsulation.
- Subinterfaces don't support configuring object groups as the tunnel destination.
- Configuring object groups as the tunnel destination feature is mutually exclusive with ACL and QoS features.
- The tunnel destination feature supports only IPv4 and IPv6 object groups.
- The router does not support changing tunnel configuration after its creation. Configure the tunnel source direct and tunnel destination with an object group while creating the tunnel only.

Prerequisites

- Define an object group including the network elements for the tunnel destination.
- Enable the tunnel source direct feature. For more information, see decapsulation using tunnel source direct.

Configuration Example

This section provides an example for configuring the tunnel destination with an object group:

Configuration

IPv4:

```
Router# configure
/* Configure the IPv4 object group */
Router(config)# object-group network ipv4 Test_IPv4
Router(config-object-group-ipv4)# 192.0.2.0/24
Router(config-object-group-ipv4)# 198.51.100.0/24
Router(config-object-group-ipv4)# 203.0.113.0/24
Router(config-object-group-ipv4)# commit
Router(config-object-group-ipv4)# exit

/* Enters the tunnel configuration mode */
Router(config)# interface tunnel TestIPv4

/* Configures the tunnel mode */
Router(config-if)# tunnel mode ipv4 decap

/* Configures the tunnel to accept all packets with destination address matching the IP
addresses on the router */
Router(config-if)# tunnel source direct

/* Configures the tunnel to accept all packets with destination address that are in the
specified object group */
Router(config-if)# tunnel destination object-group ipv4 Test_IPv4

Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit
```

IPv6:

```

Router# configure
/* Configure the IPv6 object group */
Router(config)# object-group network ipv6 Test_IPv6
Router(config-object-group-ipv6)# 2001:DB8::/32
Router(config-object-group-ipv6)# 2001:DB8::/48
Router(config-object-group-ipv6)# commit
Router(config-object-group-ipv6)# exit

/* Enters the tunnel configuration mode */
Router(config)# interface tunnel TestIPv6

/* Configures the tunnel mode */
Router(config-if)# tunnel mode ipv6 decap

/* Configures the tunnel to accept all packets with destination address matching the IP
addresses on the router */
Router(config-if)# tunnel source direct

/* Configures the tunnel to accept all packets with destination address that are in the
specified object group */
Router(config-if)# tunnel destination object-group ipv6 Test_IPv6

Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit

```

Running Configuration

```

Router# show running config object-group
object-group network ipv4 Test_IPv4
192.0.2.0/24
198.51.100.0/24
203.0.113.0/24
!
object-group network ipv6 Test_IPv6
2001:DB8::/32
2001:DB8::/48
!

Router# show interface tunnel TestIPv4
interface TunnelTestIPv4
  tunnel mode ipv4 decap
  tunnel source direct
  tunnel destination object-group ipv4 Test_IPv4
  no shutdown
!

Router# show interface tunnel TestIPv6
interface TunnelTestIPv6
  tunnel mode ipv6 decap
  tunnel source direct
  tunnel destination object-group ipv6 Test_IPv6
  no shutdown
!

```

Verification

```

Router# show tunnel ip ea database

----- node0_0_CPU0 -----
tunnel ifhandle 0x80022cc
tunnel source 161.115.1.2
tunnel destination address group Test_IPv4

```

```

tunnel transport vrf table id 0xe0000000
tunnel mode gre ipv4, encap
tunnel bandwidth 100 kbps
tunnel platform id 0x0
tunnel flags 0x40003400
IntfStateUp
BcStateUp
Ipv4Caps
Encap
tunnel mtu 1500
tunnel tos 0
tunnel ttl 255
tunnel adjacency flags 0x1
tunnel o/p interface handle 0x0
tunnel key 0x0, entropy length 0 (mask 0xffffffff)
tunnel QT next 0x0
tunnel platform data (nil)
Platform:
Handle: (nil)
Decap ID: 0
Decap RIF: 0
Decap Recycle Encap ID: 0x00000000
Encap RIF: 0
Encap Recycle Encap ID: 0x00000000
Encap IPv4 Encap ID: 0x4001381b
Encap IPv6 Encap ID: 0x00000000
Encap MPLS Encap ID: 0x00000000
DecFEC DecRcyLIF DecStatsId EncRcyLIF

```

ECMP Hashing Support for Load Balancing

The system inherently supports the n-tuple hash algorithm. The first inner header in the n-tuple hashing includes the source port and the destination port of UDP / TCP protocol headers.

The load balancing performs these functions:

- Incoming data traffic is distributed over multiple equal-cost connections.
- Incoming data traffic is distributed over multiple equal-cost connections member links within a bundle interface.
- Layer 2 bundle and Layer 3 (network layer) load-balancing decisions are taken on IPv4, and IPv6. If it is an IPv4 or an IPv6 payload, then an n-tuple hashing is done.
- An n-tuple hash algorithm provides more granular load balancing and used for load balancing over multiple equal-cost Layer 3 (network layer) paths. The Layer 3 (network layer) path is on a physical interface or on a bundle interface.
- The n-tuple load-balance hash calculation contains:
 - Source IP address
 - Destination IP address
 - IP Protocol type
 - Router ID
 - Source port

- Destination port
- Input interface
- Flow-label (for IPv6 only)

