



## Configure Generic UDP Encapsulation

Read this section to get an overview and know how to configure the Generic UDP Encapsulation.

**Table 1: Feature History Table**

Feature Name	Release Information	Feature Description
Generic UDP Encapsulation	Release 25.4.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on: <ul style="list-style-type: none"><li>• 8011-12G12X4Y-A</li><li>• 8011-12G12X4Y-D</li></ul>
Outer IP Header-Driven Hash Computation for Incoming GUE Packets	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: <ul style="list-style-type: none"><li>• 8712-MOD-M</li><li>• 8011-4G24Y4H-I</li></ul>

Feature Name	Release Information	Feature Description
Outer IP Header-Driven Hash Computation for Incoming GUE Packets	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*).</p> <p>*This feature is now supported on:</p> <ul style="list-style-type: none"> <li>• 8212-32FH-M</li> <li>• 8212-48FH-M</li> <li>• 8711-32FH-M</li> <li>• 88-LC1-12TH24FH-E</li> <li>• 88-LC1-36EH</li> <li>• 88-LC1-52Y8H-EM</li> </ul>
Outer IP Header-Driven Hash Computation for Incoming GUE Packets	Release 7.11.1	<p>We now offer you the flexibility of using only the outer IP header to calculate the hashing for incoming Generic UDP Encapsulation (GUE) packets. On enabling this feature, only the outer IP source and destination addresses are used for hashing calculations. The inner IP addresses are not considered, providing a simpler method of distribution. Previously, both inner IP and outer IP headers were used for ECMP hashing the incoming GUE packets.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">hw-module profile gue underlay-hash</a></li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• New XPath for <a href="#">Cisco-IOS-XR-npu-hw-profile-cfg.yang</a> (see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> <p>The command is supported on Q200-based ASICs.</p>

Feature Name	Release Information	Feature Description
Generic UDP Encapsulation	Release 7.3.1	<p>This feature enables you to add an additional header to packets to identify or authenticate the data using UDP. Encapsulating packets in UDP leverages the use of the UDP source port to provide entropy to Equal Cost Multi-Path (ECMP) hashing. It provides significant performance benefits for load-balancing.</p> <p>This command is introduced for this feature:</p> <p><b>decapsulate gue</b></p>

- [Understand UDP encapsulation, on page 3](#)
- [Generic UDP encapsulation, on page 3](#)
- [Flexible Assignment of UDP Port Numbers for Decapsulation, on page 13](#)
- [GUEv1 static tunnel configuration over IPv4 networks, on page 21](#)

## Understand UDP encapsulation

UDP encapsulation is a technique of adding network headers to packets and then encapsulating the packets within the User Datagram Protocol (UDP).

Encapsulating packets using UDP facilitates efficient transport across networks. By leveraging Receive Side Scaling (RSS) and Equal Cost Multipath (ECMP) routing, UDP provides significant performance benefits for load-balancing. The use of the UDP source port provides entropy to ECMP hashing and provides the ability to use the IP source or destination, and the L4 Port for load-balancing entropy.

Traditional mechanisms like Generic Routing Encapsulation (GRE) can handle only the outer Source IP address and parts of the destination address. They may not provide sufficient load balancing entropy.

## Generic UDP encapsulation

Generic UDP encapsulation (GUE) is a UDP-based network encapsulation protocol that encapsulates IPv4 and IPv6 packets. GUE provides native UDP encapsulation and defines an additional header, which helps to determine the payload carried by the IP packet. The additional header can include items such as .

- a virtual networking identifier
- security data for validating or authenticating the GUE header, and
- congestion control data.

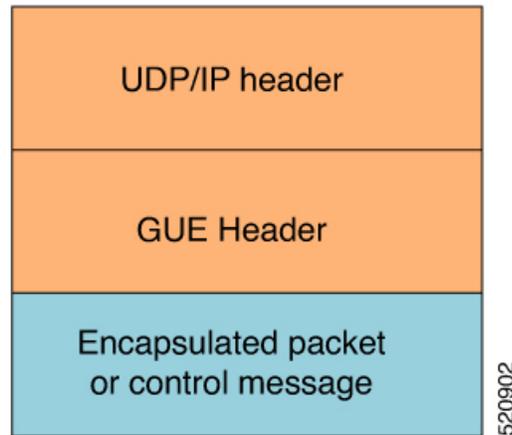
Table 2: Feature History Table

Feature Name	Release Information	Feature Description
Single UDP port for IPv4 and IPv6 packets	Release 25.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200])</p> <p>This feature enhances memory and network performance by using a single UDP port for both IPv4 and IPv6 packets in GUE variant 1 UDP encapsulation and decapsulation.</p> <p>The default UDP port for IPv4 and IPv6 is 6080.</p> <p>As part of this enhancement, the <b>hw-module profile gue udp-dest-port</b> command is deprecated.</p> <p>This feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> <li>• <b>nve overlay-encap guev1,</b></li> <li>• <b>show cef global udp-ports gue-v1</b></li> </ul>
Generic UDP Decapsulation for IPv6 Traffic	Release 25.1.1	<p>Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*)</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> <li>• 8712-MOD-M</li> <li>• 8011-4G24Y4H-I</li> </ul>
Generic UDP Decapsulation for IPv6 Traffic	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>*This feature is now supported on:</p> <ul style="list-style-type: none"> <li>• 8212-48FH-M</li> <li>• 8711-32FH-M</li> <li>• 88-LC1-12TH24FH-E</li> <li>• 88-LC1-36EH</li> <li>• 88-LC1-52Y8H-EM</li> </ul>
Generic UDP Decapsulation for IPv6 Traffic	Release 24.1.1	<p>We've now made it possible for you to decapsulate GUEv6 packets by adding an additional header to packets that identifies or authenticates the data by using User Datagram Protocol (UDP).</p> <p>In GUE, the payload is encapsulated in an IP packet that can be an IPv6 carrier. The UDP header is added to provide extra hashing parameters, and optional payload demultiplexing. At the decapsulation node, the Carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload.</p>

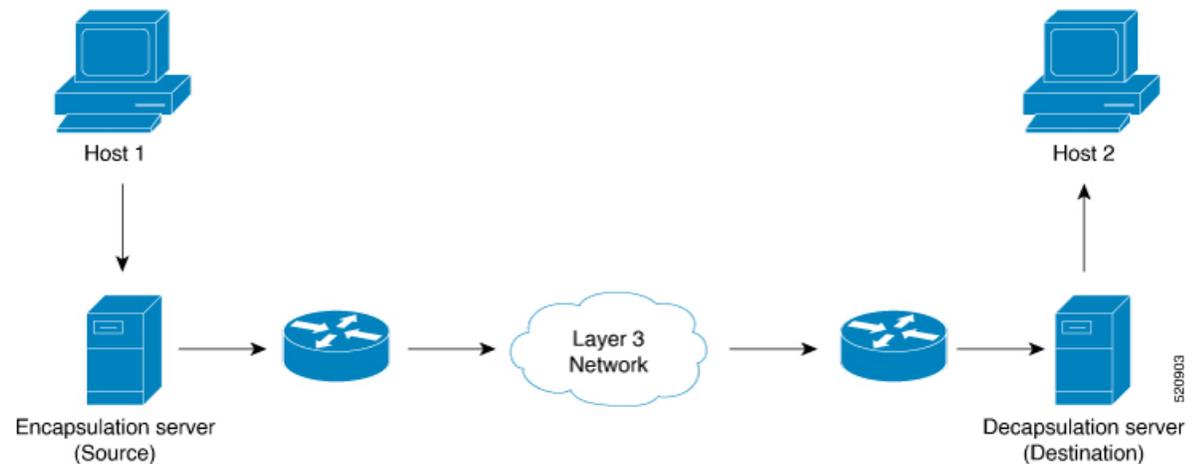
In GUE, the payload is encapsulated in an IP packet that can be IPv4 or IPv6 Carrier. The UDP header is added to provide extra hashing parameters, and optional payload demultiplexing. At the decapsulation node, the Carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload.

A GUE packet has the general format:

**Figure 1: GUE Packet Format**



For example, in this scenario, if the data stream is sent from Host 1 to Host 2. The server acts as a GUE encapsulator that sends the packets from Host 1. The server, on the other end receiving the data, validates the data for the valid carrier IP and UDP header and decapsulates the data.



GUE has various variants, but variant 1 of GUE allows direct encapsulation of IPv4 and IPv6 in UDP. This technique saves encapsulation overhead on links for the use of IP encapsulation, and also need not allocate a separate UDP port number for IP-over-UDP encapsulation.

Variant 1 has no GUE header, but a UDP packet carries an IP packet. The first two bits of the UDP payload is the GUE variant field and match with the first two bits of the version number in the IP header.

Starting from Cisco IOS XR Release 25.4.1, you can use a single UDP port for both IPv4 and IPv6 packets in GUE encapsulation and decapsulation. The default UDP port for IPv4 and IPv6 is 6080. You can modify the default UDP ports using **`nve overlay-encap guev1 udp-port destination`**.



---

**Note** The **hw-module profile gue udp-dest-port** command used to configure the UDP ports individually for IPv4, IPv6, and MPLS is deprecated from this release.

---

### Benefits of using GUE

- Allows direct encapsulation of payloads, such as IPv4 and IPv6 in the UDP packet.
  - You can use UDP port for demultiplexing payloads.
  - You can use a single UDP port, allowing systems to employ parsing models to identify payloads.
- Leverages the UDP header for entropy labels by encoding a tuple-based source port.
- Leverages source IP addresses for load-balance encoding. The destination too could be terminated based on a subnet providing additional bits for entropy.
- Avoids special handling for transit nodes because they only see an IP-UDP packet with some payload.
- Eases implementation of UDP tunneling with GUE. This is because of the direct encapsulation method of the payloads into UDP.

### Benefits of single UDP port for IPv4 and IPv6 packets

- Uses single class map per IP header.
- Allocates one counter per class map.
- Uses simpler logic to reduce duplication of end-to-end objects.

## Usage guidelines for GUE

- Receives IPv4 packets with the defined GUE port of 6080.
- Receives MPLS packets with the UDPoMPLS port of 6635.

## Restrictions

- Supports Generic UDP Decapsulation for only variant 1.
- Range of source or destination ports is not supported.
- Range, Source, or Destination addresses are not supported, but subnet mask entries are allowed.
- Terminating GRE after GUE or GUE after GRE is not supported.
- Terminating a label such as a VPN Deaggregation after GUE termination is not supported.
- Slow path support is not supported. To resolve the inner IP Adjacency, use the **cef proactive-arp-nd enable** command.
- Running the **clear all** command does not clear the interface of all its existing configurations.

- GUE for IPv6 traffic is only supported on the Cisco 8202-32FH-M router, and 88-LC0-36FH-M and 88-LC0-36FH line cards.
- GUE IPv6 decapsulation is only supported on Layer 3 ports.
- GUE IPv6 is not supported over BVI interfaces.



**Note** To use only outer IP header (L3 and L4) for calculating the hashing for incoming GUE packets, use the **hw-module profile gue underlay-hash enable** command. Otherwise, by default, both outer IP header (L3 and L4) and inner IP header (L3 and L4) are considered for calculating the hashing for incoming GUE packets. The **hw-module profile gue underlay-hash enable** command is currently not supported on the P100-based and Q100-based ASICs.

## Configure GUE IPv4

Perform the steps to configure GUE IPv4.

### Procedure

**Step 1** Configure separate GUE decapsulation tunnel UDP destination port numbers for IPv4 using the **nve overlay-encap guev1** command.

#### Example:

```
Router# configure
Router (config-if)# nve overlay-encapsulation guev1
Router (config-nve-encap-guev1)# udp-port destination ipv4 8000
Router (config-nve-encap-guev1)# commit
```

**Step 2** Configure a traffic class: Create a traffic class and specify various criteria for classifying packets using the match commands, and an instruction on how to evaluate these match commands.

#### Example:

```
Router# configure
Router (config)# class-map type traffic match-all gue-v4-udp
Router (config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router (config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router (config-cmap)# match protocol udp
Router (config-cmap)# end-class-map
Router (config)# commit
```

**Step 3** Configure a policy map: Define a policy map and associate the traffic class with the traffic policy.

#### Example:

```
Router (config)# policy-map type pbr magic-decap

Router (config-pmap)# class type traffic udp-v4
Router (config-pmap-c)# decapsulate gue variant 1
Router (config-pmap-c)# exit
```

```

Router(config-pmap)# class type traffic udp-v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-mpls1
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit

```

**Step 4** Apply the policy for each VRF, and apply this policy on all the interfaces that are part of the VRF.

**Example:**

```

Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv4 policy type pbr input magic-decap
Router(config-vrf-policy)# commit

```

**Step 5** Execute this command to view the set of counter values accumulated for the packets that match the class-map:

**Example:**

```

Router# show policy-map type pbr addr-family ipv4 statistics

```

```

VRF Name:      default
Policy-Name:   pmap
Policy Type:   pbr
Addr Family:   IPv4

Class:  cmap-loop1
  Classification statistics (packets/bytes)
  Matched      :           0/0
  Transmitted statistics (packets/bytes)
  Total Transmitted :           0/0

Class:  cmap-loop6
  Classification statistics (packets/bytes)
  Matched      :           0/0
  Transmitted statistics (packets/bytes)
  Total Transmitted :           0/0

Class:  cmap-loop2
  Classification statistics (packets/bytes)
  Matched      :           0/0
  Transmitted statistics (packets/bytes)
  Total Transmitted :           0/0

Class:  cmap-loop3
  Classification statistics (packets/bytes)
  Matched      : 198325306/17849277540
  Transmitted statistics (packets/bytes)
  Total Transmitted : 198325306/17849277540

Class:  cmap-loop4
  Classification statistics (packets/bytes)
  Matched      :           0/0
  Transmitted statistics (packets/bytes)
  Total Transmitted :           0/0

```

**Step 6** Execute this command to view the default configuration if the GUE variant 1 encapsulation is not configured for IPv4 packets.

**Example:**

```
Router# show nve global
Fri Jul 11 16:11:10.172 UTC
NVE Global details
  VNI Scope Local : No
  VxLAN l3vni bring up mode: V1
  GUEv1 UDP Destination Port (IPv4): 6080
  GUEv1 UDP Destination Port (IPv6): 6615
  GUEv1 UDP Destination Port (mpls): 6635
  Count of NVE interfaces with mpls-udp encap: 0
  Global system mac: 0033.3ebe.8f00
```

**Step 7** Execute the **show cef global udp-ports gue-v1** command to view the UDP ports for IPv4, IPV6, and MPLS packets in GUE variant 1 encapsulation.

**Example:**

```
Router# show cef global udp-ports gue-v1
Router# show cef global udp-ports gue-v1
Fri Jul 11 16:11:05.391 UTC
UDP ports for gue-v1 [0x309c3a70f8]
  IPv4: 6080 (default)
  IPv6: 6080 (default)
  MPLS: 6635 (default)
Policy update time: Not Yet Recorded
Platform update time: Not Yet Recorded
```

**Step 8** Execute the **clear vrf** command to clear the policy-map counters for each class-map rule.

**Example:**

```
Router# clear vrf default address-family ipv4 statistics
```

## Configure GUE IPv6

Perform the steps to configure GUE IPv6.

### Procedure

**Step 1** Configure separate UDP port numbers for IPv4, IPv6, and MPLS using the **nve overlay-encap guev1** command.

**Example:**

```
Router# configure
Router (config-if)# nve overlay-encapsulation guev1
Router (config-nve-encap-guev1)# udp-port destination ipv6 9000
Router (config-nve-encap-guev1)# commit
```

**Step 2** Configure a traffic class.

**Example:**

```
Router# configure
```

```
Router(config)# class-map type traffic match-all gue-v6-udp
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv6 11:1:1::1/128
Router(config-cmap)# end-class-map
Router(config)# commit
```

**Step 3** Define a policy map, and associate the traffic class with the traffic policy.

**Example:**

```
Router(config)# policy-map type pbr guev6_decap
Router(config-pmap)# class type traffic guev6_v4
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic guev6_v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic guev6_mpls
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit
```

**Step 4** Apply the policy for each VRF.

**Example:**

```
Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv6 policy type pbr input guev6_decap
Router(config-vrf-policy)# commit
```

**Step 5** Execute this command to view the set of counter values accumulated for the packets that match the class-map.

**Example:**

```
Router# show policy-map type pbr vrf default addr-family ipv6 statistics
```

```
VRF Name:          default
Policy-Name:       guev6_decap
Policy Type:       pbr
Addr Family:       IPv6

Class:      guev6_v4
  Classification statistics      (packets/bytes)
  Matched          :              190/24320
  Transmitted statistics      (packets/bytes)
  Total Transmitted  :              190/24320

Class:      guev6_v6
  Classification statistics      (packets/bytes)
  Matched          :              175/24150
  Transmitted statistics      (packets/bytes)
  Total Transmitted  :              175/24150

Class:      guev6_mpls
  Classification statistics      (packets/bytes)
```

```

    Matched          :          160/22720
Transmitted statistics (packets/bytes)
    Total Transmitted :          160/22720

Class:      class-default
Classification statistics (packets/bytes)
    Matched          :          0/0
Transmitted statistics (packets/bytes)
    Total Transmitted :          0/0

```

**Step 6** Execute this command to view the default configuration if the GUE variant 1 encapsulation is not configured for IPv4 packets.

**Example:**

```

Router# show nve global
Fri Jul 11 16:11:10.172 UTC
NVE Global details
  VNI Scope Local : No
  VxLAN l3vni bring up mode: V1
  GUEv1 UDP Destination Port (IPv4): 6080
  GUEv1 UDP Destination Port (IPv6): 6615
  GUEv1 UDP Destination Port (mpls): 6635
  Count of NVE interfaces with mpls-udp encap: 0
  Global system mac: 0033.3ebe.8f00

```

**Step 7** Execute the `show cef global udp-ports gue-v1` command to view the UDP ports for IPv4, IPV6, and MPLS packets in GUE variant 1 encapsulation.

**Example:**

```

Router# show cef global udp-ports gue-v1
Router# show cef global udp-ports gue-v1
Fri Jul 11 16:11:05.391 UTC
UDP ports for gue-v1 [0x309c3a70f8]
  IPv4: 6080 (default)
  IPv6: 6080 (default)
  MPLS: 6635 (default)
Policy update time: Not Yet Recorded
Platform update time: Not Yet Recorded

```

## Outer IP Header-Driven Hash Computation for Incoming GUE Packets

When multiple paths with the same cost are available for forwarding traffic, ECMP hashing is used to determine the path to select for each packet. Each packet that needs to be forwarded is processed using a hashing algorithm. The hashing algorithm considers specific packet fields such as source IP, destination IP, source port, and destination port, and generates a hash value. The generated hash value is then mapped to one of the available paths. The selected path is used to forward the packet to its destination. The goal is to distribute the traffic evenly across the available paths to prevent congestion and utilize the network resources efficiently.

Now you can use only the outer IP header (L3 and L4) for calculating the hash value for incoming GUE packets and completely ignore the usage of the inner IP header. This functionality is configurable using the CLI command `hw-module profile gue underlay-hash`. This is supported for both GUE termination (decapsulation) and GUE transit (pass-through) nodes. By default, the feature is disabled; that is, both outer IP header (L3 and L4) and inner IP header (L3 and L4) are used for calculating the hashing for GUE packets.

## Benefits

- **Load Balancing Efficiency:** By hashing only on the outer IP and L4 information, the packets with the same source and destination IP addresses and L4 ports consistently follow the same path in a load-balanced environment. This helps maintain session affinity or stickiness, as the inner IP addresses or L4 port numbers may change dynamically within the encapsulated packets.
- **Network Security:** Ignoring the inner IP helps preserve privacy and confidentiality within the encapsulated packets. By focusing on the outer IP and L4 headers, the network device does not have visibility into the inner IP addressing scheme or the specific content encapsulated within the packet, which enhances security.
- **Network Scalability:** Ignoring the inner IP reduces the complexity and overhead of packet processing, improving overall network performance and scalability, especially in high-throughput environments.

## Configure Outer IP Header-Driven Hash Computation for Incoming GUE Packets

This section describes how to configure hashing with only outer IP for GUE packets.

### Configuration Example

Use the following configuration to enable hashing with only outer IP for GUE packets:

```
Router# configure
Router# hw-module profile gue underlay-hash enable
Router# commit
```

### Running Configuration

```
RP/0/RP0/CPU0:ios(config)#show running-config
hw-module profile gue underlay-hash enable
end
```

### Verification

Following is the show command output before enabling hashing with only outer IP for GUE packets.

```
RP/0/RP0/CPU0:ios#show dpa objects sys location 0/RP0/CPU0 | include gue
uint32_t gue_ipv4_port => 0
uint32_t gue_ipv6_port => 0
uint32_t gue_mpls_port => 0
ofa_bool_t gue_underlay_hash => FALSE
```

Following is the show command output after enabling hashing with only outer IP for GUE packets.

```
RP/0/RP0/CPU0:ios#show dpa objects sys location 0/RP0/CPU0 | include gue
uint32_t gue_ipv4_port => 0
uint32_t gue_ipv6_port => 0
uint32_t gue_mpls_port => 0
ofa_bool_t gue_underlay_hash => TRUE
```

# Flexible Assignment of UDP Port Numbers for Decapsulation

**Table 3: Feature History Table**

Feature Name	Release Information	Feature Description
Flexible Assignment of UDP Port Numbers for Decapsulation	Release 25.1.1	<p>Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*)</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> <li>• 8712-MOD-M</li> <li>• 8011-4G24Y4H-I</li> </ul>
Flexible Assignment of UDP Port Numbers for Decapsulation	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>*This feature is now supported on:</p> <ul style="list-style-type: none"> <li>• 8212-48FH-M</li> <li>• 8711-32FH-M</li> <li>• 88-LC1-12TH24FH-E</li> <li>• 88-LC1-36EH</li> <li>• 88-LC1-52Y8H-EM</li> </ul>

Feature Name	Release Information	Feature Description
Flexible Assignment of UDP Port Numbers for Decapsulation	Release 7.3.3	<p>This feature gives you the flexibility to assign UDP port numbers from 1000 through 6400, through which IPv4, IPv6, and MPLS packets can be decapsulated. Such flexibility allows you to segregate the ingress traffic based on a QoS policy.</p> <p>In earlier releases, you could assign only default ports for decapsulation.</p> <p>The following command is introduced for this feature:</p> <pre>hw-module profile gue udp-dest-port ipv4 &lt;port number&gt; ipv6 &lt;port number&gt; mpls &lt;port number&gt;</pre>

This feature provides decapsulation support for GUE packets. In GUE, the payload is encapsulated in an IP packet—IPv4 or IPv6 carrier. The UDP header is added to provide extra hashing parameters and optional payload demultiplexing. At the decapsulation node, the carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload. Prior to Release 7.3.3, packets were decapsulated using UDP port numbers 6080, 6615, and 6635 for IPv4, IPv6, and MPLS payloads respectively. Starting from Release 7.3.3, you can assign UDP port numbers from 1000 through 64000 to decapsulate IPv4, IPv6, and MPLS packets. Define different port numbers for IPv4, IPv6, and MPLS.

## Guidelines for Setting up Decapsulation Using Flexible Port Numbers

Apply these guidelines while assigning flexible port numbers for decapsulation:

Packet	IPv4	IPv6	MPLS
UDP Outer Header	Configure IPv4 port on the hardware module.	Configure IPv6 port on the hardware module.	Configure MPLS port on the hardware module.
Encapsulation Outer Header	Configure an IPv4 encapsulation outer header that matches with the class map source.		
Inner Payload	Note that packets are forwarded based on the inner IPv4 payload.	Note that packets are forwarded based on the inner IPv6 payload.	Note that packets are forwarded based on the inner MPLS payload.

**Note**

- During the decapsulation of the IPv4, IPv6, and MPLS packets, the following headers are removed:
  - The UDP outer header
  - The IPv4 encapsulation outer header
- Select different values for each of these protocols. Valid port numbers are from 1000 through 64000.

## Restrictions

The following restrictions are applicable while configuring unique GUE destination port numbers to decapsulate IPv4, IPv6, and MPLS packets using UDP:

- While configuring the tunnel, select one of the following:
  - Match only 16 unique source IP addresses as shown in the example:
 

```
Router(config-cmap)#match source-address ipv4 210.100.20.0 255.255.255.255
```
  - Match a combination of 64 unique source and destination IP addresses as shown in the example:
 

```
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
```
- The Classless Inter-Domain Routing (CIDR) value in the source IP address subnet mask must be only /32.
- The destination address subnet mask supports all CIDR values. However, the destination address along with the subnet mask must be unique for all the three UDP payload types—IPv4, IPv6, and MPLS. The configuration fails when the destination IP address and the subnet mask are the same for all three payloads as seen in this example:

```
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-IPv4
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-IPv6
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-MPLS
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
```

## Configuring Port Numbers for Decapsulation

By configuring different port numbers on the destination router, you can match and direct traffic to different paths. For example, traffic for a specific video service can be decapsulated and sent through different ports. The steps that are involved in configuring port numbers for decapsulation are:

1. Configure the UDP destination ports for decapsulation of the required payloads.
2. Configure the traffic class to match the ports.
3. Define a policy map, and associate the traffic class with the traffic policy.

#### 4. Apply the policy for each VRF.



**Note** For the hardware module flexible port configuration to take effect you must reload the line card.

#### Configuration Example

```

Hw-module configuration:
=====
Router# configure
Router# hw-module profile gue udp-dest-port ipv4 1001 ipv6 1002 mpls 1003

Class-map configuration:
=====
Router# configure
Router(config)# class-map type traffic match-all udp-v4
Router(config-cmap)# match protocol udp
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1001
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-v6
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1002
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-mpls1
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1003
Router(config-cmap)# end-class-map
Router(config)# commit

Ingress Policy-map configuration:
=====
Router(config)# policy-map type pbr magic-decap
Router(config-pmap)# class type traffic udp-v4
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-mpls1
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit

```

```

Applying policy per VRF:
=====
Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv4 policy type pbr input magic-decap
Router(config-vrf-policy)# commit

```

## Running Configuration

```

!! File saved at 16:01:32 UTC Mon Feb 07 2022 by cisco
!! IOS XR Configuration 7.3.3.10I
!! Last configuration change at Mon Feb 7 15:35:11 2022 by cisco
!
logging console disable
username cisco
  group root-lr
  group cisco-support
  secret 10
$6$gHKmELYZAo71BE1.$3KYogrVodJxTRPZgYPGXUXkO4PqQMr2E6oYvJO4ngBmuaGsF2nAB/m1NP5I13zh9HTzBI/k4r8PwWSbsARsmp.
!
vrf vrf-gre
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
  !
!
line console
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
line default
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
!arp vrf default 29.0.1.2 0000.1122.2929 ARPA
call-home
  service active
  contact smart-licensing
  profile CiscoTAC-1
    active
  destination transport-method http
!
!
ipv6 access-list abf6-gre
  1 permit ipv6 any any nexthop1 ipv6 201:0:1::2
!
ipv4 access-list abf-gre
  1 permit ipv4 any any nexthop1 ipv4 201.0.1.2
!
class-map type traffic match-all udp-v4
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp
  match destination-port 1001
  end-class-map
!
class-map type traffic match-all udp-v6
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp

```

```

    match destination-port 1002
  end-class-map
!
class-map type traffic match-all udp-mpls1
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp
  match destination-port 1003
end-class-map
!
policy-map type pbr pbr-gre
  class type traffic class-default
    redirect ipv4 nexthop 202.0.1.2
  !
end-policy-map
!
policy-map type pbr magic-decap
  class type traffic udp-v4
    decapsulate gue variant 1
  !
  class type traffic udp-v6
    decapsulate gue variant 1
  !
  class type traffic udp-mpls1
    decapsulate gue variant 1
  !
  class type traffic class-default
  !
end-policy-map
!
interface Bundle-Ether25
  ipv4 address 25.0.1.1 255.255.255.0
  ipv6 address 25:0:1::1/64
  ipv6 enable
  shutdown
!
interface Bundle-Ether28
  ipv4 address 28.0.1.1 255.255.255.0
!
interface Loopback0
  ipv4 address 10.10.10.1 255.255.255.255
!
<output truncated>
interface MgmtEth0/RP0/CPU0/0
  ipv4 address dhcp
!
interface MgmtEth0/RP1/CPU0/0
  ipv4 address dhcp
!
interface BVI23
  ipv4 address 23.0.1.1 255.255.255.0
  ipv6 address 23:0:1::1/64
  ipv6 enable
  shutdown
!
interface BVI29
  ipv4 address 29.0.1.1 255.255.255.0
  ipv6 enable
  shutdown
!
interface HundredGigE0/0/0/0
  shutdown
!

```

```
<output truncated>
  l2transport
  !
  !
interface HundredGigE0/0/0/24
  service-policy type pbr input pbr-gre
  ipv4 address 24.0.1.1 255.255.255.0
  ipv6 address 24:0:1::1/64
  ipv6 enable
  !
interface HundredGigE0/0/0/24.24
  ipv4 address 24.0.24.1 255.255.255.0
  ipv6 enable
  encapsulation dot1q 24
  !
interface HundredGigE0/0/0/25
  bundle id 25 mode on
  !
interface HundredGigE0/0/0/26
  ipv4 address 26.0.1.1 255.255.255.0
  ipv6 address 26:0:1::1/64
  ipv6 enable
  !
interface HundredGigE0/0/0/27
  ipv4 address 27.0.1.1 255.255.255.0
  ipv6 enable
  !
interface HundredGigE0/0/0/27.27
  ipv4 address 27.0.27.1 255.255.255.0
  ipv6 address 27:0:27::1/64
  ipv6 enable
  shutdown
  encapsulation dot1q 27
  !
interface HundredGigE0/0/0/28
  bundle id 28 mode active
  !
interface HundredGigE0/0/0/29
  ipv4 address 29.0.1.1 255.255.255.0
  ipv6 enable
  !
<output truncated>
interface HundredGigE0/1/0/24
  ipv4 address 124.0.1.1 255.255.255.0
  ipv6 address 124:0:1::1/64
  ipv6 enable
  !
<output truncated>
!
interface HundredGigE0/1/0/30
  bundle id 28 mode active
  !
interface HundredGigE0/1/0/31
  ipv4 address 31.0.1.1 255.255.255.0
  ipv6 address 31:0:1::1/64
  shutdown
  !
<output truncated>
!
route-policy pass
  pass
end-policy
!
router static
```

```

address-family ipv4 unicast
 201.0.1.0/24 tunnel-ip1
 201.0.1.0/24 tunnel-ip2
 201.0.1.0/24 tunnel-ip3
 201.0.1.0/24 tunnel-ip4
!
address-family ipv6 unicast
 201:0:1::/64 tunnel-ip1
 201:0:1::/64 tunnel-ip2
 201:0:1::/64 tunnel-ip3
 201:0:1::/64 tunnel-ip4
!
!
router ospf 10
router-id 1.1.1.1
area 0
! interface Bundle-Ether28
interface Loopback0
!
interface HundredGigE0/0/0/26
!
!
!
! interface HundredGigE0/0/0/27
! interface HundredGigE0/0/0/27.27
router bgp 200
bgp router-id 1.1.1.1
address-family ipv4 unicast
 maximum-paths ibgp 64
!
! redistribute connected
! neighbor 26.0.1.2
! remote-as 200
! address-family ipv4 unicast
! multipath
! route-policy pass in
! route-policy pass out
! next-hop-self
neighbor 27.0.1.2
remote-as 200
address-family ipv4 unicast
 multipath
 route-policy pass in
 route-policy pass out
 next-hop-self
!
!
neighbor 28.0.1.2
remote-as 200
address-family ipv4 unicast
 multipath
 route-policy pass in
 route-policy pass out
 next-hop-self
!
!
neighbor 29.0.1.2
remote-as 200
address-family ipv4 unicast
 multipath
 route-policy pass in
 route-policy pass out
 next-hop-self
!

```

```

!
!
vrf-policy
 vrf default address-family ipv4 policy type pbr input magic-decap
!
l2vpn
 bridge group bg
  bridge-domain bd
  ! interface HundredGigE0/0/0/29
  !   static-mac-address 0000.1122.2929
  !   routed interface BVI29
 bridge group bgl
  bridge-domain bdl
  interface HundredGigE0/0/0/23
  static-mac-address 0000.1122.2323
  !
  routed interface BVI23
  !
!
!
mpls static
 interface HundredGigE0/0/0/24
 lsp gre
  in-label 35001 allocate per-prefix 202.0.1.2/32
  forward
  path 1 nexthop tunnel-ip1 out-label 35002
  path 2 nexthop tunnel-ip2 out-label 35002
  !
!
!
ssh server vrf default
hw-module profile gue udp-dest-port ipv4 1001 ipv6 1002 mpls 1003
end

```

## Verification

Run the **show ofa objects sys location 0/0/CPU0 | inc gue** command in the XR Config mode to verify that the unique GUE port numbers have been configured to decapsulate IPv4, IPv6, and MPLS payloads.

```

Router#show ofa objects sys location 0/0/CPU0 | inc gue
uint32_t gue_ipv4_port => 1001
uint32_t gue_ipv6_port => 1002
uint32_t gue_mpls_port => 1003

```

## GUEv1 static tunnel configuration over IPv4 networks

A GUEv1 static tunnel is a user-specified tunnel that

- is configured between two endpoints using static route configuration that includes a next-hop IP address.
- enables efficient, scalable, and secure deployment of GUEv1 tunnels, and
- carries IPv4 and IPv6 traffic over IPv4 networks.

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
GUEv1 static tunnel configuration over IPv4 networks	Release 25.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200]); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q200])</p> <p>This feature enables highly efficient, scalable, and secure deployment of Generic UDP Encapsulation variant 1 (GUEv1) static tunnels to carry IPv4 and IPv6 traffic over IPv4 networks.</p> <p>In addition, this feature enables the router to use GUEv1 as the overlay encapsulation type for network virtualization endpoint (NVE) and send encapsulated overlay traffic to IPv4, IPv6, and MPLS destinations using UDP ports.</p> <p>With this feature, you can also increase the GUEv1 scale up to 12K tunnels.</p> <p>This feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> <li>The <b>address-family ipv4 unicast</b> and <b>address-family ipv6 unicast</b> commands are modified to include these keywords: <b>remote-next-hop</b>, <b>tunnel</b>, <b>source-ip</b>, <b>dscp</b>, and <b>ttl</b>.</li> <li>A new command, <b>nve overlay-encap guev1 udp-port destination</b>, has been introduced.</li> </ul>

Starting from Cisco IOS XR Release 25.4.1, you can configure GUEv1 static tunnels to carry IPv4 and IPv6 traffic over IPv4 networks.

## Configuration guidelines for GUEv1 static tunnels

These are some guidelines to configure GUEv1 static tunnels:

- Configure **hw-module profile cef iptunnel scale** to increase the GUEv1 tunnel scale up to 12K tunnels. This tunnel scale may vary based on the number of ECMP tunnels used.
- By default, up to 3500 GUEv1, MPLS over UDP, and IP-in-IP tunnels including the number of L3 interfaces are supported.
- Configure the **nve overlay-encap guev1 udp-port destination** command for IPv4, IPv6, and MPLS to send encapsulated overlay traffic to IPv4, IPv6, and MPLS destinations using UDP ports 6080, 6615, and 6635 respectively.

## Restrictions for GUEv1 static tunnels

These are some restrictions specific to the GUEv1 static tunnel configuration:

- Traceroute over GUEv1 static tunnels is not supported both in uniform and pipe modes.

- Non-default VRF is not supported for static GUEv1 static tunnels.

## Configure IPv4 and IPv6 GUEv1 static tunnels

This procedure provides the steps to configure the IPv4 and IPv6 GUEv1 static tunnels.

### Procedure

**Step 1** Execute the `prefix mask remote-next-hop` command to configure the IPv4 or IPv6 GUEv1 static tunnels.

#### Example:

This example provides the IPv4 GUEv1 static tunnel configuration.

```
RP/0/RP0/CPU0:router# config
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 9.8.7.6/32 remote-next-hop 5.5.5.6 tunnel gue-v1 source-ip
 1.1.1.1 dscp 46 ttl 250
RP/0/RP0/CPU0:router(config-static-afi)# commit
```

#### Example:

This example provides the IPv6 GUEv1 static tunnel configuration.

```
RP/0/RP0/CPU0:router# config
RP/0/RP0/CPU0:router(config)# router static
RP/0/RP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-static-afi)# 11:1:1::1/128 remote-next-hop 10.10.10.11 tunnel gue-v1
source-ip 2.2.2.2 dscp 8 ttl 250
RP/0/RP0/CPU0:router(config-static-afi)# commit
```

**Step 2** Execute the `show running-config router static` command to view the running configuration for the GUEv1 static tunnel.

```
RP/0/RP0/CPU0:router# show running-config router static
Thu Mar 20 20:12:16.055 IST
router static
  address-family ipv4 unicast
    9.8.7.6/32 remote-next-hop 5.5.5.6 tunnel gue-v1 source-ip 1.1.1.1 dscp 46 ttl 250
  address-family ipv6 unicast
    11:1:1::1/128 remote-next-hop 10.10.10.11 tunnel gue-v1 source-ip 2.2.2.2 dscp 8 ttl 25
!
```

**Step 3** Execute the `show static afi-all topology` to verify the IPv4 and IPv6 static routes.

#### Example:

```
RP/0/0/CPU0:router# show static afi-all topology
Thu Mar 20 20:13:29.394 IST
IPv4 Unicast:
-----
VRF: default Table Id: 0xe0000000 AFI: IPv4 SAFI: Unicast
Prefix/Len Interface Nexthop Object
Explicit-path Metrics Local-Label
127.127.127.8/32 None 11.11.11.2 None
None [0/0/1/0/0] No label

IPv6 Unicast:
-----
```

```

VRF: default Table Id: 0xe0800000 AFI: IPv6 SAFI: Unicast
Prefix/Len Interface Nexthop Object
Explicit-path Metrics Local-Label
127:127:127::8/128 None 11.11.11.2 None
None [0/0/1/0/0] No label

```

### What to do next

Make sure that you configure the global UDP ports for GUEv1 static tunnels.

## Configure global UDP ports for GUEv1 static tunnels

This configuration enables the router to use GUEv1 as the overlay encapsulation type for NVE and to send encapsulated overlay traffic to IPv4, IPv6, and MPLS destinations using UDP ports 6080, 6615, and 6635 respectively.

Follow these steps to configure global UDP ports for GUEv1 static tunnels.

### Before you begin

Make sure that you have configured GUEv1 static tunnels.

### Procedure

**Step 1** Set the destination UDP port for the encapsulated packets to port 6080 on IPv4.

```

RP/0/RP0/CPU0:router# config
RP/0/RP0/CPU0:router(config)# nve overlay-encap guev1 udp-port destination ipv4 6080
RP/0/RP0/CPU0:router(config)# commit

```

**Step 2** Set the destination UDP port for the encapsulated packets to port 6615 on IPv6.

#### Example:

```

RP/0/RP0/CPU0:router(config)# nve overlay-encap guev1 udp-port destination ipv6 6615
RP/0/RP0/CPU0:router(config)# commit

```

**Step 3** Set the destination UDP port for the encapsulated packets to port 6635 for MPLS.

#### Example:

```

RP/0/RP0/CPU0:router(config)# nve overlay-encap guev1 udp-port destination mpls 6635
RP/0/RP0/CPU0:router(config)# commit

```

**Step 4** Execute **show running-config nve overlay-encap guev1** command to verify the configured UDP ports.

#### Example:

```

RP/0/RP0/CPU0:router# show running-config nve overlay-encap guev1
Wed Nov 19 19:47:11.595 UTC
nve
  overlay-encap guev1
    udp-port destination ipv4 6080
    udp-port destination ipv6 6615

```

```
udp-port destination mpls 6635  
!  
!
```

---

