Configure First Hop Security

This chapter contains these sections:

DHCP Snooping in VXLAN BGP EVPN

DHCP Snooping in VXLAN BGP EVPN is a process that

- validates ARP/GARP packets sent from a host, preventing ARP spoofing and malicious ARP storms,
- validates data-plane traffic from the host using IPSG, preventing malicious hosts from sending data traffic, and
- replicates DHCP Snooping entries across the VXLAN fabric, enabling DAI and IPSG to function across the fabric even after a host move.

First Hop Security

First Hop Security (FHS) is an access security feature that

- provides security to the network at the access point where the host connects to the first switch,
- authorizes and authenticates hosts, and
- protects the network by ensuring only authorized hosts are allowed network access.

The Dot1x, port-security, and DHCP Snooping are examples of access security features.

DHCP Snooping database

DHCP Snooping database (DB) is a database that

- contains the MAC address of the host, the IP address assigned to the host by the DHCP server, VLAN, and other details like the lease time,
- can contain local or remote snooping DB entries, and
- can be configured using the **ip source binding** *ip address* **vlan** *vlan-id* **interface** *interface* interface command.

Note Snooping entries added through this command are referred to as static entries, and they are distributed across all VTEPs.

Distributed DHCP Snooping database

Distributed DHCP Snooping DB is a database that

- validates ARPs/GARPs sent from the host using DAI,

Note The ARP/GARP will be dropped if there is no matching entry in the DB.

- validates data-plane traffic from the host using IPSG, and
- is replicated across the fabric.

How DHCP Snooping on VXLAN works

The DHCP Snooping on VXLAN process involves exchanges between your host and a DHCP server across a VXLAN fabric. It creates and distributes a DHCP Snooping database, which validates ARP/GARP packets and data-plane traffic.

Figure 1: DHCP Snooping on VXLAN

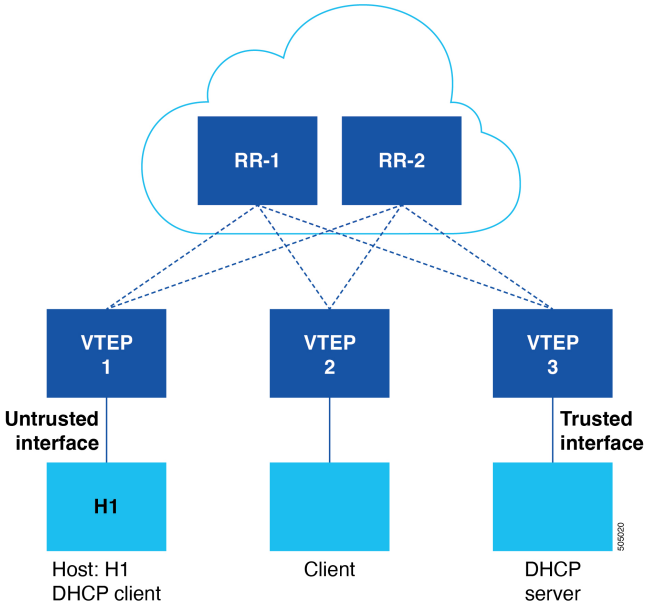

The key components involved in the process are:

- **Host**: Requests an IP address. Host H1 is attached to VTEP1
- **VTEP**: Connects hosts to the VXLAN fabric and enforces security policies.
- **DHCP Server**: Assigns IP addresses and configuration parameters to hosts. DHCP server is attached to VTEP3.
- **VXLAN Fabric**: The network infrastructure that connects VTEPs and enables communication between hosts and servers.
- **DHCP Snooping Database**: Stores IP-to-MAC address bindings and other host information.

The process includes these stages:

- The host sends a DHCP Discover message to find a DHCP server.
- The DHCP server sends a DHCP Offer message to the host, proposing an IP address.
- The host sends a DHCP Request message to accept the offered IP address.
- The DHCP server sends a DHCP Ack message to confirm the IP address assignment.

Note

- The host and the DHCP server exchange a set of messages as part of this host IP assignment procedure. These are known as Discover-Offer-Request-Ack (DORA) exchange messages.
- The DORA exchange, for a particular host (H1), must now be sent over the VXLAN fabric to reach remote DHCP servers (VTEP3).
- VTEP3 checks whether the 'Offer' and 'Ack' messages (part of a DORA sequence) coming from the DHCP server are received via a Trusted Interface.

- Upon completion of the DORA exchange, VTEP1 creates a local DHCP Snooping database entry for the host.
- The local DHCP Snooping database entry is propagated to remote VTEPs using BGP-EVPN.
- Remote VTEPs store the snooping database entry as a remote entry.
- DAI validates ARP/GARP packets against the DHCP Snooping database to prevent spoofing.
- IPSG validates data-plane traffic against the DHCP Snooping database to prevent malicious traffic.

  In the IPSG, only local DHCP clients for that VTEP are programmed. The local DHCP clients are identified with anchor flag set to true in the DHCP Snooping table. If a host moves to a different VTEP and settles down, IPSG has to reprogram the client behind the new VTEP to validate the data-traffic. On the old VTEP, IPSG has to remove this DHCP client. The anchor flag will change accordingly. The host move is triggered by the receipt of an ARP request from the host which is received on the new VTEP that the host moved to.
- The DHCP Snooping database updates to reflect the new location when a host moves to a different VTEP.

  This DHCP Snooping database will be seen as a 'Distributed Database' across the VTEPs, and the snooping entries will be in sync with all VTEPs.

The DHCP Snooping on VXLAN process ensures secure IP address assignment and prevents malicious activities by validating traffic against a distributed database of IP-to-MAC address bindings.

Guidelines and limitations for DHCP Snooping on VXLAN

DHCP Snooping on VXLAN feature has the following guidelines and limitations:

Configuration guidelines and limitations

- Ensure that the DHCP Snooping, DAI and IPSG together are enabled on all VTEPs.

Note DAI and IPSG depend on DHCP Snooping. DHCP Snooping creates the snooping DB and this DB is used by DAI and IPSG.

- The host-move is indicated by ARP/GARP/RARP receipt. In case of RARP (which contains MAC info alone), VTEPs start ARP Refreshes for the IPs learned against MAC. ARP-GARP is essentially the trigger for host-move and not any other data packet.
- In the ingress SUP region, the TCAM must be carved out to 768 entries instead of the default 512 entries to set up the ingress ACLs using the **hardware access-list tcam region ing-sup** command. Reload of a switch is required for the TCAM carving changes to reflect.
- If an I/O module lacks the resources required for an atomic update, you can disable atomic updates by using the **no hardware access-list update atomic** command; however, during the brief time required for the device to remove the preexisting ACL and implement the updated ACL, traffic that the ACL applies to is dropped by default.
- If you want to permit all traffic that an ACL applies to while it receives a nonatomic update, use the **hardware access-list update default-result permit** command.
- In case of multisite and with vPC BGW, if DHCP Snooping is enabled on the vPC BGW, ensure that DHCP clients and DHCP servers are on the same sites.

Note

- DHCP Snooping needs to be enabled (on a VTEP) for the VLAN belonging to the DHCP host that must avail the DHCP service.
- All the VLANs serviced by the DHCP server in the fabric should be enabled with DHCP Snooping on all the VTEPs of the fabric.

Supported features and platforms

- Beginning with Cisco NX-OS Release 10.4(1)F, DHCP Snooping and associated features such as Dynamic ARP Inspection (DAI) and IP Source Guard (IPSG) support is extended to VXLAN fabric on Cisco Nexus 9300-EX/FX/FX2/FX3/GX/GX2 platform switches and Cisco Nexus 9500 switches with 9700-EX/FX/GX line cards.

  Beginning with Cisco NX-OS Release 10.4(2)F, First Hop Security feature is supported on Cisco Nexus 9332D-H2R, and 93400LD-H1 switches.

  Beginning with Cisco NX-OS Release 10.4(3)F, First Hop Security feature is supported on Cisco Nexus 9364C-H1 switches.
- Beginning with Cisco NX-OS Release 10.5(2)F, First Hop Security feature is supported on Cisco Nexus 9500 Series switches with N9K-X9736C-FX3 line card.
- Only IPv4 multicast underlay is supported. However, IPv4 ingress replication underlay, IPv6 ingress replication underlay and IPv6 multicast underlay are not supported.
- Only IPv4 DHCP hosts is supported.
- For vPC VTEPs, only physical MCT is supported.
- On vPC nodes, static DHCP Snooping is supported only with vPC port-channel ports and not with orphan ports.

Unsupported features and platforms

- First Hop Security feature is not supported on EoR.
- The DHCP server cannot be deployed behind the EoR.
- This feature cannot coexist with FabricPath to VXLAN migration feature, and the counter ACL (CNT ACL) feature.

Prerequisites for DHCP Snooping

DHCP Snooping has these prerequisites:

- You should be familiar with DHCP before you configure DHCP Snooping or the DHCP relay agent.
- Make sure that the DHCP Snooping, DAI and IPSG features are enabled together on a leaf VTEP.

Enable DHCP Snooping on VXLAN

You can enable or disable DHCP Snooping on a single-box feature or enable this feature for a VLAN for the entire fabric. By default, DHCP Snooping is disabled on all VLANs.

Follow these steps to enable DHCP Snooping on VXLAN

Before you begin

- Ensure you enable the DHCP feature.
- Ensure you configure the **nv overlay evpn** command.
- Ensure you enable the DHCP Snooping, DAI, and IPSG features. For more information see the Prerequisites for DHCP Snooping, on page 6 section.
- Ensure you enable the DHCP Snooping and DAI features on all the VXLAN nodes. For more information on configuration, see **Configuring DHCP Snooping** section of Cisco Nexus 9000 Series NX-OS Security Configuration Guide.
- Ensure you enable the DHCP Snooping trust and ARP inspection trust on interfaces connected to the DHCP server nodes. For more information on configuration, see **Configuring DHCP Snooping** section of Cisco Nexus 9000 Series NX-OS Security Configuration Guide.
- Ensure you enable the IP Source Guard feature on the interfaces connected to the DHCP client nodes. For more information on configuration, see **Configuring DHCP Snooping** section of Cisco Nexus 9000 Series NX-OS Security Configuration Guide.

Procedure

Step 1 Run the **configure terminal** command to enter global configuration mode.

Example:

```
switch# configure terminal
switch(config)#
```

Step 2 Run the [**no**] **ip dhcp snooping vlan** *vlan-list* **evpn** command to enable DHCP Snooping on the VLANs specified by *vlan-list* .

Example:

```
switch(config)# ip dhcp snooping vlan 100,200,250-252 evpn
```

Beginning with Cisco NX-OS Release 10.4(1)F, the **evpn** option is provided to support host move to other interfaces on the same VTEP or other VTEPs.

Note

- When this feature is enabled with the **evpn** option, the **nve** will be implicitly added as a trusted interface.
- It is possible to have one VLAN list with **evpn** keyword and another VLAN list with **no evpn** keyword.

The **no** form of this command disables DHCP Snooping on the VLANs specified.

Step 3 (Optional) Run the **show running-config dhcp** command to validate the DHCP configuration.

Example:

```
switch(config)# show running-config dhcp
```

Step 4 (Optional) Run the **copy running-config startup-config** command to copy the running configuration to the startup configuration.

Example:

```
switch(config)# copy running-config startup-config
```

Clear the duplicate host after permanent freeze

If the MAC or MAC-IP address becomes permanently frozen, you cannot automatically recover to restart mobility or check for duplicates.

The mobility and duplicate detection logic for DHCP clients in FHS-enabled VTEPs is the same as the BGP EVPN's mobility and duplicate detection logic. However, duplicate detection may happen in any of the VTEPs in non-FHS deployments. In FHS deployments, the host duplicate will always be detected on a VTEP where the DHCP binding entry is remote.

For more information on mobility and duplicate detection, see the Duplicate Detection for IP and MAC Addresses section.

Once the MAC or MAC-IP is permanently frozen, there is no auto recovery mechanism to re-initiate mobility or duplicate check sequences. To clear MAC and MAC-IP permanent freeze state, use these commands:

- For MAC:

```
clear l2route evpn mac [mac-address] [topo] permanently-frozen-list
```

- For MAC-IP:

```
clear fabric forwarding dup-host [{ ip|ipv6 address }] [vrf {vrf-name | vrf-known-name
 | all}]
```

Verification of DHCP Snooping configuration

To validate the DHCP Snooping configuration, enter these commands:

Command	Purpose
show ip dhcp snooping binding evpn	Displays all entries from the DHCP snooping binding database.
show l2route fhs [**topology** *topology id* | **all**]	Displays all entries from the L2RIB database.

The following example shows sample output for the **show ip dhcp snooping binding evpn** command:

```
switch(config)# show ip dhcp snooping binding evpn
MacAddress         IpAddress        Lease(Sec)  Type      BD    Interface       anchor
  Freeze
------------------ -------------- ----------- ---------- ---- ----------------- ----------
---------
00:10:00:10:00:10  10.10.10.10      infinite  static      2001  Ethernet1/48         YES
     NONE
00:15:06:00:00:01  100.1.150.156    86282     dhcp-snoop  2001  Ethernet1/31         YES
     NONE
00:17:06:00:00:01  100.1.150.155    86265     dhcp-snoop  2001  nve1(peer-id: 1)     NO
     NONE
```

The following example shows sample output for the **show l2route fhs** command:

```
switch(config)# show l2route fhs all
Flags - (Stt):Static (Dyn):Dynamic (R):Remote
Topo ID  Mac Address     Host IP             Prod          Flags      Seq No     Next-Hops
---- -- ---- ----------- ------------------ ---------- ----------- ------------ -------------
2001     0015.0600.0001  100.1.150.156      DHCP_DYNAMIC  Dyn,       0          Eth1/31
2001     0017.0600.0001  100.1.150.155      BGP           Dyn,R,     0          1.13.13.13
 (Label: 0)
switch(config)#
```

The following example shows DHCP configurations for a VTEP with DHCP clients:

```
feature dhcp
service dhcp
ip dhcp snooping
ip dhcp snooping vlan 2001-2002 evpn
ip arp inspection vlan 2001-2002

interface Ethernet1/31
ip verify source dhcp-snooping-vlan
```

The following example shows DHCP configurations for a VTEP with DHCP server:

```
feature dhcp
service dhcp
ip dhcp snooping
ip dhcp snooping vlan 2001-2002 evpn
ip arp inspection vlan 2001-2002

interface Ethernet1/47
ip dhcp snooping trust
ip arp inspection trust
```