Configure VXLAN BGP-EVPN Null Route

This chapter contains these sections:

EVPN null routes

An EVPN null route is a network security mechanism that

- enables devices in an EVPN (Ethernet VPN) fabric to drop traffic destined for specific MAC or IP addresses identified as part of a Distributed Denial of Service (DDoS) attack
- is implemented by injecting null (drop) route entries into the forwarding tables of VTEPs (Virtual Tunnel Endpoints), and
- helps prevent malicious traffic from consuming network bandwidth and impacting legitimate traffic between hosts.

A null route is a network route (routing table entry) that leads nowhere. All matching packets are dropped (ignored or redirected) rather than forwarded, acting as a limited firewall. Null route filtering refers to the practice of setting such routes.

NX-OS supports configuring null/drop routes for IPv4, IPv6, and MAC addresses, and these routes should be distributed to all VTEPs in the fabric for comprehensive protection.

- For IPv4/IPv6 attacks, use these commands to configure an IPv4/IPv6 static route to the null interface:
  - **ip route x.x.x.x/y Null0**
  - **ipv6 route X:X:X::X/Y Null0**
- For MAC-based attacks, use **mac address-table static xxxx.yyyy.zzzz vlan <VLAN-ID> drop** command to drop packets at Layer 2.

- A DDoS attack on a host in an EVPN Fabric consumes the network bandwidth resources and in turn impacts legitimate traffic to other hosts.
  - The DDoS attack can originate from:
    - Host connected to a leaf switch within the local site
    - Host connected to a leaf switch in a remote site
    - External networks such as WAN
  - DDoS attacks can be intra-subnets (MAC based) or inter-subnets (Host-based – IPv4/IPv6)
  - Null route filtering has been traditionally used in mitigating DDoS attacks especially in service provider networks.

Table 1: Comparison of manual and EVPN null route configuration

	Manual configuration of null routes	**EVPN null routing feature**
Scalability	Difficult across many VTEPs/sites	Scalable via automated route injection
Ease of management	Labor-intensive; prone to errors	Centralized control with orchestrator
Attack containment	May be delayed; not uniformly applied	Immediate, edge-based traffic drop

Example

In large EVPN deployments with many VTEPs and multiple sites, manually configuring drop (null) routes for attack targets on each VTEP is challenging without an orchestrator like Nexus Dashboard Fabric Controller (NDFC). The EVPN null routing feature enables a VTEP to send null-tagged Type-2 and Type-5 routes, so other VTEPs (Borders and Leafs) automatically install drop entries in the IP or MAC tables, discarding attack traffic at network edges to preserve bandwidth.

Requirement: Configure and manage VXLAN BGP-EVPN null route MACs consistently

Follow these requirements to ensure consistent and reliable deployment of VXLAN BGP-EVPN null route MAC configurations.

- A null route (static) MAC configuration must have matching static ARP/ND configuration. You must not configure dynamic ARP/ND with MACs set as null route MACs.
- If you use only L2 services and have no configuration that leads to dynamic ARP/ND learning, a “mac drop” configuration alone is allowed. In all other cases, always configure static ARP/ND alongside the “mac drop” configuration.
- In vPC environments, always configure the null route (MAC, mac-ip, prefix) on both vPC boxes (VMCT and PMCT). The behavior is undefined if this is not configured on both boxes, including during unconfiguration. vPC consistency checker support is not available for this feature.

- Always apply the route-map on remote VTEPs. This ingress route-map is essential for Type-5 routes.
- Do not expect feature interaction with multicast traffic.
- If remote static is present on a VTEP and you configure the same MAC as a local static (either with a valid interface or as “mac drop”/null route), a syslog message will warn of duplicate configuration and require correction, although the configuration is not rejected. Local static configuration takes precedence over remote static on the same VTEP.
- If you change a local static MAC with a valid interface on a VTEP to a null route MAC, the null route MAC takes effect.
- Although remote dynamic MAC routes allow any remote MAC route derived from MAC-IP route split to overwrite its entry and propagate to the MAC manager, remote static MAC routes do not honor these derived MACs for overwriting. The MAC entry remains unchanged until you delete the remote static MAC.
- Treat the null route MAC as a form of static MAC configuration only.

Configure static MAC addresses

Manually add static drop MAC addresses to ensure they override dynamically learned MAC addresses on selected interfaces.

Follow these steps to configure static MAC addresses:

Before you begin

- Identify the MAC addresses and VLAN IDs you want to configure.
- Ensure you have administrative access to the switch CLI.

Procedure

Step 1 Enter global configuration mode.

Example:

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

Step 2 Define the static MAC address, specifying the VLAN and action (drop, interface, or port-channel).

Example:

```
switch(config)# mac address-table static 3001.3010.99aa vlan 3001 drop
switch(config)#
```

Step 3 Exit the configuration mode.

Example:

```
switch# exit
switch#
```

The static MAC address is successfully added to the Layer 2 MAC address table and will override any dynamically learned MAC address entries for the specified interface or VLAN.

Configure ARP/ND

Establish static ARP and Neighbor Discovery (ND) entries on a switch SVI for IPv4 and IPv6 hosts, ensuring stable MAC-IP bindings and correct forwarding behavior.

Use this procedure to prevent MAC-IP mobility by ensuring both DROP MAC and MAC-IP entries originate from the same VTEP and to support static host configuration in Layer 2/Layer 3 integrated network environments.

Before you begin

- Identify the VLAN, VRF, IP addresses, and MAC addresses needed for configuration.
- Preconfigure any static MAC-IP entries as drop entries to avoid mobility issues.

Procedure

Step 1 **configure terminal**

Example:

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

Step 2 Specify the VLAN interface.

Example:

```
switch(config)# interface Vlan 3001
switch(config-if)#
```

Step 3 Assign the VLAN interface to the tenant VRF.

Example:

```
switch(config-if)# vrf member cgw_3001_3050
switch(config-if)#
```

Step 4 Assign static IPv4 and IPv6 addresses.

a) Static IPv4 address

Example:

```
switch(config-if)# ip address 192.0.2.1/16
```

b) Static IPv6 address

Example:

```
switch(config-if)# ipv6 address 2001:DB8::1/64
```

Step 5 Disable IPv4 and IPv6 redirects.

Example:

For IPv4

```
switch(config-if)# no ip redirects
switch(config-if)#
```

For IPv6

```
switch(config-if)# no ipv6 redirects
switch(config-if)#
```

Step 6 Configure static IPv6 neighbor.

Example:

```
switch(config-if)# switch(config-if)# ipv6 neighbor 2001:DB8::99 3001.3010.99aa
```

Step 7 Associate an IP address with a MAC address as a static entry.

Example:

```
switch(config-if)# switch(config-if)# ip arp 192.0.2.99 3001.3010.99aa
```

Step 8 Enable anycast gateway forwarding for the VLAN.

Example:

```
switch# fabric forwarding mode anycast-gateway
```

The VLAN interface (SVI) is configured with static ARP and ND bindings, and IP redirects are disabled. This ensures consistent forwarding and prevents unwanted MAC-IP mobility.

Configure a prefix-null route on a local VTEP

Use this task when you need to discard specific destination prefixes within a tenant VRF and ensure static null routes are advertised through BGP.

Follow these steps to configure a prefix-null route and redistribute it into BGP:

Before you begin

- Make sure you have administrative access.
- Gather relevant VRF name, prefix, mask, route tag, route-map name, and BGP AS number.
- Confirm that BGP and tenant VRF exist on the device.

Procedure

Step 1 Enter global configuration mode.

Example:

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

Step 2 Access the tenant VRF context.

Example:

```
switch(config)# vrf context tenant-0001
switch(config-vrf)#
```

Step 3 Configure a static Null0 route for the required prefix with the matching tag (IPv4 or IPv6 as needed).

Example:

For IPv4

```
switch(config-vrf)# ip route 203.0.113.0/24 Null0 tag 6666
switch(config-vrf)#
```

For IPv6

```
switch(config-vrf)# ipv6 route 2001:DB8::1/120 Null0 tag 6666
switch(config-vrf)#
```

Step 4 Create or enter the route-map and configure an entry.

Example:

```
switch(config)# route-map SET_BHC permit 10
switch(config-route-map)#
```

Use seq to order the entries in a route map.

Step 5 In the route-map entry, match the configured tag.

Example:

```
switch(config-route-map)# match tag 6666
switch(config-route-map)#
```

Step 6 Set the weight and the community attribute to blackhole in the route-map.

Example:

```
switch (config-route-map)# set weight 65535
switch(config-route-map)# set community blackhole
```

We recommend to set the **set weight** value to maximum value, to give the highest precedence to the null routes. The maximum value of **set weight** is 65535.

Step 7 Enter BGP configuration mode and specify the AS number.

Example:

```
switch(config)# router bgp 100
switch(config-router)#
```

The range of as-num is 1–65535.

Step 8 Configure BGP for the correct tenant VRF and address-family (IPv4/IPv6 unicast).

Example:

```
switch(config-router)# vrf tenant-0001
switch(config-router-vrf)# address-family ipv4 unicast
switch(config-router-vrf-af)#
```

This configuration is required for IPv4/IPv6 over VXLAN with IPv4/IPv6 underlay.

Step 9 Redistribute static routes into BGP using the created route-map.

Example:

```
switch(config-router-vrf-af)# redistribute static route-map SET_BHC
switch(config-router-vrf-af)#
```

Redistributes the prefix-null static route into BGP using the configured route-map.

The local VTEP advertises the configured static prefix-null route in BGP as a null route for the specified tenant VRF, causing matching traffic to be discarded.

Configure RPM route-map on remote VTEP

Assign highest precedence to null routes on a remote VTEP by using a community-list and a route-map.

Apply this configuration when you need to ensure routes marked with the well-known “blackhole” community are preferred by the VTEP.

Follow these steps to configure the RPM route-map:

Before you begin

- Ensure you have access to the remote VTEP.
- Confirm BGP is already enabled on the device.

Procedure

Step 1 Enter global configuration mode.

Example:

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

Step 2 Create a standard community-list to permit routes with the "blackhole" community.

Example:

```
switch (config)# ip community-list standard BH seq 10 permit blackhole
switch(config)#
```

Configures a community list and permits routes that have the well-known "blackhole" community value.

Beginning with Cisco NX-OS Release 10.3(2)F, the blackhole (well-known community) is added to the existing IP community list.

Step 3 Create a route-map and enter route-map configuration mode.

Example:

```
switch(config)# route-map PREFER_BHC permit 10
switch(config-route-map)#
```

Step 4 Match the community-list in the route-map.

Example:

```
switch(config-route-map)# match community BH
switch(config-route-map)#
```

The BGP routes are matched using the community list.

Step 5 Set the weight for the matched routes to the maximum value to ensure highest precedence.

Example:

```
switch (config-route-map)# set weight 65535
switch(config-route-map)#
```

We recommend to set the **set weight** value to maximum value, to give the highest precedence to the null routes. The maximum value of **set weight** is 65535.

Step 6 Add a fallback permit clause in the route-map to allow other routes.

Example:

```
switch(config-route-map)# route-map PREFER_BHC permit 20
switch(config-route-map)#
```

Step 7 Enter BGP configuration mode.

Example:

```
switch(config)# router bgp 100
switch(config-router)#
```

Enables a routing process. The range of as-num is from 1 to 65535.

Step 8 Apply the route-map to the BGP neighbor in the inbound direction.

Example:

```
switch(config-router-neighbor-af)# route-map PREFER_BHC in
```

The RPM route-map is implemented, giving highest priority to null routes (routes with the blackhole community) and optimizing traffic handling on the remote VTEP.

Null route configuration options

A null route (also known as a blackhole route) is a network routing option used to intentionally drop traffic that matches specified prefixes, MAC addresses, or MAC-IP combinations. Null routes are most commonly deployed to enhance security, mitigate attacks, and segment traffic in data center fabrics such as VXLAN EVPN.

Configuration options

- Prefix null route

  Use case: Drop traffic matching specified IPv4/IPv6 prefixes.

  - Local VTEP configuration:

    On local VTEP (Border leaf switch) where the Type-5 null route is to be advertised, perform these steps:

    - Configure static IPv4/IPv6 address with Null0 adjacency

```
vrf context tenant-0001
  vni 3100001
  ip route 50.1.0.0/24 Null0 tag 6666
  ipv6 route 50::1:0/120 Null0 tag 6666
```

- Configure route-map to set null route community on static route and redistribute into BGP

```
route-map SET_BHC permit 10
  match tag 6666
  set community blackhole
router bgp 100
  router-id 10.1.0.21
  vrf tenant-0001
    address-family ipv4 unicast
      redistribute static route-map SET_BHC
    address-family ipv6 unicast
      redistribute static route-map SET_BHC
```

- Remote VTEP configuration:

  Configure route-map to match the null route community and set weight to highest value to ensure null route is always preferred.

```
ip community-list standard BH seq 10 permit blackhole
route-map PREFER_BHC permit 10
  match community BH
  set weight 65535
route-map PREFER_BHC permit 20
router bgp 100
 router-id 10.1.0.13
 address-family l2vpn evpn
 template peer LEAF_to_FABRIC_IBGP_OVERLAY
    remote-as 100
    address-family l2vpn evpn
      send-community
      send-community extended
      route-map PREFER_BHC in
```

- Configuration – MAC/MAC-IP Drop

  Use case: Drop traffic matching specified MAC addresses or MAC-IP pairs.

  - Local VTEP configuration:

    On local VTEP where Type-2 null route is to be advertised, perform the following steps:

    - Configure static MAC address with drop adjacency

```
mac address-table static 0013.e001.0001 vlan 2 drop
```

    - Configure static ARP/ND neighbor for same address

```
interface Vlan2
  no shutdown
  vrf member tenant-0001
  ip address 5.0.63.254/18
  ipv6 address 5::3f7f/114
  ipv6 neighbor 5::17fe 0013.e001.0001
  no ipv6 redirects
  ip arp 5.0.23.254 0013.e001.0001
  fabric forwarding mode anycast-gateway
```

  - Remote VTEP configuration:

    On all other remote VTEPs, perform the following step:

    - Configure route-map to match the blackhole community and set weight to highest value to ensure null route is always preferred.

```
ip community-list standard BH seq 10 permit blackhole
route-map PREFER_BHC permit 10
  match community BH
  set weight 65535
route-map PREFER_BHC permit 20
router bgp 100
 router-id 10.1.0.13
 address-family l2vpn evpn
 template peer LEAF_to_FABRIC_IBGP_OVERLAY
    remote-as 100
    address-family l2vpn evpn
      send-community
      send-community extended
      route-map PREFER_BHC in
  neighbor 10.1.0.31
    inherit peer LEAF_to_FABRIC_IBGP_OVERLAY
```

EVPN null route verification commands

EVPN null route verification commands provide information required to confirm the configuration and operation of EVPN null routes on a network device. These commands display routing details, ARP records, adjacency entries, MAC address tables, community lists, and route maps, which are useful for troubleshooting and validation.

Table 2: EVPN Null Route Verification Commands

Command	Purpose
show bgp l2vpn evpn	Displays routing table information.
show ip arp static vlan <*vlan-id*> **vrf** <*vrf-name*>	Displays local ARP information.
show ip arp static remote vlan <*vlan-id*> **vrf** <*vrf-name*>	Displays remote ARP information.
show ip adjacency vlan <*vlan-id*> **detail vrf** <*vrf-name*>	Displays local adjacency information.
show ipv6 icmp neighbour static remote [*vlan* <*id*>] [*vrf* <*name*>]	Displays remote static neighbor information.
show mac address-table static vlan <*vlan-id*>	Displays local/remote MAC information.
show ip community-list *name*	Displays information about a IP community list.
show route-map *name*	Displays information about a route map.

Sample outputs:

These output samples illustrate key attributes to review when verifying EVPN null route configuration, (such as the "Community: blackhole" flag), routing entry details, and path indicators.

```
switch# show bgp l2vpn evpn 1111.1111.1111
BGP routing table information for VRF default, address family L2VPN EVPN
```

```
Route Distinguisher: 192.0.2.53:32769    (L2VNI 1000002)
BGP routing table entry for [2]:[0]:[0]:[48]:[1111.1111.1111]:[32]:[192.0.2.51]/272, version
 23
Paths: (1 available, best #1)
Flags: (0x000102) (high32 00000000) on xmit-list, is not in l2rib/evpn
Multipath: eBGP iBGP
  Advertised path-id 1
  Path type: local, path is valid, is best path, no labeled nexthop, has esi_gw
  AS-Path: NONE, path locally originated
  192.0.2.53 (metric 0) from 0.0.0.0 (192.0.2.53)
  Origin IGP, MED not set, localpref 100, weight 32768
  Received label 1000002 1000100
  Community: Blackhole
  Extcommunity: RT:23456:1000002 RT:23456:1000100 ENCAP:8
  Router MAC:0476.b0f0.8157
  Path-id 1 advertised to peers:
  192.0.2.1
```

```
switch# sh bgp ipv4 uni 192.0.2.0 vrf 100
BGP routing table information for VRF 100, address family IPv4 Unicast
BGP routing table entry for 192.0.2.0/24, version 6
Paths: (1 available, best #1)
Flags: (0x80c0002) (high32 0x000020) on xmit-list, is not in urib, exported, has label
vpn: version 5, (0x00000000100002) on xmit-list
local label: 492287

Advertised path-id 1, VPN AF advertised path-id 1
Path type: redist, path is valid, is best path, no labeled nexthop, is extd
AS-Path: NONE, path locally originated
0.0.0.0 (metric 0) from 0.0.0.0 (44.192.0.2)
Origin incomplete, MED 0, localpref 100, weight 32768
Community: blackhole
Extcommunity: RT:23456:1000100

VRF advertise information:
Path-id 1 not advertised to any peer

VPN AF advertise information:
Path-id 1 not advertised to any peer

switch# sh bgp l2 e 192.0.2.0
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 192.0.2.53:4 (L3VNI 1000100)
BGP routing table entry for [5]:[0]:[0]:[24]:[192.0.2.0]/224, version 5
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn
Multipath: eBGP iBGP

Advertised path-id 1
Path type: local, path is valid, is best path, no labeled nexthop, has esi_gw
Gateway IP: 0.0.0.0
AS-Path: NONE, path locally originated
192.0.2.53 (metric 0) from 0.0.0.0 (192.0.2.53)
Origin incomplete, MED 0, localpref 100, weight 32768
Received label 1000100
Community: blackhole
Extcommunity: RT:23456:1000100 ENCAP:8 Router MAC:0476.b0f0.8157

Path-id 1 advertised to peers:
192.0.2.1
```

```
switch# sh bgp l2 e 192.0.2.0
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 192.0.2.53:4
BGP routing table entry for [5]:[0]:[0]:[24]:[192.0.2.0]/224, version 2
```

```
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn, is not in HW
Multipath: eBGP iBGP

Advertised path-id 1
Path type: external, path is valid, is best path, no labeled nexthop, has esi_gw
Imported to 2 destination(s)
Imported paths list: 100 L3-1000100
Gateway IP: 0.0.0.0
AS-Path: 4241653625 , path sourced external to AS
192.0.2.53 (metric 2) from 198.51.100.1 (192.0.2.53)
Origin incomplete, MED 0, localpref 100, weight 0
Received label 1000100
Community: blackhole
Extcommunity: RT:11000:1000100 Route-Import:192.0.2.53:100
Source AS:4241653625:0 SOO:50529024:00000000 ENCAP:8
Router MAC:0476.b0f0.8157
Path-id 1 not advertised to any peer

switch# show bgp ipv4 uni 192.0.2.0 vrf 100
BGP routing table information for VRF 100, address family IPv4 Unicast
BGP routing table entry for 192.0.2.0/24, version 3
Paths: (1 available, best #1)
Flags: (0x8008001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in
HW
vpn: version 3, (0x00000000100002) on xmit-list

Advertised path-id 1, VPN AF advertised path-id 1
Path type: external, path is valid, is best path, no labeled nexthop, in rib, has esi_gw

Imported from 192.0.2.53:4:[5]:[0]:[0]:[24]:[192.0.2.0]/224
AS-Path: 4241653625 , path sourced external to AS
192.0.2.53 (metric 2) from 198.51.100.1 (192.0.2.53)
Origin incomplete, MED 0, localpref 100, weight 0
Received label 1000100
Community: blackhole
Extcommunity: RT:11000:1000100 Route-Import:192.0.2.53:100
Source AS:4241653625:0 SOO:50529024:00000000 ENCAP:8
Router MAC:0476.b0f0.8157

VRF advertise information:
Path-id 1 not advertised to any peer
```

When reviewing these commands, check for:

- Presence of "Community: blackhole" for null route validation
- Path type and advertised path ID to confirm route propagation
- VRF and VLAN identifiers matching your configuration