



Configuring ND Suppression

This chapter contains the following sections:

- [ND Suppression on the Overlay, on page 1](#)
- [Guidelines and Limitations for ND Suppression, on page 1](#)
- [Configuring ND Suppression, on page 2](#)
- [Verifying the ND Suppression Configuration, on page 4](#)

ND Suppression on the Overlay

Multicast Neighbor Solicitation packets from host to another host are flooded over the BGP/EVPN VXLAN Core when hosts are behind two different VXLAN peers.

The ND Suppression cache is built by:

- Snooping NS request in the hosts and populating the ND Suppression cache with source IP and MAC bindings in the request.
- Learning IPv6-Host or MAC address information through BGP EVPN MAC route advertisements.

With ND Suppression, for host to host communication behind two different VXLAN peers, if the remote host is not learned in the suppression cache initially, then NS packets are flooded over the BGP/EVPN VXLAN Core. However, once the ND Suppression cache on a switch S1 is populated with the remote host, any subsequent Neighbor Solicitation request packet for the remote host in the hosts behind S1 are proxied by the Switch S1 thereby preventing the flooding of Neighbor Solicitation packet over the BGP-EVPN/VXLAN core.

For ND Suppression cache scale values, see *Cisco Nexus 9000 Series NX-OS Verified Scalability Guide*.

Guidelines and Limitations for ND Suppression

ND suppression has the following configuration guidelines and limitations:

- Beginning with Cisco NX-OS Release 10.3(1)F, the Cisco Nexus 9300-X Cloud Scale switches supports the ND Suppression feature only on plain BGP EVPN.
- ND Suppression is not supported with BGP-EVPN feature variants like Multisite, Virtual MCT, IRB, Centralized Gateway, Firewall Clustering, vPC.

- For link-local addresses of hosts, ND Suppression is not supported and instead multicast NS for link local address of hosts are flooded over the core of BGP EVPN VXLAN network.
 - ND Suppression gets enabled on all VNIs on which suppress-arp is enabled.
 - ND Suppression CLI knob must be enabled only under the following conditions:
 - The suppress-arp must be enabled on a VNI and there must be an SVI associated with this VNI/VLAN. Also, this SVI must be in up state and must have both IPv4 and IPv6 address enabled.
 - ND Suppression will not work in the following conditions:
 - If SVI not present for the VLAN/VNI on which suppress-arp/suppress nd is enabled.
 - If SVI associated with VLAN VNI on which suppress-arp/suppress nd is enabled is down.
 - If SVI associated with VLAN/VNI on which suppress-arp/suppress nd is enabled has only IPv4 and no IPv6 address.
 - If SVI associated with VLAN/VNI on which suppress-arp/suppress nd is enabled has only IPv6 and no IPv4 address.
- In all the above conditions, host to host traffic can potentially be dropped.

- For ND Suppression VACL to work, increase the SUP TCAM size to 768 or above using the **hardware access-list team region sup-team 768** command.
- If the installed Cisco NX-OS switch does not support ND suppression, ensure that Anycast Gateway MAC addresses across sites are identical.

Configuring ND Suppression

This procedure describes how to enable/disable the ND suppression feature on the NVE interface.

Before you begin

Ensure that ARP suppression is enabled.

SUMMARY STEPS

1. **configure terminal**
2. **hardware access-list team region ing-sup 768**
3. **copy running-config startup-config**
4. **reload**
5. **configure terminal**
6. **interface nve 1**
7. **[no]suppress nd**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enters global configuration mode.
Step 2	hardware access-list tcam region ing-sup 768 Example: switch# hardware access-list tcam region ing-sup 768	Carves the Ingress SUP TCAM size to 768.
Step 3	copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.
Step 4	reload Example: switch# reload	Reloads the switch.
Step 5	configure terminal Example: switch# configure terminal	Enters global configuration mode.
Step 6	interface nve 1 Example: switch(config)# interface nve 1 switch(config-if-nve)#	Enters interface nve configuration mode.
Step 7	[no]suppress nd Example: switch(config-if-nve)# suppress nd	Configures ND Suppression for all ARP enabled VNIs. Option no disables the ND Suppression for all ARP enabled VNIs.



Note

- When global **suppress arp** command is configured, ND Suppression is enabled on all VNIs.
- When global **suppress arp** command is not configured and instead per VNI **suppress arp** command is configured, then ND Suppression is enabled on all VNIs on which ARP suppression is configured.
- When enabling suppress arp command on a vPC pair, ensure steps 1-4 on both peers are complete before enabling the feature.

Verifying the ND Suppression Configuration

To display the ND Suppression configuration information, enter one of the following commands:

Command	Purpose
show run nv overlay	Displays the ND suppression configuration status.
show nve vni	Displays whether the ND suppression config has been enabled for ARP enabled VNIs.
show nve internal export nve	Displays whether the ND suppression config has been enabled or not in SDB.
show nve internal export vni	Displays the ND suppression state per VNI in SDB.
show ipv6 nd suppression-cache detail command.	Displays the ICMPv6 cache entries that are present in local.
show ipv6 nd suppression-cache remote	Displays the ICMPv6 cache entries that are present in remote.
show ipv6 nd suppression-cache summary	Displays the IPv6 cache entries summary of both local and remote.
show ipv6 nd suppression-cache statistics	Displays the IPv6 ND suppression cache statistics.
show ipv6 nd suppression-cache vlan "vlan_id"	Displays the details of IPv6 ND Suppression cache entries for a particular VLAN.

The following example shows sample output for the **show run nv overlay** command:

```
switch(config-if-nve)# sh run nv overlay
!Command: show running-config nv overlay
!Running configuration last done at: Sat Mar 19 01:07:49 2022
!Time: Sat Mar 19 01:10:00 2022

version 10.2(3) Bios:version 07.68
feature nv overlay

vlan 101-110,200-203,500-501

interface nvel
  no shutdown
  host-reachability protocol bgp
  suppress nd
  global suppress-arp
```

The following example shows sample output for the **show nve vni** command:

```
switch(config-if-nve-vni)# sh nve vni
Codes: CP - Control Plane      DP - Data Plane
       UC - Unconfigured        SA - Suppress ARP
       S-ND Suppress ND
       SU - Suppress Unknown Unicast
       Xconn - Crossconnect
       MS-IR - Multisite Ingress Replication
       HYB - Hybrid IRB mode
```

Interface	VNI	Multicast-group	State	Mode	Type	[BD/VRF]	Flags
nve1	5000	239.2.0.2	Up	CP	L2	[500]	SA S-ND

The following example shows sample output for the **show nve internal export nve** command:

```
switch(config-if-nve-vni)# sh nve internal export nve

NVE Interface information.
+-----+
Interface: nve1, Admin State: Up,
    State: nve-intf-add-complete, Encap: vxlan
    Source interface: loopback3, VRF: default,
    Anycast-interface: <none>
    Mcast-routing src intf <none>
    Primary IP: 4.4.4.4, Secondary IP: 0.0.0.0,
    VNI-VRF: default, Allow-Src-Lpbk-Down: No,
    Advertise MAC route: No,
    Virtual-rMAC: 0000.0000.0000,
    Mcast-routing Primary IP: 0.0.0.0
    Suppress ND: 1
    Host-reachability: CP
    unknown-peer-forwarding-mode: disable
    VNI assignment mode: n/a
    Multisite bgw-if: <none> (ip: 0.0.0.0, admin/oper state: Down/Down)
        src-node-last-notify: None
        anycast-node-last-notify: None
        mcast-src-node-last-notify: None
        multi-src-node-last-notify: None
+-----+
```

The following example shows sample output for the **show nve internal export vni** command:

```
switch(config-if-nve-vni)# sh nve internal export vni

NVE VNI Information.
+-----+
VNI: 5000 [500] Mgroup: 239.2.0.2 Provision-State: vni-add-complete
    Primary: 4.4.4.4 Secondary: 0.0.0.0 SRC-VRF: default
    Encap: vxlan Repl-mode: Mcast
    Suppress ARP: SP Suppress ND: Enabled Mode: CP, VNI-VRF: <FALSE> [vrf-id 0] [vrf flags 0x0]
    Suppress Unknown-Unicast: FALSE
    X-connect : Disabled
    [VNI local configs] SA : TRUE, Mcast-group : TRUE, IR proto BGP: FALSE
    Config Src: CLI, VNI flags: 0x0
    Spine-AGW: Disabled, HYBRID: Disabled
    Multisite optimized IR: Disabled
    Multisite DCI Group Unknown Address
+-----+
```

The following example shows sample output for the **show ipv6 nd suppression-cache detail** command:

```
switch(config)# show ipv6 nd suppression-cache detail

Flags: + - Adjacencies synced via CFSoE
      L - Local Adjacency
      R - Remote Adjacency
      L2 - Learnt over L2 interface
      PS - Added via L2RIB, Peer Sync
      RO - Dervied from L2RIB Peer Sync Entry

IPv6 Address          Age       Mac Address     Vlan Physical-ifindex   Flags      Remote Vtep
```

Configuring ND Suppression

```
Addrs

172:11:1:1::51 00:00:18 acf2.c5f6.7641 11 Ethernet1/51      L
172:11:1:1::201 00:06:14 0000.0011.1111 11 (null)           R      30.100.1.1
172:11:1:1::101 00:06:14 74a0.2fld.d481 11 (null)           R      10.10.11.11
```

The following example shows sample output for the **show ipv6 nd suppression-cache local** command:

```
switch(config)# show ipv6 nd suppression-cache local

Flags: + - Adjacencies synced via CFSOE
      L - Local Adjacency
      R - Remote Adjacency
      L2 - Learnt over L2 interface

Ip Address      Age      Mac Address      Vlan Physical-ifindex      Flags
172:11:1:1::51 00:00:23 acf2.c5f6.7641 11 Ethernet1/51      L

The following example shows sample output for the show ipv6 nd suppression-cache remote command:
```

```
switch(config)# show ipv6 nd suppression-cache remote

Flags: + - Adjacencies synced via CFSOE
      L - Local Adjacency
      R - Remote Adjacency
      L2 - Learnt over L2 interface
      PS - Added via L2RIB, Peer Sync
      RO - Derived from L2RIB Peer Sync Entry

IPv6 Address      Age      Mac Address      Vlan Physical-ifindex      Flags      Remote Vtep
Addrs

172:11:1:1::201 00:06:24 0000.0011.1111 11 (null)           R      30.100.1.1
172:11:1:1::101 00:06:24 74a0.2fld.d481 11 (null)           R      10.10.11.11
```

The following example shows sample output for the **show ipv6 nd suppression-cache statistics** command:

```
switch(config)# show ipv6 nd suppression-cache statistics
```

```
ND packet statistics for suppression-cache
```

```
Suppressed:
```

```
Total: 1
L3 mode :      Requests 1, Replies 1
              Flood ND Probe 0
```

```
Received:
```

```
Total: 1
L3 mode:      NS 1, Non-local NA 0
              Non-local NS 0
```

```
Mobility Requests:
```

```
Total: 0
L3 mode:      Remote-to-local 0, Local-to-remote 0
              Remote-to-remote 0
```

```
RARP Signal Refresh: 0
```

```
ND suppression-cache Local entry statistics
Adds 3, Deletes 0
```

The following example shows sample output for the **show ipv6 nd suppression-cache summary** command:

```
switch(config)# show ipv6 nd suppression-cache summary
```

```
IPV6 ND suppression-cache Summary
Remote :2
Local :1
Total :3
```

The following example shows sample output for the **show ipv6 nd suppression-cache vlan "vlan_id"** command:

```
switch(config)# show ipv6 nd suppression-cache vlan 11
```

```
Flags: + - Adjacencies synced via CFSOE
      L - Local Adjacency
      R - Remote Adjacency
      L2 - Learnt over L2 interface
      PS - Added via L2RIB, Peer Sync
      RO - Derived from L2RIB Peer Sync Entry
```

IPv6 Address Addrs	Age	Mac Address	Vlan	Physical-ifindex	Flags	Remote Vtep
172:11:1:1::51	00:00:40	acf2.c5f6.7641	11	Ethernet1/51	L	
172:11:1:1::201	00:06:36	0000.0011.1111	11	(null)	R	30.100.1.1
172:11:1:1::101	00:06:36	74a0.2f1d.d481	11	(null)	R	10.10.11.11

