



# Configuring Microsoft Network Load Balancing (NLB)

---

This chapter describes how to configure the Microsoft Network Load Balancing (NLB) feature on Cisco NX-OS devices.

- [About Network Load Balancing \(NLB\), on page 1](#)
- [Guidelines and Limitations for NLB, on page 2](#)
- [Prerequisites for Microsoft Network Load Balancing \(NLB\), on page 3](#)
- [Multicast Mode, on page 3](#)
- [IGMP Multicast Mode, on page 4](#)
- [Unicast Mode NLB and OTV Configuration Considerations, on page 6](#)
- [Verifying the NLB Configuration, on page 6](#)

## About Network Load Balancing (NLB)

Network Load Balancing (NLB) technology is used to distribute client requests across a set of servers. There are three primary modes of NLB: unicast, multicast, and Internet Group Management Protocol (IGMP) multicast:

- **Unicast mode** assigns the cluster a virtual IP and virtual MAC address. This method relies on unknown unicast flooding. Because the virtual MAC address is not learned on any switchports, traffic that is destined to the virtual MAC address is flooded within the VLAN. This means that all clustered servers receive traffic destined to the virtual MAC address. One downside to this method is that all devices in the VLAN receive this traffic. The only way to mitigate this behavior is to limit the NLB VLAN to only the NLB server interfaces in order to avoid flooding to interfaces that should receive the traffic.
- **Multicast mode** assigns a unicast IP address to a non-Internet Assigned Numbers Authority (IANA) multicast MAC address (03xx.xxxx.xxxx). IGMP snooping does not dynamically program this address, which results in flooding of the NLB traffic in the VLAN. Not requiring a PIM-enabled SVI or the IGMP snooping querier means that NLB works with custom non-IP multicast applications. For more information see, [Multicast Mode, on page 3](#)
- **IGMP multicast mode** assigns the cluster a virtual unicast IP address and a virtual multicast MAC address within the IANA range (01:00:5E:XX:XX:XX). The clustered servers send IGMP joins for the configured multicast group, and thus the switch dynamically populates its IGMP snooping table to point toward the clustered servers, which prevents unicast flooding. See [IGMP Multicast Mode, on page 4](#) for configuration examples.

This section describes how to configure a Cisco Nexus 9000 series switches for multicast and IGMP multicast mode NLB. As previously referenced, multicast NLB requires that you have a unicast IP address that is mapped to a multicast MAC address.

- Static Address Resolution Protocol (ARP) multicast.
- MAC address to a unicast IP address, but the traffic to that IP address floods the VLAN.
- Configure the system to constrain these packets to only those interfaces that require them. You can use several methods to configure the system, each with pros and cons.




---

**Note** Release 9.2(1) or later is required for unicast mode NLB to exist at multiple sites across an Overlay Transport Virtualization (OTV) overlay. See the Unicast Mode NLB and OTV Configuration Consideration Transport Virtualization (OTV) overlay. See the [Unicast Mode NLB and OTV Configuration Considerations, on page 6](#) section for further information.

---

## Guidelines and Limitations for NLB

Network Load Balancing (NLB) has the following configuration guidelines and limitations:

- Beginning with Cisco NX-OS Release 9.3(5), Multicast NLB is supported on Cisco Nexus 9300-FX3 platform switches.
- Beginning with Cisco NX-OS Release 10.2(1q)F, Multicast NLB is supported on Cisco Nexus N9KC9332D-GX2B platform switches.
- Multicast NLB is supported on Cisco Nexus 9300-EX, Cisco Nexus 9300-FX, Nexus 9300-FX2 platform switches, Cisco Nexus 9500 platform switches with N9K-X9700-EX line cards, N9K-X9700-FX line cards, Cisco Nexus 9500 platform switches with N9K-C9500-FM-E fabric cards and N9K-C9500-FM-E2 fabric cards. Beginning with Cisco NX-OS Release 9.3(6), Multicast NLB is supported on Cisco Nexus 9300-GX platform switches.
  - Multicast NLB is not supported on the Cisco Nexus 9500 modules with N9K-C9508-FM-2.
  - Multicast NLB is not supported on the Cisco Nexus 9300 and 9364C switches.
  - L2 (switched multicast) and L3 (routed multicast) is not supported to, from or inside of a VLAN that is configured for multicast NLB. This includes link local multicast groups as well, thus control plane protocols that use these groups are not supported to be configured on these VLANs.
  - Note that HSRP and VRRP are not included in the above mentioned limitations.
- Flooding for Microsoft Network Load Balancing (NLB) unicast mode is not supported on Cisco Nexus 9000 switches. A static ARP entry must be configured to map the NLB virtual IP address to the NLB virtual MAC address. Furthermore, a static MAC address entry must be configured to map the NLB virtual MAC address to a specific egress interface.
- FEX HIF interfaces cannot receive a multicast NLB flow.
- If none of the ports in the interface set is UP, the traffic floods to all ports in the VLAN.

- L2 and L3 regular multicast is not supported to, from or inside the NLB VLAN.
- NLB traffic that enters the NLB VLAN may be looped back to the source interface. This looped back NLB traffic time-to-live (TTL) is decremented even though it is intra-VLAN.
- Multicast Mode - If servers/firewalls move, the administrator must update the static multicast MAC table configuration.
- IGMP Multicast Mode - If servers/firewalls move, the administrator must update the static-group configuration.
- NLB in the unicast, multicast, and IGMP multicast modes is not supported on Cisco Nexus 9000 Series based VXLAN VTEPs. The work around is to move the NLB cluster behind intermediary device (which supports NLB in the respective mode) and inject the cluster IP address as external prefix into VXLAN fabric.
- Do not configure **feature nv overlay** and **feature nlb** together.

## Prerequisites for Microsoft Network Load Balancing (NLB)

Microsoft Network Load Balancing (NLB) has the following prerequisites:

- You are logged into the device.
- For global configuration commands, you are in the correct virtual routing and forwarding (VRF) mode. The default configuration mode shown in the examples in this chapter applies to the default VRF.
- Multicast NLB requires that you have a unicast IP address mapped to a multicast MAC address.

## Multicast Mode

Multicast mode assigns a unicast IP address to a non-Internet Assigned Numbers Authority (IANA) multicast MAC address (03xx.xxxx.xxxx). IGMP snooping does not dynamically program this address, which results in flooding of the NLB traffic in the VLAN. Refer to Option 2A for an example of how to configure for this mode. The following example shows how to configure for IGMP Multicast Mode:

### Example 1: Static ARP + MAC-based L2 Multicast Lookups + Static Joins + Non-IP Multicast MAC

This option does not require a PIM-enabled SVI or the IGMP snooping querier; works with non-IP multicast applications (custom applications).



---

**Note** The **hardware profile multicast nlb** CLI must be enabled on the switch to support Multicast Mode.

---

1. Configure a static ARP entry that maps the unicast IP address to a multicast MAC address, but this time in the non-IP address multicast range:

```
interface Vlan10
no shutdown
ip address 10.1.2.1/24
ip arp 10.1.2.200 03bf.0000.1111
```

2. Enable MAC-based Layer 2 multicast lookups in the VLAN (by default, multicast lookups are based on the destination multicast IP address):



**Note** You must use MAC-based lookups in VLANs where you want to constrain IP address unicast packets with multicast MAC addresses.

```
vlan configuration 10
layer-2 multicast lookup mac
```

3. Configure static MAC address-table entries that point to the interfaces connected to the NLB server and any redundant interface:

```
mac address-table multicast 03bf.0000.1111 vlan 10 interface Ethernet8/2
mac address-table multicast 03bf.0000.1111 vlan 10 interface Ethernet8/4
mac address-table multicast 03bf.0000.1111 vlan 10 interface Ethernet8/7
```

## IGMP Multicast Mode

**IGMP multicast mode** assigns the cluster a virtual unicast IP address and a virtual multicast MAC address within the IANA range (01:00:5E:XX:XX:XX). The clustered servers send IGMP joins for the configured multicast group, and thus the switch dynamically populates its IGMP snooping table to point toward the clustered servers, which prevents unicast flooding. The following describes three examples of how to configure for IGMP Multicast Mode:

### Option 1: Static ARP + MAC-based L2 Multicast Lookups + Dynamic Joins

This option allows servers and firewalls to dynamically join or leave the corresponding group; enables or disables reception of the target traffic (for example, maintenance mode).



**Note** The **hardware profile multicast nlb** CLI must be enabled on the switch to support IGMP Multicast Mode.

1. Configure a static ARP entry that maps the unicast IP address to a multicast MAC address in the IP address multicast range on a Protocol Independent Multicast (PIM)-enabled interface:

```
interface Vlan10
no shutdown
ip address 10.1.2.1/24
ip pim sparse-mode
ip arp 10.1.2.200 0100.5E01.0101
```

2. Enable MAC-based Layer 2 multicast lookups in the VLAN (by default, multicast lookups are based on the destination multicast IP address):

```
vlan configuration 10
layer-2 multicast lookup mac
```

### Option 2: Static ARP + MAC-based L2 Multicast Lookups + Dynamic Joins with IGMP Snooping Querier

Option 2 does not require PIM-enabled SVI and allows servers and firewalls to dynamically join or leave the corresponding group; enables or disables reception of the target traffic (for example, maintenance mode).



**Note** The **hardware profile multicast nlb** CLI must be enabled on the switch to support IGMP Multicast Mode.

1. Configure a static ARP entry like in Option 1, but do not enable PIM on the switch virtual interface (SVI).

```
interface Vlan10
no shutdown
ip address 10.1.2.1/24
ip arp 10.1.2.200 0100.5E01.0101
```

2. Enable MAC-based Layer 2 multicast lookups in the VLAN, and enable the Internet Group Management Protocol (IGMP) snooping querier:

```
vlan configuration 10
ip igmp snooping querier 10.1.1.254
layer-2 multicast lookup mac
```

### Option 3: Static ARP + MAC-based L2 Multicast Lookups + Static Joins + IP Multicast MAC

Option three does not require a PIM-enabled SVI or the IGMP snooping querier.



**Note** The **hardware profile multicast nlb** CLI must be enabled on the switch to support IGMP Multicast Mode.

1. Configure a static ARP entry that maps the unicast IP address to a multicast MAC address in the IP address multicast range:

```
interface Vlan10
no shutdown
ip address 10.1.2.1/24
ip arp 10.1.2.200 0100.5E01.0101
```

2. Enable MAC-based Layer 2 multicast lookups in the VLAN (by default, multicast lookups are based on the destination multicast IP address):

```
vlan configuration 10
layer-2 multicast lookup mac
```

You must use MAC-based lookups in VLANs where you want to constrain IP address unicast packets with multicast MAC addresses.

3. Configure static IGMP snooping group entries for the interfaces connected to the NLB server that needs the traffic:

```
vlan configuration 10
ip igmp snooping static-group 239.1.1.1 interface Ethernet8/2
ip igmp snooping static-group 239.1.1.1 interface Ethernet8/4
ip igmp snooping static-group 239.1.1.1 interface Ethernet8/7
```

**If you are using multicast NLB where traffic is ingress and egress same port-channel and you have members of that port-channel on different modules, perform the following steps:**

1. Execute the following commands in global mode:

```
hardware profile multicast nlb
hardware profile multicast nlb port-Channel
clear ip igmp snooping groups (For ingress and egress VLANs)
```

2. Continue with the configuration steps of the respective options as mentioned above.

## Unicast Mode NLB and OTV Configuration Considerations

OTV allows the advertising of MAC addresses between the OTV edge devices, as well as the mapping of MAC address destinations to IP next hops that are reachable through the network transport. The consequence is that the OTV edge device starts to behave like a router instead of a Layer 2 bridge, because it forwards Layer 2 traffic across the overlay if it has previously received information on how to reach that remote MAC destination.



**Note** Multicast and IGMP multicast mode are treated as broadcasts over the OTV overlay. They work across OTV without additional configuration.

When the OTV edge device receives a frame destined to a MAC across the overlay, by default it performs a Layer 2 lookup in the MAC table. Because it does not have information for the MAC, the traffic is flooded out the internal interfaces (because they behave as regular Ethernet interfaces) but not via the overlay. In releases earlier than 6.2(2), unicast mode NLB only works if the servers are on a single side of the OTV overlay. The OTV VDC at the site that these servers is placed is configured in this manner:

```
mac address-table static 02bf.0000.2222 vlan 10 interface <internal-interface>
```

Unicast mode NLB servers can exist on both sides of the OTV overlay. This is done through use of the selective unicast flood command on the OTV VDCs at all sites where the server exists:

```
otv flood mac 02bf.0000.2222 vlan 10
```



**Note** When you use NLB for an OTV extended VLAN, you must disable ARP ND cache "no otv suppress-arp-nd" on the Overlay.

## Verifying the NLB Configuration

To display the NLB configuration information, perform one of the following tasks.

Command	Description
<code>show ip arp virtual-address</code>	Displays the ARP table.
<code>show ip igmp snooping groups [source [group]   group [source]] [vlan vlan-id] [detail]</code>	Displays IGMP snooping information about groups by VLAN.
<code>show ip igmp snooping mac-oif vlan vlan-id</code>	Displays IGMP snooping static MAC addresses.