Microsoft Network Load Balancing on Nexus 7000 Configuration Example

Document ID: 116150

Contributed by Mahir Rabbi and Clark Dyson, Cisco TAC Engineers.
Feb 06, 2014

Contents

Introduction

This document describes how to configure Microsoft Network Load Balancing (NLB) on Nexus 7000.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on Cisco NX−OS Software, Release 5.2(x) or later.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Configure

Note: Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

Overview of 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 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. Refer to Option 2A for an example of how to configure for this 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 towards the clustered servers, which prevents unicast flooding. Refer to Option 1, Option 1A, and Option 2 for examples of how to configure for this mode.

This document covers how to configure Nexus 7000 series switches for multicast and IGMP multicast mode NLB. As previously referenced, multicast NLB requires that you have a unicast IP address mapped to a multicast MAC address. If you have a Catalyst switch, you can follow the configuration in Catalyst Switches for Microsoft Network Load Balancing Configuration Example. The Nexus 7000 follows the same concept, but the configurations are different.

The Nexus 7000 needs to be able to run Release 5.2(x) or later in order to perform these configurations:

- In NX−OS Release 4.2 and later, you can map a static Address Resolution Protocol (ARP) multicast MAC address to a unicast IP address, but the traffic to that IP address floods the VLAN.
- In NX−OS Release 5.2 and later, you can 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 6.2(2) 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 section for further information.

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

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
```

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

When hosts (load balancing [LB] servers or firewalls) join an IP address multicast group that corresponds to the MAC address of the ARP entry, the system installs a snooping entry that constrains traffic destined to that group's MAC address to only those ports where a join was received.

Pros of Option 1: allows servers/firewalls to dynamically join/leave the corresponding group; enables/disables reception of the target traffic (for example, maintenance mode).

Cons of Option 1: constraint can only occur if at least one server/firewall is joined to the group address; if the last device leaves the group, the traffic floods to all ports in the VLAN.

Option 1A: Static ARP + MAC−based L2 Multicast Lookups + Dynamic Joins with IGMP Snooping Querier

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
```

Pros of Option 1A: does not require PIM−enabled SVI. Otherwise, the pros are the same as in Option 1.

Cons of Option 1A: same as in Option 1.

Option 2: Static ARP + MAC−based L2 Multicast Lookups + Static Joins + IP Multicast MAC

1. In this option, you again 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
```

Pros of Option 2: does not require a PIM−enabled SVI or the IGMP snooping querier.

Cons of Option 2: constraint can only occur if at least one server/firewall port is in the UP state (link up); if none of the ports in the static−group interface set is UP, the traffic floods to all ports in the VLAN. If servers/firewalls move, the administrator must update the static−group configuration.

Option 2A: Static ARP + MAC−based L2 Multicast Lookups + Static Joins + Non−IP Multicast MAC

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):

```
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 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
```

Note: A static MAC entry should be applied on any device that shares the NLB VLAN that points to the server and redundant links. The specific configuration varies for each platform.

Pros of Option 2A: does not require a PIM−enabled SVI or the IGMP snooping querier; works with non−IP multicast applications (custom applications).

Cons of Option 2A: constraint can only occur if at least one server/firewall port is in the UP state (link up); if none of the ports in the interface set is UP, the traffic floods to all ports in the VLAN. If servers/firewalls move, the administrator must update the static multicast MAC table configuration.

Unicast Mode NLB and OTV Configuration Considerations

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

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.

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>
```

In Release 6.2(2) and later, 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
```

Caveats

There are a few caveats related to NLB on the Nexus 7000:

- Cisco Bug ID CSCtw73595: IGMP mode floods routed traffic on M1 and M2 modules. This is a hardware limitation.
- Cisco Bug ID CSCtv00148: Multicast mode floods routed traffic. This issue is fixed in Releases 5.2(3a), 6.0(2), and later.

Supported Platforms

This document was written specifically for the Nexus 7000. However, only these NX−OS platforms currently have support for NLB:

- Nexus 7000
- Nexus 6000
- Nexus 5000
- Nexus 9500 (unicast only; see Cisco Bug ID CSCup90853)

Here is some additional information in regards to NLB support:

- Support for NLB on the 3548 Series platform is tracked by Cisco Bug ID CSCup43205.
- Support for NLB on the 30xx and 31xx Series platforms is tracked by Cisco Bug IDs CSCup92860 and CSCui82585.
- Support for NLB on the Nexus 9300/9500 Series platforms is tracked by Cisco Bug IDs CSCuq14783 and CSCuq03168.

Verify

Note: The Output Interpreter Tool (registered customers only) supports certain ***show*** commands. Use the Output Interpreter Tool in order to view an analysis of ***show*** command output.

Static ARP can be verified with this command:

```
show ip arp <Virtual IP>
```

IGMP snooping entries can be verified with this command:

```
show ip igmp snooping groups <multicast group> vlan <VLAN>
```

Static MAC address table entries can be verified with this command:

```
show ip igmp snooping mac-oif vlan <VLAN>
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Updated: Feb 06, 2014 Document ID: 116150