Configure Catalyst Switches for Microsoft NLB

Contents

Introduction

This document describes how to configure Cisco Catalyst switches to interact with Microsoft Network Load Balancing (NLB).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based and configured on these software and hardware versions:

- Catalyst 6500 Sup2T switch running Cisco IOS Software 15.1(1)SY1
- Catalyst 4948 switch running Cisco IOS Software 15.0(2)SG7
- Microsoft Windows Servers

Note: Consult the appropriate configuration guide for the commands that are used in order to

enable these features on other Cisco platforms.

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.

Background Information

The NLB technology can be used in order to distribute client requests across a set of servers. In order to ensure that clients always experience acceptable performance levels, Microsoft Windows NLB provides the ability to add additional servers so that you can scale out stateless applications, such as IIS-based web servers, as client load increases. In addition, it reduces downtime that is caused by servers that malfunction.

Microsoft Windows NLB is a clustering technology that is offered as a part of all Windows 2000 Server and Windows 2003 Server family operating systems. It provides a single virtual IP address for all clients as the destination IP address for the entire cluster.

You can configure NLB to work in one of these three modes:

- Unicast mode
- Multicast mode
- Internet Group Management Protocol (IGMP) mode

Unicast Mode

Here are some notes about the use of NLB in Unicast mode:

- In Unicast mode, NLB replaces the actual Media Access Control (MAC) address of each server in the cluster with a common NLB MAC address. When all of the servers in the cluster have the same MAC address, all of the packets that are forwarded to that address are sent to all of the members in the cluster. The NLB creates a fictitious MAC address and assigns it to each server in the NLB cluster. The NLB assigns each NLB server a different fictitious MAC address, based on the host ID of the member. This address appears in the Ethernet frame header.

- The MAC address is used in the Address Resolution Protocol (ARP) header, not the Ethernet header. The switch uses the MAC address in the Ethernet header, not the ARP header. This causes an issue when a packet is sent to the NLB cluster with the destination MAC address as the cluster MAC address 00-bf-ac-10-00-01. The switch views the Content Addressable Memory (CAM) table for the MAC address 00-bf-ac-10-00-01, and since there is no port registered with the NLB cluster MAC address 00-bf-ac-10-00-01, the frame is delivered to all of the switch ports. This introduces *unicast flooding.* In order to avoid flooding, Cisco recommends that you use a dedicated VLAN for NLB so that the flooding is constrained.

Multicast Mode

Here are some notes about the use of NLB in Multicast mode:

- In Multicast mode, the system administrator clicks the Multicast button in the Microsoft NLB configuration GUI. This choice instructs the cluster members to respond to the ARPs for their virtual address with the use of a multicast MAC address, such as 0300.5e01.0101.

- The ARP process does not complete for multicast MAC addresses (this breaks RFC 1812). A static MAC address is required in order to reach the cluster outside of the local subnet.

- The virtual IP address is 10.100.1.99 and the multicast MAC address is 0300.5e01.0101. Enter this command in order to populate the ARP table statically:

```
arp 10.100.1.99 0300.5e01.0101
```

- Since the incoming packets have a unicast destination IP address and multicast destination MAC the Cisco device ignores this entry and unicast floods each cluster-bound packets. In order to avoid this flooding, insert a static **mac-address-table** entry as given below in order to switch cluster-bound packets in hardware.

```
mac address-table static 0300.5e01.0101 vlan 200 interface TenGigabitEthernet1/4
TenGigabitEthernet1/5 disable-snooping
```

> **Note**: When you statically map a MAC address to multiple ports, it is only supported by the software on the Cisco Catalyst 4500 Series switch. Also, the use of this configuration on the Catalyst 4500 Series switch might cause high CPU. In order to avoid this problem, you can isolate the NLB to a specific VLAN, add only the static ARP entries, and allow flooding on that VLAN.

> **Note**: For Cisco Catalyst 6000/6500 Series switches, you must add the disable-snooping parameter to constrain traffic to the specified ports only. When you configure a static connection, enter the disable-snooping keyword to prevent multicast traffic addressed to the statically configured multicast MAC address from also being sent to other ports in the same VLAN. (This command is not required on other platforms).

IGMP Mode

Here are some notes about the use of NLB in IGMP mode:

- The use of NLB in IGMP mode requires the least amount of manual configuration. The virtual MAC address falls within the Internet Assigned Numbers Authority (IANA) range and starts with **0100.5exx.xxxx**. Since the MAC address now conforms to IANA specifications, the Cisco switches can dynamically program the MAC address with the use of IGMP snooping. This removes the need to manually program the MAC address to the port maps that are required in Multicast mode in order to prevent flooding to the VLAN.

- The IGMP snooping programs the virtual MAC address for you once the switch receives a membership report from a member in the cluster. An Mrouter port must also be programmed for the NLB VLAN with the use of either Protocol Independent Multicast (PIM) or the IGMP querier feature.

- Since the virtual IP address uses a multicast MAC address, it is unreachable outside of the

local subnet. In order to address this, you must configure a static ARP entry on each device with a Layer 3 (L3) interface in the cluster VLAN. Complete this in the same fashion as with Multicast mode. For example, if the virtual IP address is 10.100.1.99 and the multicast MAC address is 0100.5e01.0101, use this command in order to populate the ARP table statically:

```
arp 10.100.1.99 0100.5e01.0101
```

IGMP Mode Caveats

Here are important notes to keep in mind when you use NLB in IGMP mode:

Note: As tracked by Cisco bug ID CSCsw72680 you cannot use PIM on the NLB VLANSwitch Virtual Interface (SVI) with certain versions of code. View the bug details for the code releases that address this issue or use the IGMP snooping querier feature.

Note: As tracked by Cisco bug ID CSCsy62709 packets are duplicated for all of the traffic thats is routed to the NLB servers in IGMP mode. View the bug details for the affected code versions.

Note: Due to a hardware limitation that is tracked by Cisco nug ID CSCug49149 NLB traffic cannot be sent across a Distributed Etherchannel on the same 6708 line card when switch is running in either PFC3B or PFC3C mode. The port-channel must be cabled so all member links are on the same forwarding engine.

Network Diagram

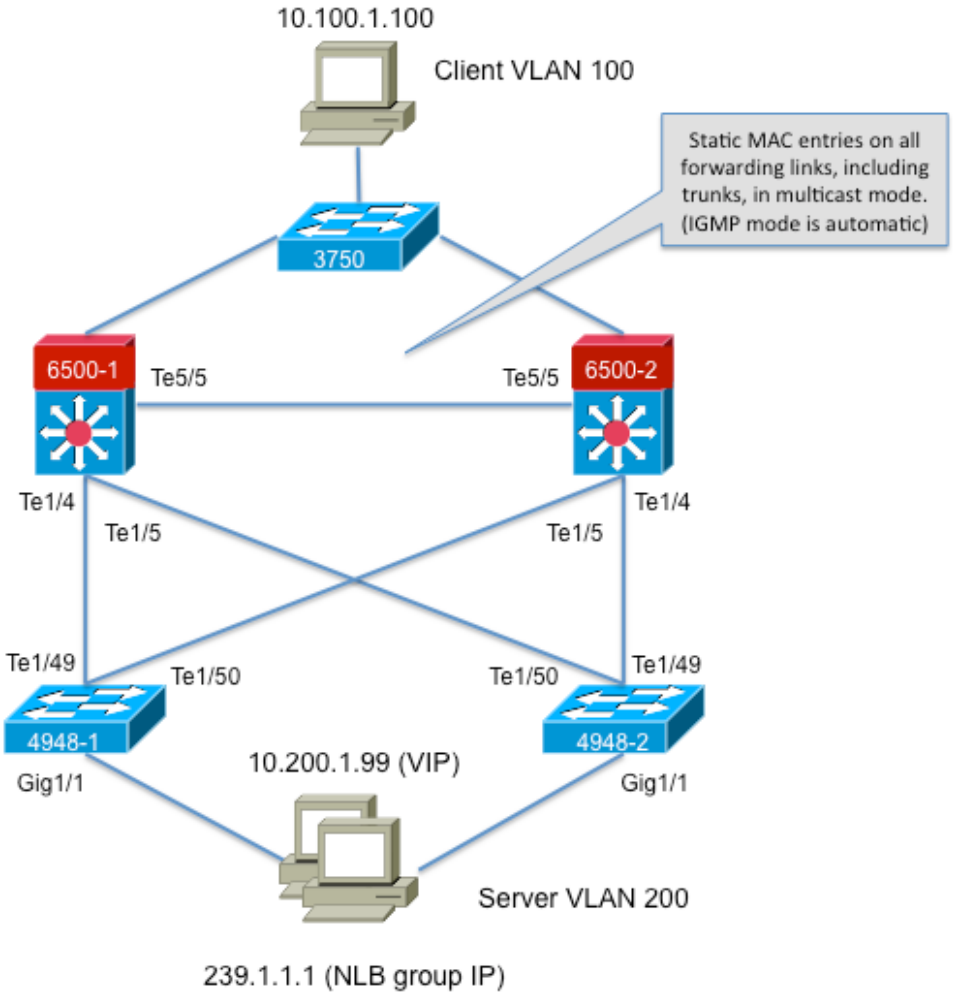

Configure

This section describes how to configure NLB for the Cisco Catalyst 6500 and 4948 Series platforms that run in Multicast or IGMP mode.

Configuration for Multicast Mode

This section describes how to configure NLB for the Cisco Catalyst 6500 and 4948 Series platforms that run in Multicast mode:

```
6500-1#show running-config
Building configuration...
!
hostname 6500-1
!
boot system flash disk0:s2t54-adventerprisek9-mz.SPA.151-1.SY1
!
interface TenGigabitEthernet1/4
switchport
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
interface TenGigabitEthernet1/5
switchport
switchport trunk allowed vlan 1,100,200
```

```
switchport mode trunk
!
interface Vlan100
ip address 10.100.1.1 255.255.255.0
!
!
interface Vlan200
ip address 10.200.1.1 255.255.255.0
!
!
arp 10.100.1.88 0300.5e01.0101 ARPA
!
!
mac address-table static 0300.5e01.0101 vlan 200 interface TenGigabitEthernet1/4
TenGigabitEthernet1/5 TenGigabitEthernet5/5
!
```

Here are some important notes about this configuration:

- The**interface Vlan100 ip address**value configures the user VLAN.
- The**interface Vlan200 ip address**value configures the NLB cluster VLAN. It is important that you configure the default gateway of the Microsoft Server to this address.
- The**arp 10.100.1.88 0300.5e01.0101 ARPA**includes all of the L3 interfaces in the VLAN and is the virtual IP address of the NLB cluster servers.
- The**mac address-table static 0300.5e01.0101 vlan 200 interface**creates a static MAC entry to port mapping in the switch for the multicast virtual MAC address.

Note: Ensure that you use Multicast mode on the NLB cluster. Cisco recommends that you do not use multicast MAC addresses that begin with 01 because they are known to have a conflict with the IGMP setup.

```
4948-1#show running-config
Building configuration...
!
hostname 4948-1
!
boot system bootflash:cat4500-entservices-mz.150-2.SG7
!
interface GigabitEthernet1/1
switchport access vlan 200
!
interface TenGigabitEthernet1/49
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
interface TenGigabitEthernet1/50
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
mac address-table static 0300.5e01.0101 vlan 200 interface Gi1/1 Te1/49 Te1/50
!
!
end
```

Note: The **mac address-table static 0300.5e01.0101 vlan 200 interface** creates a static entry in the switch for the multicast virtual MAC address. It is important to remember that all of the trunk interfaces that carry NLB traffic between the switches must be added. Once a

static MAC address is defined, flooding is constrained. If you forget to include an interface, the NLB cluster breaks.

Configuration for IGMP Mode

This section describes how to configure NLB for the Cisco Catalyst 6500 and 4948 Series platforms that run in IGMP mode

```
6500-1#show running-config
Building configuration...
!
hostname 6500-1
!
boot system flash disk0:s2t54-adventerprisek9-mz.SPA.151-1.SY1
!
ip igmp snooping querier
!
vlan configuration 1,100
no ip igmp snooping querier
!
vlan configuration 200
ip igmp snooping querier address 10.200.1.1
!
interface TenGigabitEthernet1/4
switchport
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
interface TenGigabitEthernet1/5
switchport
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
interface Vlan100
ip address 10.100.1.1 255.255.255.0
!
interface Vlan200
ip address 10.200.1.1 255.255.255.0
!
arp 10.100.1.99 0100.5e01.0101 ARPA
!
end
```

Here are some important notes about this configuration:

- The**ip igmp snooping querier**enables the snooping querier feature.
- The**ip igmp snooping querier address 10.200.1.1**configures the snooping querier for the NLB VLAN.
- The user VLAN is**interface Vlan100**.
- The NLB cluster VLAN is**interface Vlan200**. It is important that you configure the default gateway of the Microsoft Server to this address (**ip address 10.200.1.1 255.255.255.0**).
- The**arp 10.100.1.99 0100.5e01.0101 ARPA**is the virtual IP address of the NLB cluster servers. The static ARP must be on all of the L3 interfaces in the VLAN.

```
4948-1#show running-config
Building configuration...
```

```
!
hostname 4948-1
!
boot system bootflash:cat4500-entservices-mz.150-2.SG7
!
interface GigabitEthernet1/1
switchport access vlan 200
!
interface TenGigabitEthernet1/49
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
interface TenGigabitEthernet1/50
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,200
switchport mode trunk
!
end
```

> **Note**: There is no need to configure static entries, as IGMP snooping does this dynamically in this mode. Also, no special configuration for this mode is required on the downstream Layer 2 (L2) switches.

Verify

Use this section to confirm that your configuration works properly.

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

Multicast Mode Verification

Enter the**show ip arp**command in order to view the ARP cache:

```
6500-1#sh ip arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.100.1.99 - 0300.5e01.0101 ARPA
```

Enter the **show mac address-table static** command in order to view a specific MAC address table static and dynamic entry or the MAC address table static and dynamic entries on a specific interface or VLAN:

```
6500-1#show mac address-table static add 0300.5e01.0101

vlan mac address type learn age ports
----+---------------+-------+-----+------+-----------------
200  0300.5e01.0101  static  No     -     Te1/4 Te1/5 Te5/5

4948-1#show mac address-table static add 0300.5e01.0101

Multicast Entries
vlan mac address type ports
------+---------------+-------+--------------------
200    0300.5e01.0101  static  Gi1/1,Te1/49,Te1/50
```

IGMP Mode Verification

Enter the**show ip arp**command in order to view the ARP cache:

```
6500-1#show ip arp

Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.100.1.99 - 0100.5e01.0101 ARPA
```

Enter the**show ip igmp snooping mrouter**in order to view the Mrouter port that is programmed by the queries received from the upstream snooping querier:

```
4948-1#show ip igmp snooping mrouter

Vlan ports
---- -----
200 Te1/49(dynamic)
```

Enter the **show mac address-table multicast igmp-snooping** in order to view the dynamically-added MAC address that is learned from IGMP snooping and the member ports:

```
4948-1#show mac address-table multicast igmp-snooping

Multicast Entries
vlan mac address type ports
------+---------------+------+-------------
200    0100.5e01.0101  igmp   Gi1/1,Te1/49
```

Enter the**show ip igmp snooping groups**in order to view the port list of cluster members that joined the multicast group:

```
4948-1#show ip igmp snooping groups

Vlan Group Version Port List
----------------------------------------
200     239.1.1.1    v2           Gi1/1
```

Troubleshoot

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