Catalyst Switches for Microsoft Network Load Balancing Configuration Example

TAC **Document ID: 107995**

Contributed by Shashank Singh, Cisco TAC Engineer.
Dec 19, 2013

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 on these software and hardware versions:

- Catalyst 6500 switch which runs Cisco IOS® Software
- Catalyst 4500 switch which runs Cisco IOS Software
- Catalyst 3550 switch which runs Cisco IOS Software
- Catalyst 3560 switch which runs Cisco IOS Software
- Catalyst 3750 switch which runs Cisco IOS Software
- Microsoft Windows 2000/2003 Servers

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

NLB technology can be used to distribute client requests across a set of servers. In order to make sure clients always experience acceptable performance levels, Windows NLB is often used to ensure that you can add additional servers to scale out stateless applications, such as IIS−based web servers, as client load increases. In addition, it reduces downtime caused by servers that malfunction. End users will never know that a particular member server in the Windows NLB is or has been down.

Network Load Balancing is a clustering technology offered by Microsoft as part of all Windows 2000 Server and Windows Server 2003 family operating systems. NLB uses a distributed algorithm to load balance network traffic across a number of servers.

NLB bundles the servers into one multicast group and tries to use the standard multicast IP and MAC address. At the same time, it provides a single virtual IP for all clients as the destination IP, which means servers join the same multicast group, and the clients will not know anything about it. They use normal unicast access to the VIP.

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

- Unicast Mode
- Multicast Mode

Unicast Mode

The NLB default setting is unicast mode. In unicast mode, NLB replaces the actual MAC address of each server in the cluster to a common NLB MAC address. When all the servers in the cluster have the same MAC address, all packets forwarded to that address are sent to all members of the cluster. However, a problem with this configuration is when the servers NLB cluster is connected to the same switch. You cannot have two ports on the switch register the same MAC address. In order to solve this problem, NLB masks the cluster MAC address. The switch looks at the source MAC address in the Ethernet frame header in order to learn which MAC addresses are associated with its ports. NLB creates a bogus MAC address and assigns that bogus MAC address to each server in the NLB cluster. NLB assigns each NLB server a different bogus MAC address based on the host ID of the member. This address appears in the Ethernet frame header.

For example, the NLB cluster MAC address is 00−bf−ac−10−00−01. NLB in unicast mode takes the cluster MAC address and, for each cluster member, NLB changes the second octet so that it consists of the NLB member?s host ID. For example, server number 1 has the bogus MAC address 00−01−ac−10−00−01, host ID number 2 has the bogus MAC address 00−02−ac−10−00−01, so on. If a unique MAC address is registered on each switch port, packets are not delivered to all members of the array. Rather, packets should still be sent to the individual switch ports based on the MAC address assigned to that port. In order to have frames delivered to all members of the NLB cluster when each switch port connected to an NLB cluster member registers a different MAC address, Address Resolution Protocol (ARP) broadcast is used. When the router sends an ARP request for the MAC address of the virtual IP address, the reply contains an ARP header with the actual NLB cluster MAC address 00−bf−ac−10−00−01, as per the example given above and not the bogus MAC address.

The clients use the MAC address in the ARP header, not the Ethernet header. The switch uses the MAC address in the Ethernet header, not the ARP header. The issue is when a client sends a packet to the NLB cluster with the destination MAC address as cluster MAC address 00−bf−ac−10−00−01. The switch looks at the Content Addressable Memory (CAM) table for the MAC address 00−bf−ac−10−00−01. Since there is no port registered with the NLB cluster MAC address 00−bf−ac−10−00−01, the frame is delivered to all switch ports. This introduces *switch flooding*. Switch flooding causes issues when a significant amount of traffic flows and also when there are other servers on the same switch. A solution to switch flooding is to put a simple hub in front of the NLB cluster members and then uplink the hub to a switch port. This solution does not even need to mask the NLB cluster MAC address because the single switch port connected to the hub learns the NLB cluster MAC address. This avoids the problem of two switch ports that register the same MAC address. When the client sends packets to the NLB cluster MAC address, the packets go directly to the switch port connected to the hub and then to the NLB cluster members.

Multicast Mode

Another solution is to use multicast mode in the MS NLB configuration GUI instead of unicast mode. In multicast mode, the system administrator clicks the Internet Group Management Protocol (IGMP) multicast button in the MS NLB configuration GUI. This choice instructs the cluster members to respond to ARPs for their virtual address with a multicast MAC address (for example, 0100.5e11.1111) and to send IGMP Membership Report packets. If IGMP snooping is enabled on the local switch, it snoops the IGMP packets that pass through it. In this way, when a client ARPs for the cluster?s virtual IP address, the cluster responds with multicast MAC (for example, 0100.5e11.1111). When the client sends the packet to 0100.5e11.1111, the local switch forwards the packet out each of the ports connected to the cluster members. In this case, there is no chance of flooding the ARP packet out of all the ports. The issue with the multicast mode is the virtual IP address becomes unreachable when accessed from outside the local subnet because Cisco devices do not accept an ARP reply for a unicast IP address that contains a multicast MAC address. So the MAC portion of the ARP entry shows as *incomplete*. (Enter the command ***show arp*** to view the output.) As there is no MAC portion in the ARP reply, the ARP entry never appears in the ARP table. It eventually quits ARPing and returns an ICMP Host unreachable message to the clients. In order to override this, use a static ARP entry to populate the ARP table as shown here. In theory, this allows the Cisco device to populate its MAC address table. For example, if the virtual IP address is 172.16.63.241 and the multicast MAC address is 0100.5e11.1111, enter this command in order to populate the ARP table statically:

```
arp 172.16.63.241 0100.5e11.1111
```

However, since the incoming packets have a unicast destination IP address and multicast destination MAC, the Cisco device ignores this entry and process−switches each cluster−bound packet. In order to avoid this process switching, insert a static MAC address table entry as shown here in order to switch cluster−bound packets in hardware.

```
mac−address−table static 0100.5e11.1111 vlan 200 interface fa2/3 fa2/4
```

Note: Statically mapping a MAC address to multiple ports is supported only in software on the Catalyst 4500 switch. This configuration might cause high CPU on the Catalyst 4500 switch.

Note: For Cisco Catalyst 6500 Series switches, you must add the *disable−snopping* parameter. For example:

mac−address−table static 0100.5e11.1111 vlan 200 interface fa2/3 fa2/4 disable−snooping

The ***disable−snooping*** parameter is essential and applicable only for Cisco Catalyst 6500 Series switches. Without this statement, the behavior is not affected.

Please note that on Cisco Catalyst 6500 Series switches the disable snooping option is available only for 0100.5exx.xxxx and 3333.xxxx.xxxx ranges of multicast MAC addresses.

Configure

In this section, you are presented with the information to configure the features described in this document.

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

Network Diagram

This document uses this network setup:

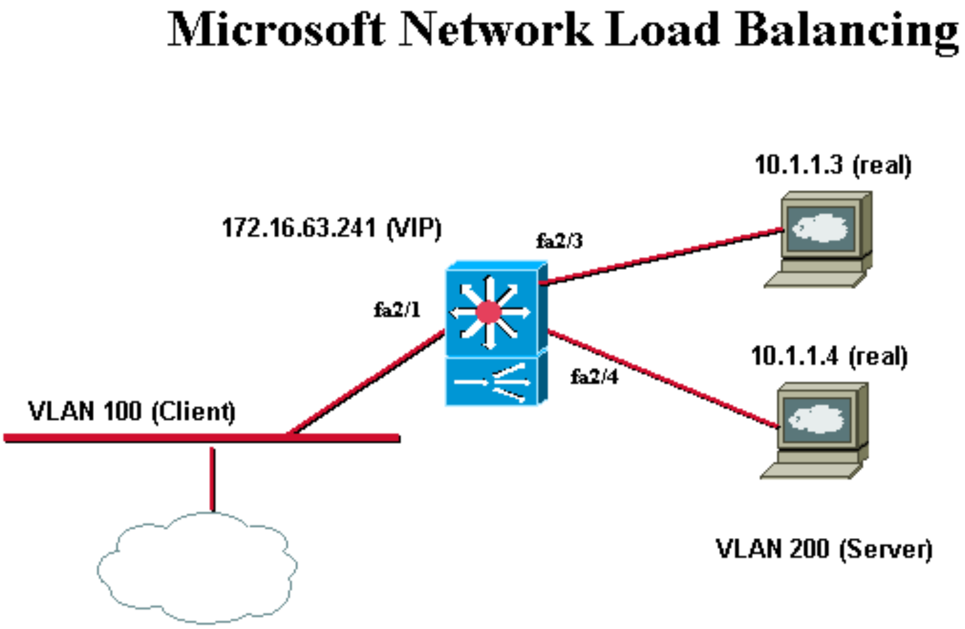

Configurations

This document uses the Catalyst 6509 configuration described in this section.

```
Cat6K#show running-config
 Building configuration...
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Cat6K
!
boot buffersize 126968
boot system flash slot0:c6sup11-jsv-mz.121-8a.E.bin
!
redundancy
 main-cpu
  auto-sync standard
ip subnet-zero
!
!
interface GigabitEthernet1/1
 no ip address
 shutdown
```

```
!
interface GigabitEthernet1/2
 no ip address
 shutdown
!
interface FastEthernet2/1
 description "Uplink to the Default Gateway"
 no ip address
 switchport
 switchport access vlan 100
!
interface FastEthernet2/2
 no ip address
 shutdown
!
interface FastEthernet2/3
 description "Connection to Microsoft server"
 no ip address
 switchport
 switchport access vlan 200
!
interface FastEthernet2/4
 description "Connection to Microsoft server"
 no ip address
 switchport
 switchport access vlan 200
!
interface FastEthernet2/5
 no ip address
 shutdown
!
interface FastEthernet2/48
 no ip address
 shutdown
!
interface Vlan1
 no ip address
 shutdown
!
mac-address-table static 0100.5e11.1111 vlan 200 interface fa2/3 fa2/4 disable-snooping

! --- Creating a static entry in the switch for the multicast virtual mac.

! --- fa2/3 & fa2/4 are the ports connected to server.

!--- The disable-snooping is applicable only for Cisco Catalyst 6500 series switches

arp 172.16.63.241 0100.5e11.1111

! --- 172.16.63.241 is the Virtual IP of 2 servers

interface Vlan100
 ip address 172.17.63.240 255.255.255.192

!--- Client Side Vlan

!
 interface Vlan200
 ip address 10.1.1.250 255.255.255.0

!--- Server Vlan

!--- Important: Configure the default gateway

!--- of the Microsoft Server to this address.
```

```
!
ip classless
ip route 0.0.0.0 0.0.0.0 172.16.63.193
no ip http server
!
line con 0
line vty 0 4
 login
!
end
```

Note: Ensure that you use the 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.

Verify

Use this section to confirm that your configuration works properly.

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

- ***show mac−address−table*** − Displays a specific MAC address table static and dynamic entry or the MAC address table static and dynamic entries on a specific interface or VLAN.

  ```
  Cat6K#show mac−address−table 0100.5e11.1111
            Mac Address Table
  -------------------------------------------
  Vlan    Mac Address       Type        Ports
  ----    -----------      --------     -----
  200    0100.5e11.1111     STATIC    Fa2/3 Fa2/4
  ```
- ***show ip arp*** − Displays the ARP cache.

  ```
  Cat6K#show ip arp
  Protocol  Address          Age (min)  Hardware Addr   Type   Interface
  Internet  172.16.1.1            −     0100.5e11.1111  ARPA   Vlan200
  ```

Troubleshoot

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

Related Information

- ***Technical Support & Documentation − Cisco Systems***

Updated: Dec 19, 2013 Document ID: 107995