Ethernet Routing Switch
Engineering

Technical Configuration Guide for Microsoft Network Load Balancing

Avaya Data Solutions
Document Date: June 2010
Document Number: NN48500-593
Document Version: 3.0

Abstract

The document provides an overview on how to configure Nortel Ethernet & Ethernet Routing Switches to support Microsoft's Network Load Balancing (NLB) server clustering technology.

Table of Contents

List of Figures

Document Updates

- January 22, 2010
  - Updated section 1.2.1 – corrected figure 1.2.1.3
  - Figure 2.2
  - Updated section 3.1
  - Removed reference to VLAN 1 in configuration examples
  - Added additional verification commands and results to each configuration example

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols:

Tip – Highlights a configuration or technical tip.

Note – Highlights important information to the reader.

Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text:

Bold text indicates emphasis.

Italic text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

ERS5520-48T# ***show running-config***

Output examples from Nortel devices are displayed in a Lucinda Console font:

ERS5520-48T# ***show running-config***

```
! Embedded ASCII Configuration Generator Script
! Model = Ethernet Routing Switch 5520-24T-PWR
! Software version = v5.0.0.011
enable
configure terminal
```

1. Overview: Network Load Balancing

Network Load Balancing is a clustering technology available with Microsoft Windows 2000 / Windows 2003 Server family of operating systems. Network Load Balancing uses a distributed algorithm to load balance TCP/IP network traffic across a number of hosts, enhancing the scalability and availability of mission critical, IP based services, such as Web, VPN, Streaming Media, Firewalls, etc. Network Load Balancing also provides high availability by detecting host failures and automatically redistributing traffic to remaining operational hosts.

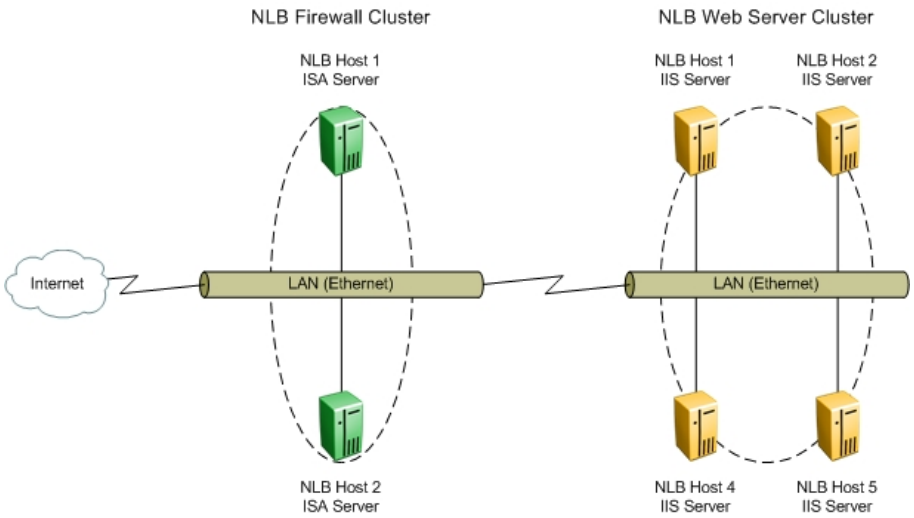

Figure 1.1 – Network Load Balancing Cluster

With Network Load Balancing, each host runs separate copies of the desired server applications, such as Web Server, FTP Server, or ISA Firewall. Network Load Balancing distributes incoming client requests to the hosts in the cluster group. The load weight to be handled by each host can be configured by the administrator and hosts can be dynamically added or removed from the cluster as necessary. In addition, Network Load Balancing can direct all traffic to a designated single host, called the default host.

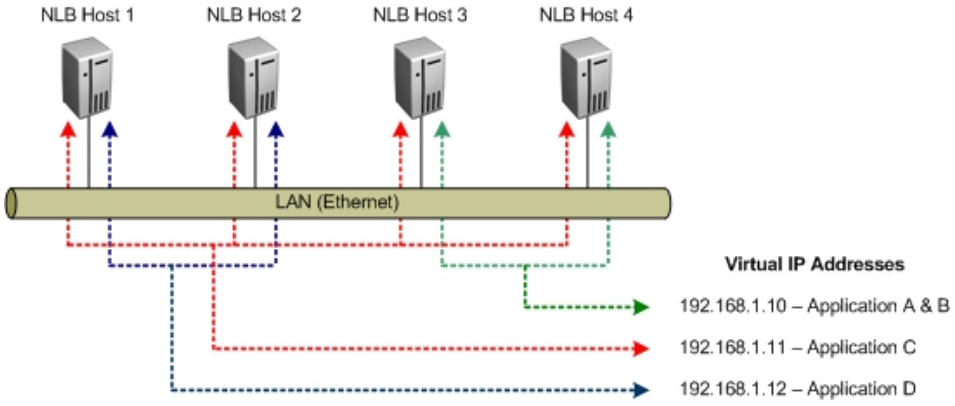

Figure 1.2 – Example Network Load Balancing Cluster

1.1 Architecture

Network Load Balancing uses fully distributed software architecture and an identical copy of the Network Load Balancing driver runs in parallel on each cluster host. The drivers arrange for all cluster hosts on a single subnet to concurrently detect incoming network traffic for the cluster's virtual IP address. On each cluster host, the driver acts as a filter between the network adapter's driver and the TCP/IP stack, allowing a portion of the incoming network traffic to be received by the host. By this means incoming client requests are partitioned and load-balanced among the Network Load Balancing cluster hosts.

Network Load Balancing runs as a network driver logically situated beneath higher-level application protocols, such as HTTP and FTP. Figure 1.1.1 shows the implementation of Network Load Balancing as an intermediate driver in the Windows 2000/2003 network stack.

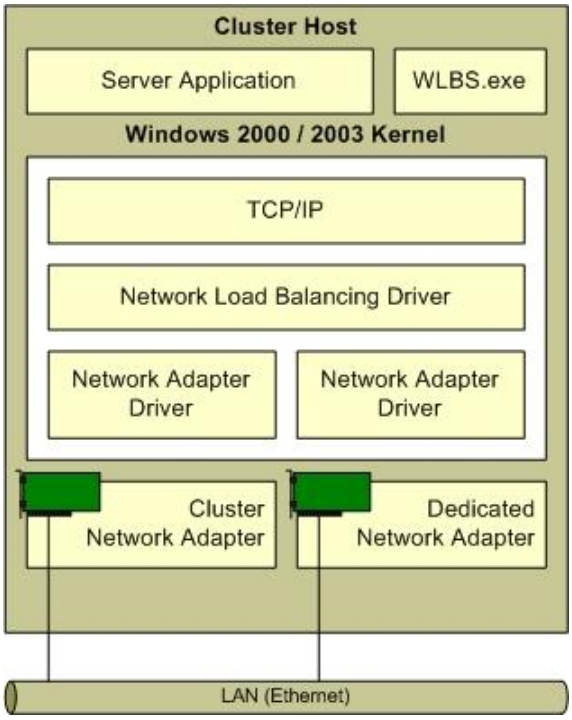

Figure 1.1.1 –Network Load Balancing Stack

The Network Load Balancing architecture maximizes throughput by using the broadcast domain to deliver incoming network traffic to all cluster hosts and by eliminating the need to route incoming packets to individual cluster hosts. Since filtering unwanted packets is faster than routing packets. As network and server speeds grow, its throughput also grows proportionally, thus eliminating any dependency on a particular hardware routing implementation.

Network Load Balancing architecture takes advantage of Ethernet switching architecture to simultaneously deliver incoming network traffic to all cluster hosts. However, this approach may increase the burden on switches by occupying additional port bandwidth. This is usually not a concern in most intended applications, such as Web services and streaming media, since the percentage of incoming traffic is a small fraction of total network traffic. However, if the client-side

network connections to the switch are significantly faster than the server-side connections, incoming traffic can occupy a prohibitively large portion of the server-side port bandwidth. The same problem arises if multiple clusters are hosted on the same switch and measures are not taken to setup virtual LANs for individual clusters.

1.2 Operation

Microsoft Network Load Balancing can be deployed in unicast (default), multicast and IGMP-multicast modes. These modes are configured on the MSNLB server cluster. The following sections highlight the three options for MSNLB configuration.

1.2.1 Unicast

Unicast mode is the default option for Network Load Balancing. With unicast mode, Network Load Balancing replaces the network adapter's real MAC address with a cluster virtual MAC address. All Network Load Balancing cluster host adapters share a common virtual MAC address and Virtual IP address and all frames forwarded to the cluster are received by all hosts in the cluster.

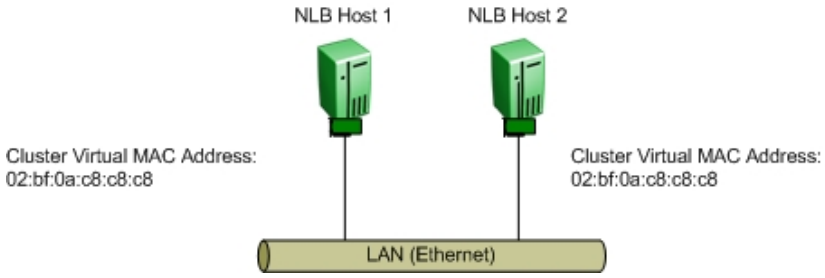

Figure 1.2.1.1 – Unicast Virtual MAC Assignment

Sharing a common MAC address amongst multiple hosts' works fine in shared media such as repeaters (hubs) but can cause issues in Ethernet switched environments.

An Ethernet switch forwards frames to hosts based on MAC addresses. An Ethernet switch does this by learning the MAC addresses of hosts connected to each of its ports. The Ethernet switch builds a forwarding database which provides a logical mapping of a MAC address to the port it was learned on. A switch expects that a MAC address is unique, only connected to one port, and therefore will not associate a MAC address with multiple ports of the switch.

As described above, unicast mode creates a cluster virtual MAC address that is common to all cluster hosts and an Ethernet switch would learn the clusters virtual MAC address on multiple ports. Since the switch only associates a MAC address to a single port and not many ports, Network Load Balancing will not function correctly.

Network Load Balancing solves this problem by masking the cluster virtual MAC address. When unicast mode is enabled, Network Load Balancing binds a bogus MAC address on each hosts adapter which starts with 02 and contains the host ID in the second octet. The bogus MAC address will appear in the Ethernet frame header and will be learned by the Ethernet switch rather than the clusters virtual MAC address. This ensures that the Ethernet switch will not learn the clusters virtual MAC addresses across multiple ports and will instead learn the unique MAC addresses for each cluster host.

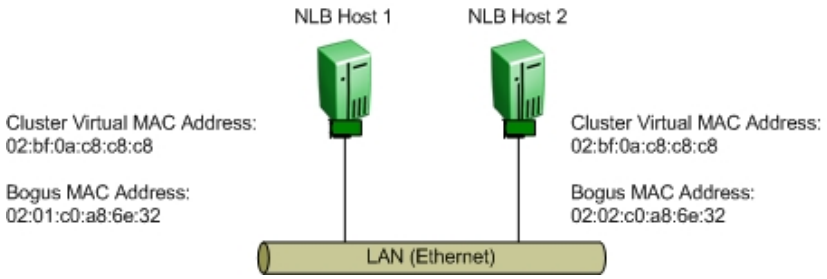

Figure 1.2.1.2 – Unicast Bogus MAC Assignment

If each network adapters MAC address is unique, how are frames delivered to all members of the cluster?

Microsoft Network Load Balancing solves this problem with IP. A client will learn the clusters MAC address using Address Resolution Protocol (ARP). When a client sends an ARP request for the MAC address of the clusters virtual IP address, the ARP response will contain cluster MAC virtual address and not the bogus MAC addresses.

Frames from the client will then be forwarded to the clusters virtual IP address with a destination MAC address set to the cluster MAC address. On receipt of the frames, the Ethernet switch will perform a lookup and will not have a forwarding entry for the clusters virtual MAC address. The switch will then flood the frames to all active ports in the broadcast domain so that all hosts in the cluster will receive the frames.

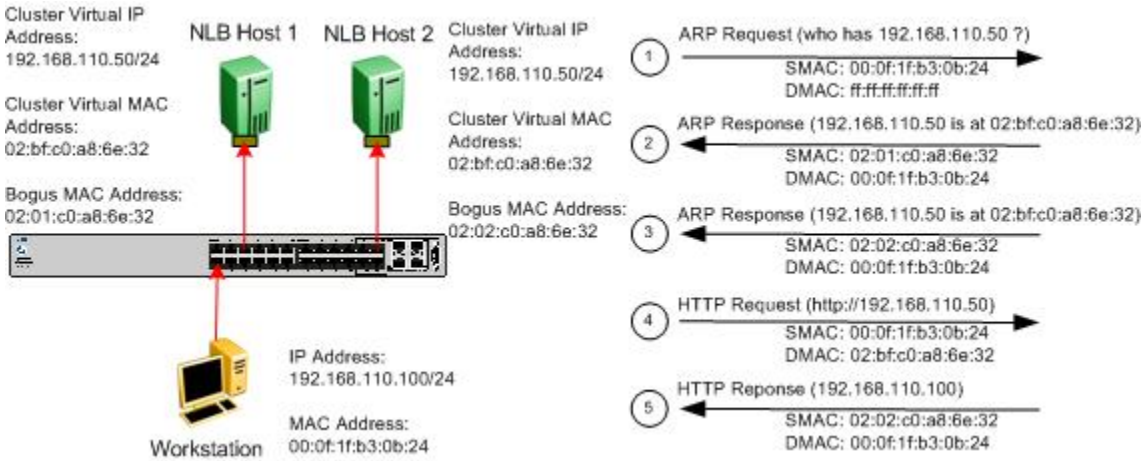

Figure 1.2.1.3 – Unicast Traffic Flow

Please refer to the Microsoft support bulletins 898867 and 193602 in reference to NLB Unicast operation.

http://support.microsoft.com/kb/898867/en-us

http://support.microsoft.com/kb/193602.

Assuming the above switch is an ERS 5520 with NLB VLAN 1300, we can view the MAC address table by using the command shown below. Notice the NLB cluster virtual MAC address is never learned by the layer 2 switch. Hence, when the client forwards traffic to NLB cluster, the packet will be flooded as the NLB cluster virtual MAC is unknown to the switch.

```
5520T-PWR#show mac-address-table vid 1300
Mac Address Table Aging Time: 300
Number of addresses: 4

   MAC Address    Vid  Source         MAC Address    Vid  Source
----------------- ---- -------     ----------------- ---- -------
00-0f-1f-b3-0b-24 1300 Port:2      02-01-C0-A8-6E-32 1300 Port:5
02-02-C0-A8-6E-32 1300 Port:23
```

1.2.2 Multicast / IGMP-Multicast

Multicast and IGMP-multicast modes are optional modes for Network Load Balancing. With multicast mode, a multicast virtual MAC address with the prefix 03-bf is bound to all cluster hosts but the network adapter's real MAC address is retained. The multicast MAC address is used for client-to-cluster traffic and the adapter's real MAC address is used for network traffic specific to the host computer.

MACs starting with odd numbers are multicast.

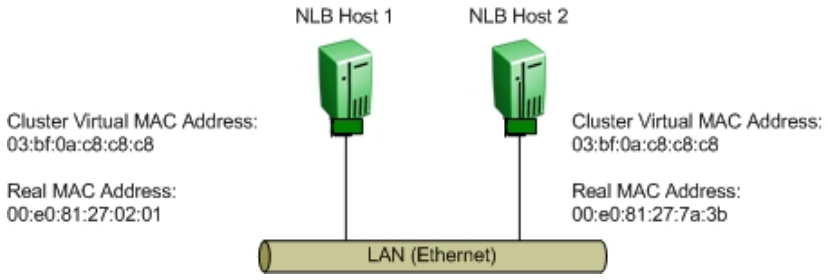

Figure 1.2.2.1 – Multicast MAC Assignment

With IGMP-multicast mode, a multicast virtual MAC address with the prefix 01-00 is bound to all cluster hosts and the network adapter's real MAC address is retained. The multicast MAC address is used for client-to-cluster traffic and the adapter's real MAC address is used for network traffic specific to the host computer.

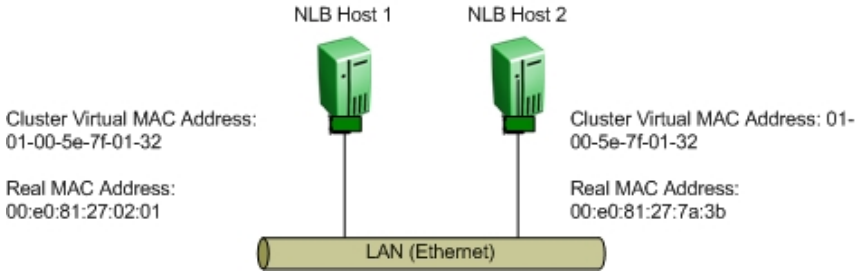

Figure 1.2.2.2 – IGMP-Multicast MAC Assignment

Both multicast and IGMP-multicast modes operate by all cluster hosts receiving the frames from the clients. With multicast mode all traffic forwarded to the clusters virtual IP address is flooded to all ports in the broadcast domain which ensures that all hosts in the cluster will receive the frames.

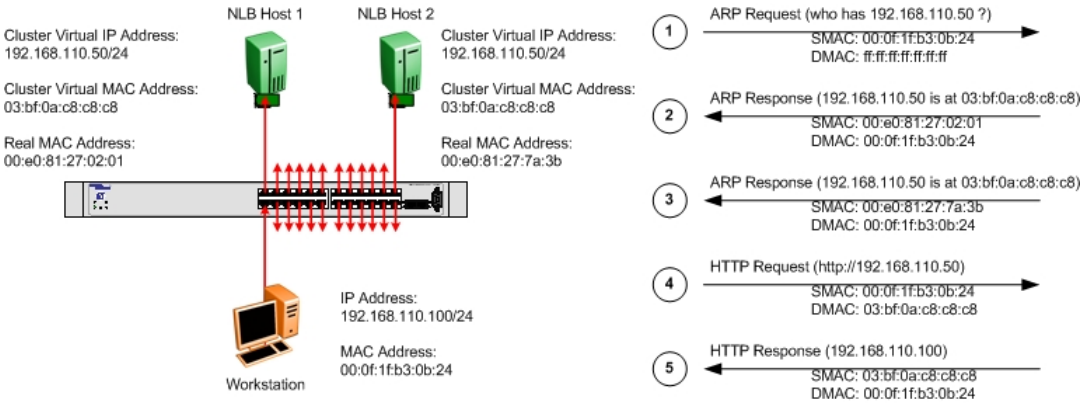

Figure 1.2.2.3 – Multicast Traffic Flow

IGMP-multicast mode implements IGMP and all hosts in the cluster forward IGMPv1 group membership reports. IGMP allows the Ethernet switches to prune the multicast traffic and limit the flooding to only the ports that connect to the cluster hosts. When IGMP-multicast mode is enabled, traffic is pruned.

Frames from clients are forwarded to the clusters virtual IP address with a destination MAC address set to the clusters virtual multicast MAC address. Depending on the multicast mode, the frames are ether flooded to all ports in the broadcast domain or forwarded to only the ports that the cluster hosts are connected to.

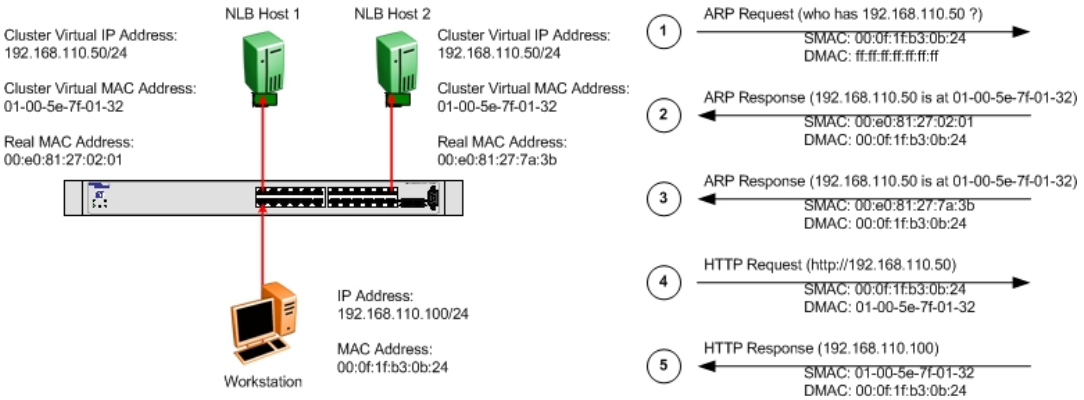

Figure 1.2.2.4 – IGMP-Multicast Traffic Flow

1.3 Load Balancing Algorithm

Network Load Balancing employs a fully distributed filtering algorithm to map incoming clients to the cluster hosts. The distributed algorithm enables cluster hosts to make load-balancing decisions independently and quickly for each incoming packet. The distributed algorithm is optimized to statistically load balance traffic for large client populations and is less effective when the client population is small or the client connections produce widely varying loads on the server.

Network Load Balancing balances incoming client requests by directing a selected percentage of new requests to each cluster host. The algorithm does not respond to changes in the load on each cluster host (such as the CPU load or memory usage). However, the mapping is modified when the cluster membership changes, and load percentages are renormalized accordingly.

When inspecting an arriving packet, all cluster hosts simultaneously perform a statistical mapping to quickly determine which host should handle the packet. The mapping uses a randomization function that calculates a host priority based on the client's IP address, port, and other state information. The corresponding host forwards the packet up the network stack to TCP/IP, and the other cluster hosts discard it. The mapping does not vary unless the membership of cluster hosts changes, ensuring that a given client's IP address and port will always map to the same cluster host. The particular cluster host to which the client's IP address and port map cannot be predetermined since the randomization function takes into account the current and past cluster's membership to minimize re-mappings.

1.4 Convergence

Network Load Balancing hosts periodically exchange multicast or broadcast heartbeat messages within the cluster. This allows the hosts to monitor the status of the cluster. When the state of the cluster changes (such as when hosts fail, leave, or join the cluster), Network Load Balancing invokes a process known as convergence, in which the hosts exchange heartbeat messages to determine a new, consistent state of the cluster and to elect the host with the highest host priority as the new default host.

During convergence, the hosts continue to handle incoming network traffic as usual, except that traffic for a failed host does not receive service. Client requests to surviving hosts are unaffected. Convergence terminates when all cluster hosts report a consistent view of the cluster membership for several heartbeat periods. If a host attempts to join the cluster with inconsistent port rules or an overlapping host priority, completion of convergence is inhibited. This prevents an improperly configured host from handling cluster traffic.

At the completion of convergence, client traffic for a failed host is redistributed to the remaining hosts. If a host is added to the cluster, convergence allows this host to receive its share of load-balanced traffic. Expansion of the cluster does not affect ongoing cluster operations and is achieved in a manner transparent to both Internet clients and to server programs. However, it may affect client existing sessions because clients may be remapped to different cluster hosts between connections.

In unicast, multicast and IGMP-multicast modes, each cluster host generates heartbeat messages. Each heartbeat message occupies one Ethernet frame and is tagged with the cluster's primary IP address so that multiple clusters can reside on the same subnet. Network Load Balancing's heartbeat messages are assigned an ether type-value of hexadecimal 886F

and by default are forwarded every second. During convergence, the exchange period is reduced by half in order to expedite the convergence process.

Network Load Balancing assumes that a host is functioning properly within the cluster as long as it participates in the normal heartbeat exchange among the cluster hosts. If other hosts do not receive a heartbeat message from any member for several periods of message exchange, they initiate convergence. The number of missed heartbeat messages is set to five by default.

1.5 MAC Address Formats

Microsoft Network Load Balancing can be implemented in unicast, multicast or IGMP-multicast modes and the MAC address formats used by the cluster hosts will depend on the cluster mode. The following section describes the IEEE formatting of Ethernet MAC addresses as well as the MAC address formats for each Network Load Balancing mode.

In Ethernet there are four types of MAC addresses defined by IEEE:

MAC Address Type	MAC Address Range
Globally Unique	x0-xx-xx-xx-xx-xx
	x4-xx-xx-xx-xx-xx
	x8-xx-xx-xx-xx-xx
	xC-xx-xx-xx-xx-xx
Locally Administered	x2-xx-xx-xx-xx-xx
	x6-xx-xx-xx-xx-xx
	xA-xx-xx-xx-xx-xx
	xE-xx-xx-xx-xx-xx
Multicast	x1-xx-xx-xx-xx-xx
	x3-xx-xx-xx-xx-xx
	x5-xx-xx-xx-xx-xx
	x7-xx-xx-xx-xx-xx
	x9-xx-xx-xx-xx-xx
	xB-xx-xx-xx-xx-xx
	xD-xx-xx-xx-xx-xx
	xF-xx-xx-xx-xx-xx (exception broadcast address)
Broadcast	FF-FF-FF-FF-FF-FF

1.5.1 Globally Unique

Globally unique addresses are allocated by the IEEE in blocks containing 2^24 (16,777,216) addresses and start with even numbers. In each allocation, the first 3 octets are fixed (e.g. 00-12-83 is Nortel) and the last three octets are variable (e.g. 00-00-00 through FF-FF-FF). The fixed portion of the allocation is known formally as the Organizationally Unique Identifier (OUI) and is used informally as the Vendor ID.

1.5.2 Locally Administered

Locally administered addresses are MAC addresses which have the second least significant bit of the first octet is set to '1' (for example, 'xxxxxx1x'. Locally administered addresses enable administrators to assign MAC addresses using their own scheme.

1.5.3 Multicast

Multicast addresses have the least significant bit of the first octet set to '1' and start with an odd number. Ethernet multicast addressing is used by protocols which require efficient communication among groups of hosts.

1.5.4 Broadcast

Broadcast address is a special case where all bits of the MAC address are set to '1' (e.g. FF-FF-FF-FF-FF-FF).

When an adapter receives a packet with a destination broadcast address, it always passes it to the operating system for further processing.

1.5.5 Network Load Balancing Unicast

When NLB is deployed in unicast mode, the globally unique MAC address on each cluster hosts network adaptor is replaced with a locally administered MAC address assigned by Microsoft. The locally administered MAC address starts with a 02:xx prefix and the second octet will contain the host-id of the host in the cluster.

The clusters virtual MAC address is also a locally administered MAC address and starts with a 02:bf prefix.

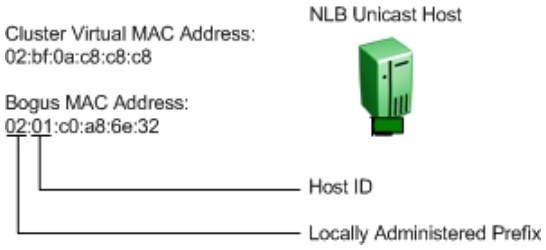

Figure 1.5.1 – Unicast MAC Format

1.5.6 Network Load Balancing Multicast / IGMP-Multicast

When Microsoft Network Load Balancing is deployed in multicast or IGMP-multicast modes, the globally unique MAC address on the hosts network adaptor is retained.

The clusters virtual MAC address is multicast MAC address assigned by Microsoft and will start with a 03:bf prefix for multicast mode or 01:00 prefix for IGMP-multicast mode. All the hosts in cluster will be configured with the same multicast virtual MAC address.

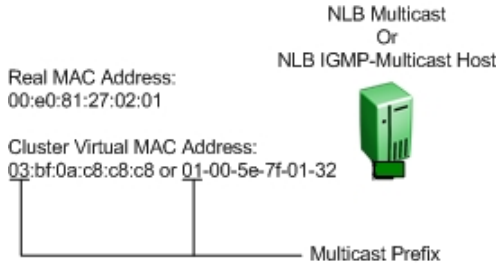

Figure 1.5.2 – Multicast / IGMP-Multicast MAC Format

1.6 Implementation Models

Microsoft's Network Load Balancing can be deployed using one of four models. This section provides a brief overview of the supported models and provides advantages and disadvantages of each.

1.6.1 Single Network Adapter in Unicast Mode

The single network adapter unicast model is suitable for a cluster in which ordinary network communication among cluster hosts is not required and there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts.

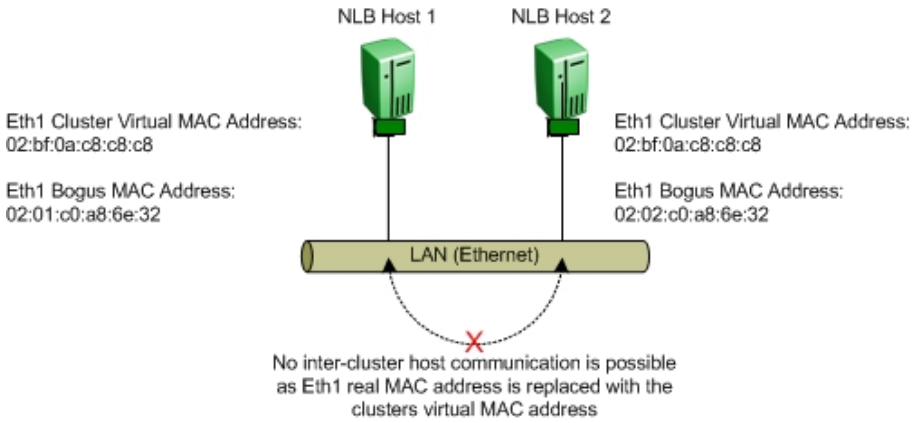

Figure 1.6.1 – Single Adapter Unicast Mode

Advantages	Disadvantages
One network adapter per cluster host is required.	Network communication between cluster hosts is not possible.
Minimum configuration is required.	All traffic from clients to cluster hosts will be flooded throughout the broadcast domain.
Works with all routers and L2 switches.	Not supported by all L3 switches.

1.6.2 Single Network Adapter in Multicast / IGMP-Multicast Mode

The single network adapter multicast / IGMP-multicast model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable, but in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts.

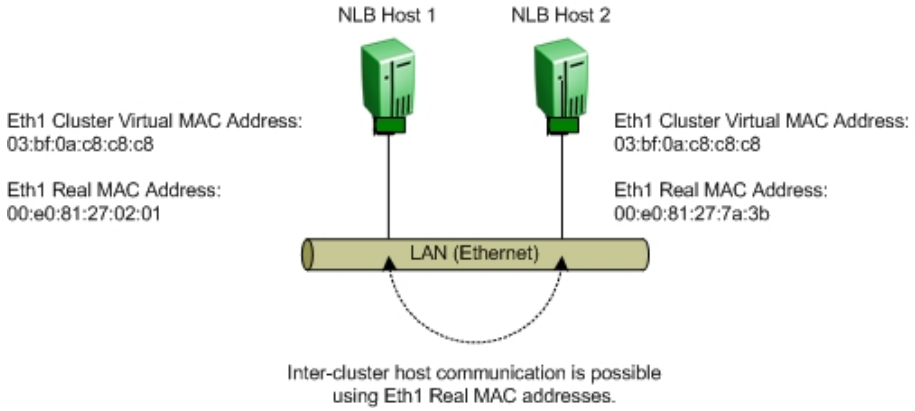

Figure 1.6.2 – Single Adapter Multicast / IGMP-Multicast Mode

Advantages	Disadvantages
One network adapter per cluster host is required.	Some Routers or Routing Switches may not support the ability to map a unicast IP address with a multicast MAC address.
Network communication between cluster hosts is permitted.	Some Routers or Routing Switches may not be able to dynamically learn the clusters virtual MAC address.
Flood suppression is available with IGMP-multicast mode.	

1.6.3 Multiple Network Adapters in Unicast Mode

The multiple network adapter unicast model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable. It is also appropriate when you want to separate the traffic used to manage the cluster from the traffic occurring between the cluster and client computers.

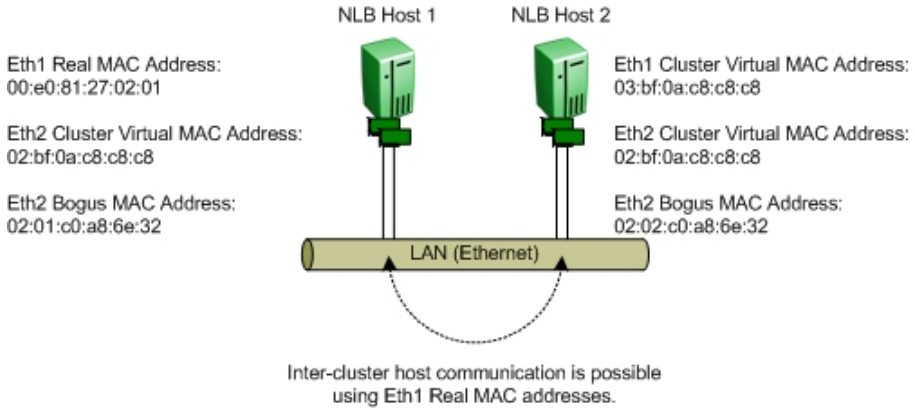

Figure 1.6.3 – Multiple Adapters Unicast Mode

Advantages	Disadvantages
Network communication between cluster hosts is permitted.	This model requires a second network adapter.
This model works with all routers and L2 switches.	All traffic from clients to cluster hosts will be flooded to the broadcast domain. Not supported by all L3 switches.

1.6.4 Multiple Network Adapters in Multicast / IGMP-Multicast Mode

The multiple network adapter multicast model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary and in which there is heavy dedicated traffic from outside the cluster subnet to specific cluster hosts.

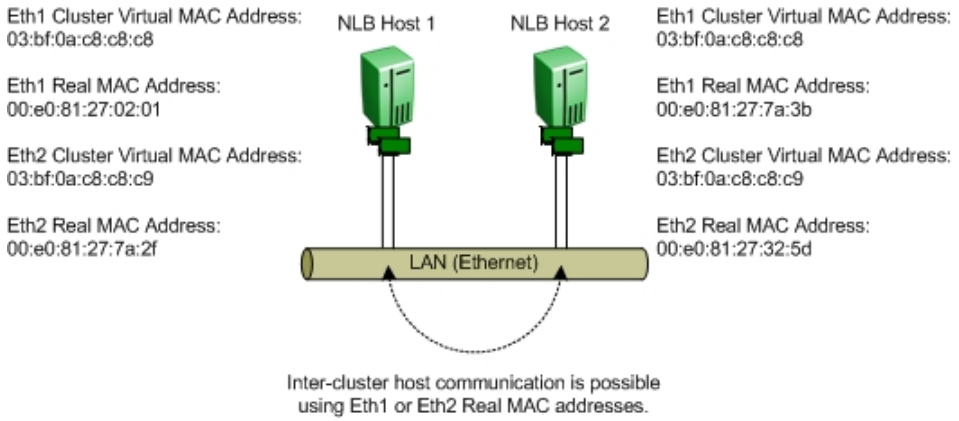

Figure 1.6.4 – Multiple Adapters Multicast / IGMP-Multicas Mode

Advantages	Disadvantages
Network communication between cluster hosts is permitted.	This model requires a second network adapter.
Cluster performance may be enhanced.	Some Routers or Routing Switches may not support the ability to map a unicast IP address with a multicast MAC address. Some Routers or Routing Switches may not be able to dynamically learn the clusters virtual MAC address.

There are no restrictions on the number of network adapters that can be bound to network load balancing on a host computer. Each host may have a different number of adapters, but you can never have more than one adapter on a host be part of the same cluster.

Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

2. Supported Topologies & Releases

The following section outlines the tested and supported Network Load Balancing topologies on Avaya Ethernet switching platforms. This section provides information on specific releases of software that may be required as well as any features that may need to be enabled on Avaya Ethernet switching platforms to support Microsoft Network Load Balancing clusters.

This section assumes that the reader has configuring experience with parameters such as VLANs, IP interfaces, MLT, SMLT and RSMLT. Step-by-step configuration examples on how to configure these parameters is out of the scope of this document. For assistance with configuring these parameters please refer to the product documentation, technical configuration guides and technical solution guides available on Avaya's technical support Web site.

An example of how to create a Microsoft Network Load Balancing cluster for unicast, multicast or IGMP-multicast mode is provided for convenience to the reader in Section 3.

2.1 Single Layer 2 Switch

The following topology is supported on all Avaya Ethernet switching platforms. Using this topology, a customer can deploy Network Load Balancing clusters in unicast, multicast and IGMP-multicast modes. It's important to note that with this topology no IP routing is enabled on the Ethernet switching platform.

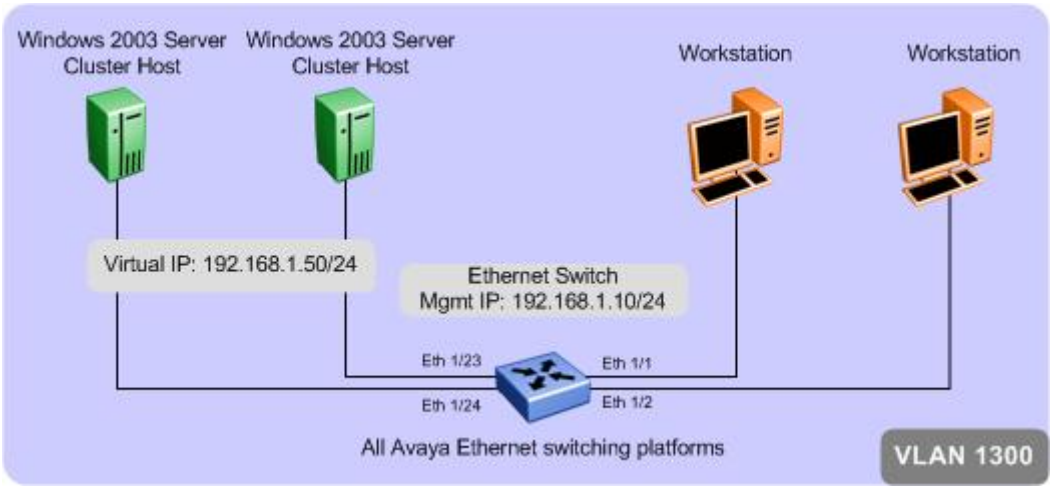

Figure 2.1 – Single Layer 2 Switch

2.1.1 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet switching platforms that can be deployed to support this topology:

Avaya Switch Model	Unicast Mode	Multicast Mode	IGMP-Multicast Mode
ERS 8600	Yes	Yes	Yes
ERS 8300	Yes	Yes	Yes
ERS 5000	Yes	Yes	Yes
ERS 1600	Yes	Yes	Yes
ERS 4500	Yes	Yes	Yes
ERS 2500	Yes	Yes	Yes

Table 2.1.1 – Supported Avaya Ethernet switch platforms

2.1.2 Configuration

To support this topology the following configuration steps need to be performed on the Ethernet switching platform:

Mandatory Configuration Steps

No mandatory configuration steps need to be performed. By default Avaya Ethernet switching platforms will flood Network Load Balancing cluster traffic with no additional configuration being required.

Optional Configuration Steps

If Network Load Balancing clusters are deployed using IGMP-multicast mode, administrators may optionally enable IGMP snooping and proxy to eliminate the flooding of cluster traffic to non cluster hosts.

2.1.2.1 Configuration Steps

The following CLI commands create VLAN 1300 with option to enable / disable IGMP snooping and IGMP proxy for NLB IGMP-Multicast:

Step 1 – Create VLAN 1300
4550T-PWR(config)#***vlan configcontrol automatic***
4550T-PWR(config)#***vlan create 1300 name NLB type port 1***
4550T-PWR(config)#***vlan members add 1300 1-24***
Step 2 – Enable IGMP Snoop and Proxy if using NLB Multicast
4550T-PWR(config)#***vlan igmp 1300 snooping enable***
4550T-PWR(config)#***vlan igmp 1300 proxy enable***

2.1.2.1.1 Verify Operations

Step 1 – The following CLI command displays the IGMP configuration for VLAN 1300:
4550T-PWR#***show vlan igmp 1300***
Result:

```
        Snooping:  Enabled
        Proxy:  Enabled
        Robust Value:  2
        Query Time:  125 seconds
        IGMPv1 Static Router Ports: NONE
        IGMPv2 Static Router Ports: NONE
        Querier Port: NONE
        Multicast Router Expiration: 0 seconds
```

Step 2 – The following CLI command displays the multicast group membership for VLAN 1300. In this example the cluster hosts are connected to ports 1/23 & 1/24:
4550T-PWR# ***show vlan multicast membership 1300***
Result:

```
        Number of groups: 1
        Multicast Group Address Port
        ----------------------- ----
        239.255.1.50                            23
        239.255.1.50                            24
```

2.2 Centralized Routing Switch – ERS 5000 or ERS 4500

The following topology is supported when an Ethernet Routing Switch 5000, 4500 or 1600 is used to route between server and client VLANs. The Network Load Balancing cluster hosts must be connected to a Layer 2 subtended Ethernet Switch. The clients may be connected directly to the Core switch or to a Layer 2 subtended Ethernet Swtich. The subtended Ethernet switch can use a single uplink port or a multi-port MLT/DMLT trunk. This topology supports Network Load Balancing clusters in unicast, multicast and IGMP-multicast modes.

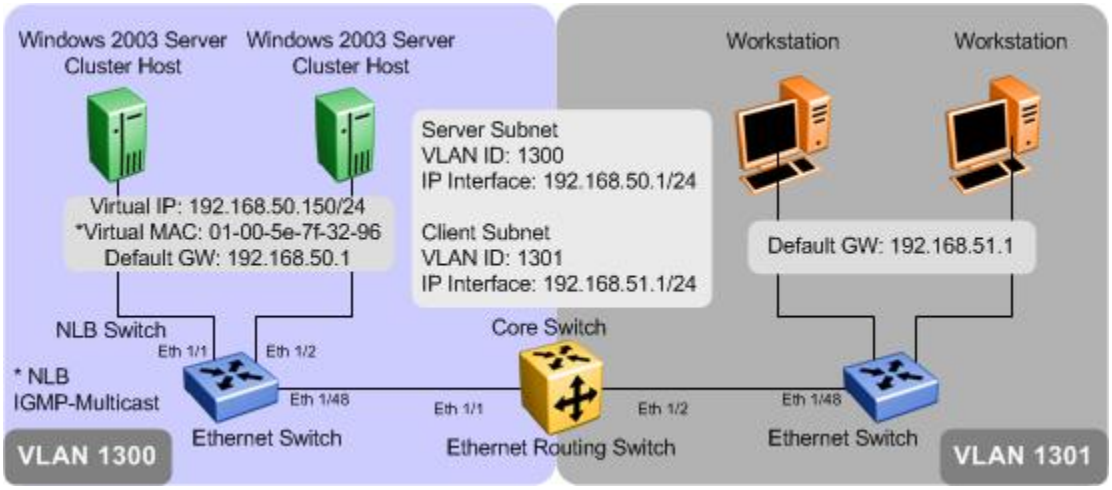

Figure 2.2 – Centralized Routing Switch

2.2.1 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a centralized Routing switches to support this topology:

Avaya Switch Model	Microsoft Server NLB Mode		
	Unicast Mode	Multicast Mode	IGMP-Multicast Mode
ERS 5000	Yes	Yes	Yes
ERS 4500	Yes	Yes	Yes

Table 2.2.1 – Supported Avaya Ethernet Routing switch platforms

2.2.2 Configuration

To support this topology the following configuration steps need to be performed on the Ethernet Routing switching platform:

2.2.2.1 Central Switch

The following show the configuration steps for an ERS 5530 switch. Since the workstation and the NLB switches are directly connected to the core switch, there is no need to enable either a dynamic routing protocol or static routes on the core switch resulting in a very simple configuration.

Step 1 – Create VLAN 1300 and 1301 and add port members

```
5530-24TFD(config)#vlan configcontrol automatic
5530-24TFD(config)#vlan create 1300 name NLB type port 1
5530-24TFD(config)#vlan create 1301 name Client type port 1
5530-24TFD(config)#vlan members add 1300 1
5530-24TFD(config)#vlan members add 1301 2
```

Step 2 – Add IP address to each VLAN and enable IP routing

```
5530-24TFD(config)#ip routing
5530-24TFD(config)#interface vlan 1300
5530-24TFD(config-if)#ip address 192.168.50.1 255.255.255.0
5530-24TFD(config-if)#exit
5530-24TFD(config)#interface vlan 1301
5530-24TFD(config-if)#ip address 192.168.51.1 255.255.255.0
5530-24TFD(config-if)#exit
```

Step 3 – Create a Static ARP Entry on the ERS 5530 if the NLB is running in multicast or IGMP-multicast mode

```
5530-24TFD(config)#ip arp 192.168.50.150 01:00:5e:7f:32:96 1/1 vid 1300
```

2.2.2.1.1 Verify Operations

Step 1 – The following CLI command displays the arp entry for 192.168.1.50:

```
ERS5530-24TFD(config)#show ip arp 192.168.50.150
```

Result:

```
================================================================================
                                    IP ARP
================================================================================
IP Address      Age (min) MAC Address       VLAN-Unit/Port/Trunk Flags

--------------------------------------------------------------------------------
192.168.50.150  0         01:00:5e:7f:32:96 VLAN#1300-1          S
Total ARP entries : 1
--------------------------------------------------------------------------------
Flags Legend:
S=Static, D=Dynamic, L=Local, B=Broadcast
```

2.2.2.2 NLB Edge Switch

Assuming the NLB edge switch is an ERS 4548GT-PWR, enter the configuration shown below. The following CLI commands create VLAN 1300 with option to enable / disable IGMP snooping and IGMP proxy for NLB Multicast:

Step 1 – Create VLAN 1300

```
4548GT-PWR(config)#vlan configcontrol automatic
4548GT-PWR(config)#vlan create 1300 name NLB type port 1
4548GT-PWR(config)#vlan members add 1300 1/1-24
```

Step 2 – Enable IGMP Snoop and Proxy if using NLB Multicast

```
4548GT-PWR(config)#vlan igmp 1300 snooping enable
4548GT-PWR(config)#vlan igmp 1300 proxy enable
```

2.2.2.2.1 Verify Operation

Step 1 – The following CLI command displays the IGMP configuration for VLAN 1300:

```
4548GT-PWR#show vlan igmp 1300
```

Result:

```
        Snooping:  Enabled
        Proxy:  Enabled
        Robust Value:  2
        Query Time:  125 seconds
        IGMPv1 Static Router Ports: NONE
        IGMPv2 Static Router Ports: NONE
        Querier Port: NONE
        Multicast Router Expiration: 0 seconds
```

Step 2 – The following CLI command displays the multicast group membership for VLAN 1300:

```
4548GT-PWR# show vlan multicast membership 1300
```

Result:

```
        Number of groups: 1
        Multicast Group Address Unit Port
        ----------------------- ---- ----
        239.255.50.150          1    1
        239.255.50.150          1    2
```

2.3 Single L3 Switch - Ethernet Routing Switch 8300/8600

The following topology is supported when an Ethernet Routing Switch 8600 is used to route between server and client VLANs when both Network Load Balancing cluster hosts and clients are directly connected to the Avaya Ethernet Routing Switch 8600 and IP routing is enabled. This topology supports Network Load Balancing clusters in unicast, multicast and IGMP-multicast modes.

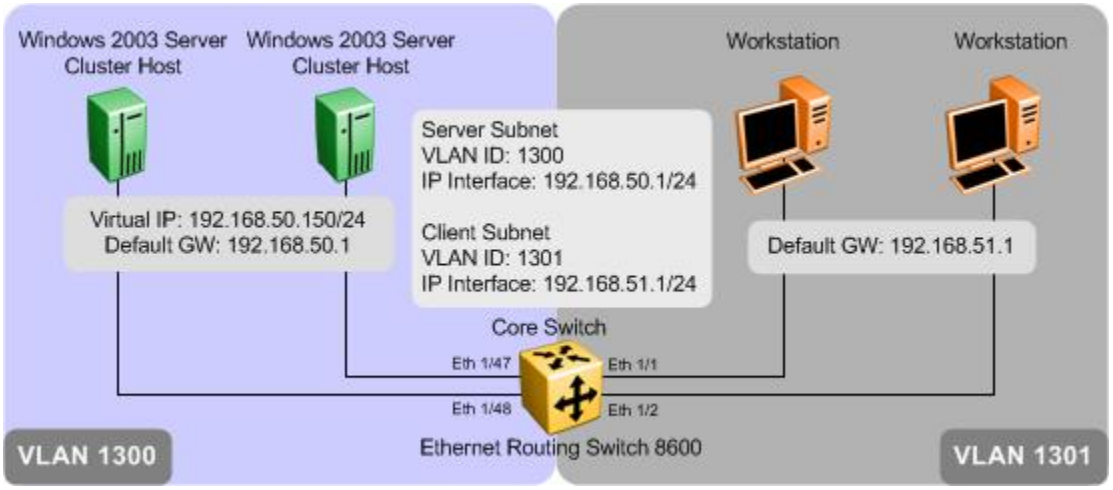

Figure 2.3 – Single Ethernet Routing Switch 8600

2.3.1 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed to support this topology:

Avaya Switch Model With NLB Support	Microsoft Server NLB		
	Unicast Mode	Multicast Mode	IGMP-Multicast Mode
ERS 8600	Yes	Yes	Yes
ERS 8300	Yes	Yes	Yes

Table 2.3.1 – Supported Avaya Ethernet Routing switch platforms

2.3.2 Configuration

To support this topology the following configuration steps need to be performed on the Ethernet Routing Switch 8600:

Mandatory Configuration Steps
1. The per Avaya Ethernet Routing Switch VLAN NLB mode must match the Microsoft Server NLB mode. The ERS 8600 must have 4.1.1 or higher and the ERS 8300 must have 4.0 or higher.
Optional Configuration Steps
None

Table 2.3.2.1 – Configuration Steps

2.3.2.1 Configuration Steps – Per VLAN NLB

Step 1 – The following PPCLI command enables / disables per VLAN NLB unicast, multicast or IGMP-multicast support for VLAN 1300
ERS-8600# ***config vlan 1300 nlb-mode <disable/igmp-mcast/multicast/unicast>***
Step 2 – If IGMP-multicast NLB is enabled, also enable IGMP snoop and proxy using the following command for VLAN 1300
ERS-8600# ***config vlan 1300 ip igmp snoop enable*** ERS-8600# ***config vlan 1300 ip igmp proxy-snoop enable***

2.3.2.1.1 Verify Operations

Step 1 – The following PPCLI displays the status of the per VLAN NLB support showing the status when NLB unicast support is enabled for VLAN 1300 and cluster hosts are connected to port 1/47 & 1/48. For this example, NLB IGMP-multicast is enabled.
ERS8600# ***show vlan info nlb-mode***
Result:

```
================================================================================
                                  Vlan Nlb
================================================================================
VLAN_ID  NLB_ADMIN_MODE NLB_OPER_MODE  PORT_LIST                    MLT_GROUPS
--------------------------------------------------------------------------------
1300     igmp-mcast     igmp-mcast    1/47-1/48

Total Entries: 1
```

Step 2 – If Network Load Balance IGMP-multicast support is enabled for VLAN 1300 and the cluster hosts are connected to port 1/47 & 1/48, you can view the member and group address by issuing the following command:

ERS-8600# ***show ip igmp group***

Result:

```
================================================================================
                            IGMP Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
239.255.50.150  V1300-1/47      192.168.50.150  222        Dynamic
239.255.50.150  V1300-1/48      192.168.50.150  253        Dynamic

2 out of 2 group Receivers displayed

Total number of unique groups 1
```

2.3.2.2 Configuration Step – Global ARP Multicast MAC Flooding

In older software releases prior to 3.7.15, Network Load Balancing can be deployed in multicast mode by enabling the global ARP multicast MAC flooding feature.

Step 1 – The following PPCLI command enables / disables the global ARP multicast MAC flooding feature.

ERS-8600# ***config ip arp multicast-mac-flooding <enable|disable>***

2.3.2.2.1 Verify Operations

Step 1 – Verify that the NLB mode configured is set to multicast. The NLB operational state should also display multicast if configured correctly with uplink port to the NLB server.

ERS8600# ***config ip arp info***

Result:

```
multicast-mac-flooding :  enable
aging : 360 (min)
arpreqthreshold : 500
delete : N/A
add :
```

2.4 Switch Clustering - Topologies

Section 2.4 will cover various SMLT cluster topologies and the type of configuration required on the SMLT cluster switches. Section 2.5 will cover the configurations details.

2.4.1 Switch Clustering – Topology 1

The following topology is supported when Ethernet Routing Switch 8600s are deployed as a SMLT core where Network Load Balancing cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster. This topology supports Network Load Balancing clusters in unicast and multicast modes.

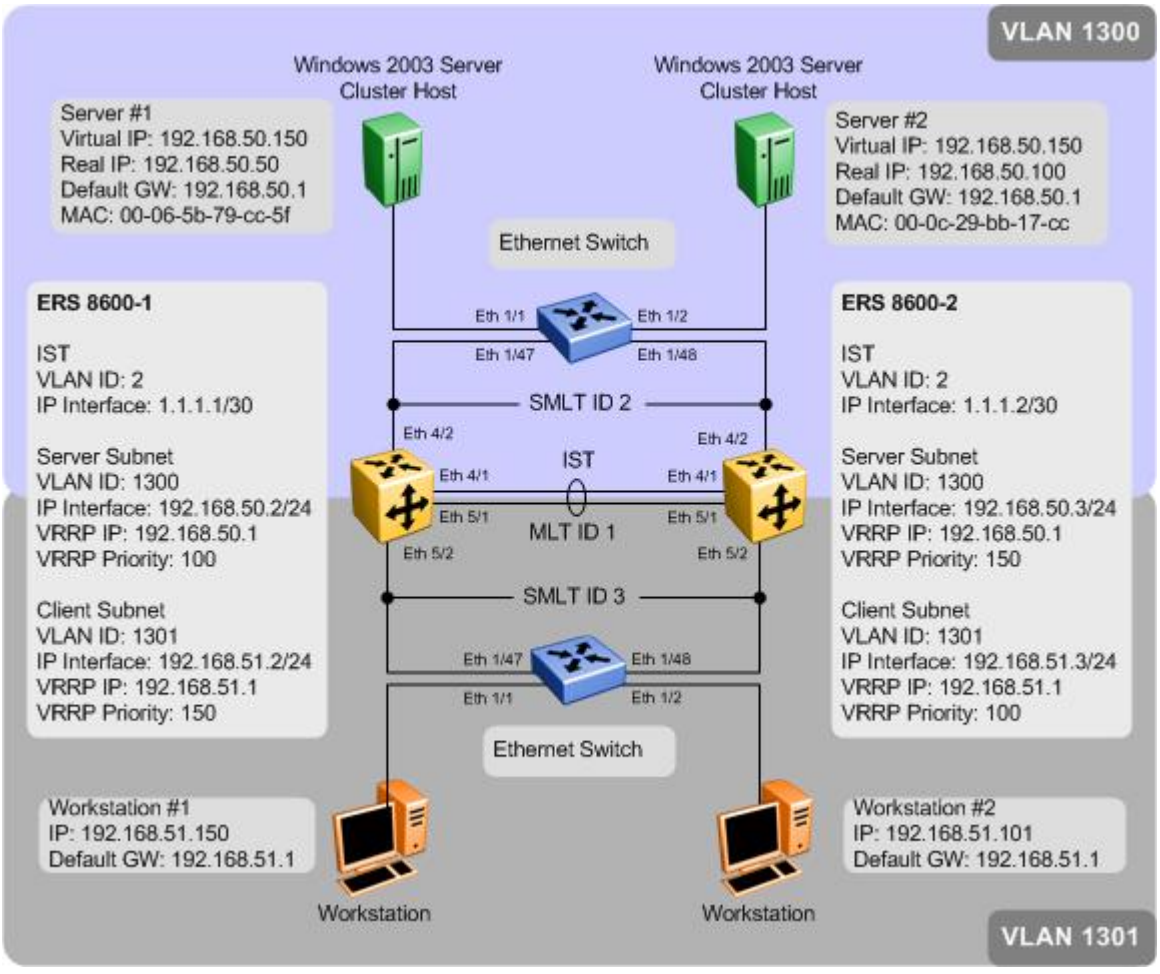

Figure 2.4 – Switch Clustering – Topology 1

2.4.2 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a SMLT core to support this topology:

Avaya Switch		Microsoft NLB Server Configuration		
Switch	**Configuration**	**Unicast Mode**	**Multicast Mode**	**IGMP-Multicast Mode**
ERS 8600 ERS 8300	NLB Unicast	Yes	No	No
ERS 8600 ERS 8300	Static Multicast entry	No	Yes[1]	No
ERS 8600 ERS 8300	ARP Multicast Flooding	No	Yes	No
ERS 8300	NLB Multicast	No	Yes[2]	No

Table 2.4.1 – Supported Avaya Ethernet Routing switch platforms

Note 1 – Required the system flag enhanced-operational-mode to be enabled and is not supported on legacy I/O modules

Note 2 – Normally, only one of the cluster switches will register the NLB Server ARP and forwarding port entries. Traffic will be forwarded on this cluster switch to the active NLB server either via the local SMLT or via the IST connection. If this node should fail, the peer SMLT cluster switch will forward traffic to the Microsoft NLB server via the SMLT connection.

2.4.3 Switch Clustering – Topology 2

The following topology is supported when Ethernet Routing Switch 8600s are deployed as a SMLT core and Network Load Balancing cluster hosts are directly connected and distributed between the ERS 8600s and the clients are connected to Layer 2 Ethernet switch that is SMLT connected to the SMLT cluster. This topology supports Network Load Balancing clusters in unicast and multicast modes.

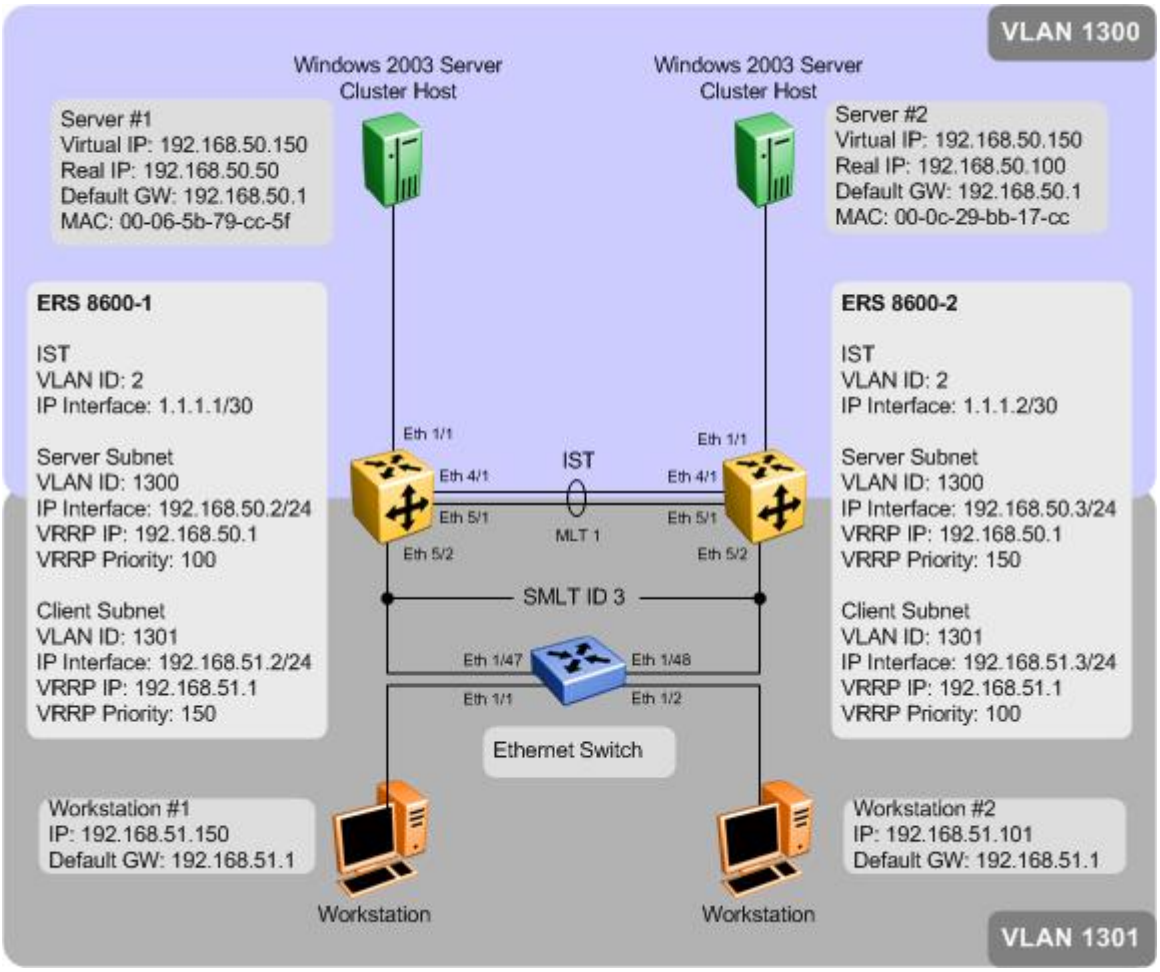

Figure 2.5 – Switch Clustering – Topology 2

2.4.4 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a SMLT core to support this topology:

Avaya Switch		Microsoft NLB Server Configuration		
Switch	**Configuration**	**Unicast Mode**	**Multicast Mode**	**IGMP-Multicast Mode**
ERS 8600 ERS 8300	NLB Unicast	Yes	No	No
ERS 8600	NLB Multicast[2]	No	Yes	No
ERS 8600 ERS 8300	Static Multicast entry	No	Yes[1]	No
ERS 8600 ERS 8300	Arp Multicast Flooding	No	Yes	No

Table 2.5.1 – Supported Avaya Ethernet Routing switch platforms

Note 1 – Required the system flag enhanced-operational-mode to be enabled and is not supported on legacy I/O modules in reference to the ERS8600

Note 2 – Please note that the per VLAN NLB mode of multicast does not work with the ERS 8300 on all failure scenarios, hence, is only supported on the ERS 8600

2.4.5 Switch Clustering – Topology 3

The following topology is supported when Ethernet Routing Switch 8600s are deployed as a SMLT core and Network Load Balancing cluster hosts and clients are directly connected and distributed between the ERS 8600s in the SMLT cluster. This topology supports Network Load Balancing clusters in unicast and multicast modes.

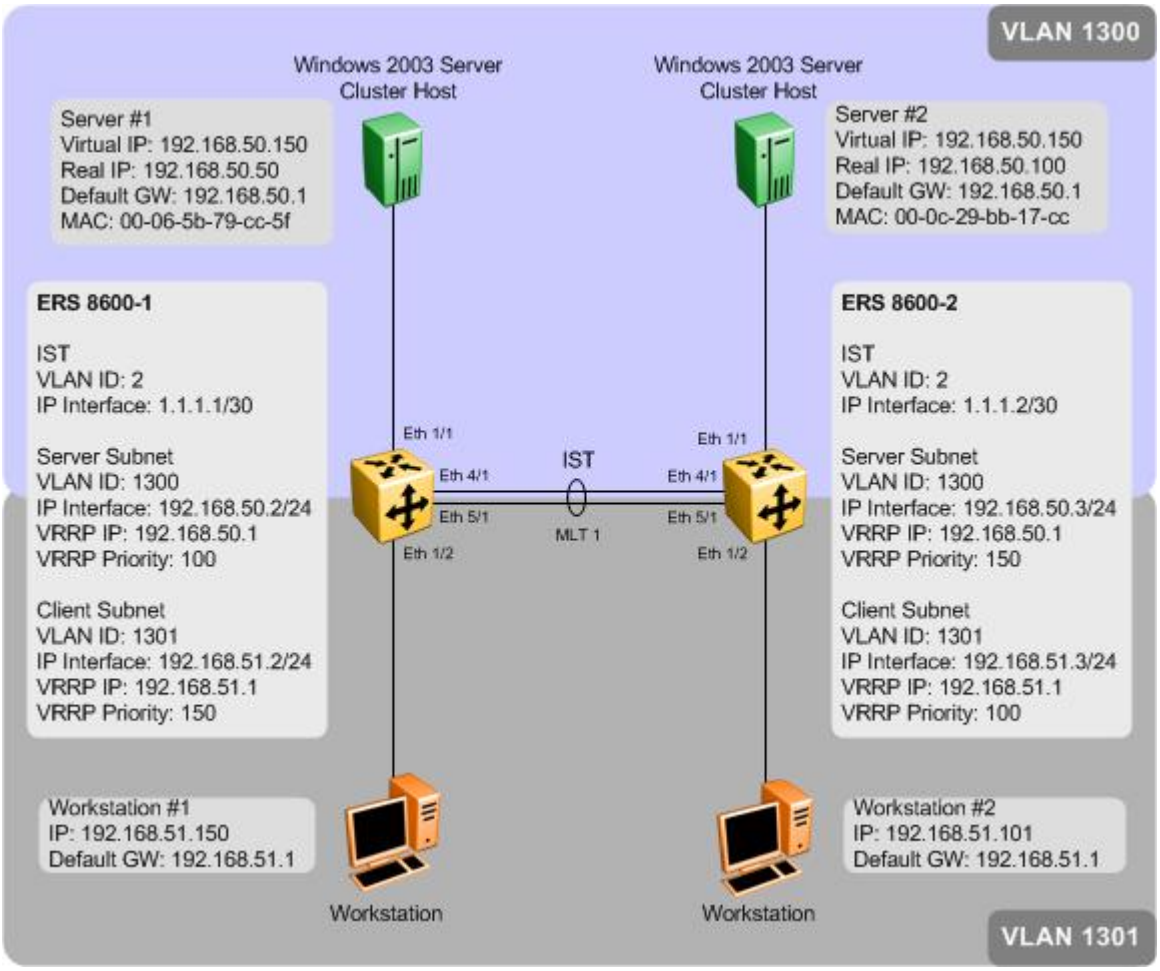

Figure 2.6 – Switch Clustering – Topology 3

2.4.6 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a SMLT core to support this topology:

Avaya Switch		Microsoft NLB Server Configuration		
Switch	**Configuration**	**Unicast Mode**	**Multicast Mode**	**IGMP-Multicast Mode**
ERS 8600 ERS 8300	NLB Unicast	Yes	No	No
ERS 8600	NLB Multicast[2]	No	Yes	No
ERS 8600 ERS 8300	Static Multicast entry	No	Yes[1]	No
ERS 8600 ERS 8300	Arp Multicast Flooding	No	Yes	No

Table 2.6.1 – Supported Avaya Ethernet Routing switch platforms

Note 1 – Required the system flag enhanced-operational-mode to be enabled and is not supported on legacy I/O modules in reference to the ERS 8600

Note 2 – Please note that the per VLAN NLB mode of multicast does not work with the ERS 8300 on all failure scenarios, hence, is only supported on the ERS 8600

2.4.7 Switch Clustering – Topology 4

The following topology is supported when Ethernet Routing Switch 8600s are deployed as a SMLT core and Network Load Balancing cluster hosts and clients are connected to a Layer 2 Ethernet switch that is SMLT connected to the SMLT cluster core. This topology supports Network Load Balancing clusters in unicast and multicast modes.

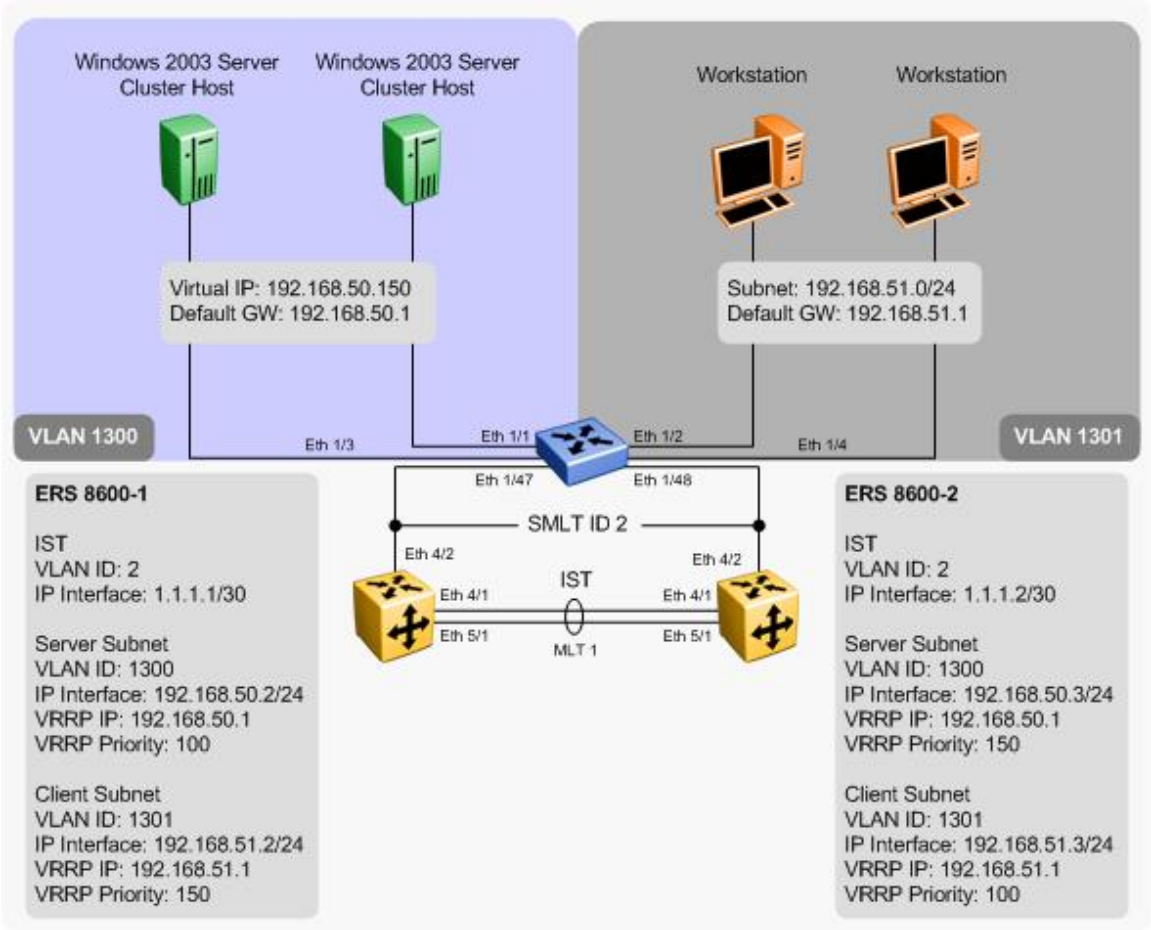

Figure 2.7 – Switch Clustering – Topology 4

2.4.8 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a SMLT core to support this topology:

Avaya Switch		Microsoft NLB Server Configuration		
Switch	**Configuration**	**Unicast Mode**	**Multicast Mode**	**IGMP-Multicast Mode**
ERS 8600 ERS 8300	NLB Unicast	Yes	No	No
ERS 8600 ERS 8300	Static Multicast entry	No	Yes[1]	No
ERS 8600 ERS 8300	ARP Multicast Flooding	No	Yes	No
ERS 8300	NLB Multicast	No	Yes[2]	No

Table 2.7.1 – Supported Avaya Ethernet Routing switch platforms

Note 1 – Required the system flag enhanced-operational-mode to be enabled and is not supported on legacy I/O modules

Note 2 – Normally, only one of the cluster switches will register the NLB Server ARP and forwarding port entries. Traffic will be forwarded on this cluster switch to the active NLB server either via the local SMLT or via the IST connection. If this node should fail, the peer SMLT cluster switch will forward traffic to the Microsoft NLB server via the SMLT connection.

2.4.9 Switch Clustering – Topology 5 (RSMLT Edge)

The following topology is supported when Ethernet Routing Switch 8600s are deployed as a SMLT core using RSMLT edge where Network Load Balancing cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster core. This topology supports Network Load Balancing clusters in unicast and multicast modes.

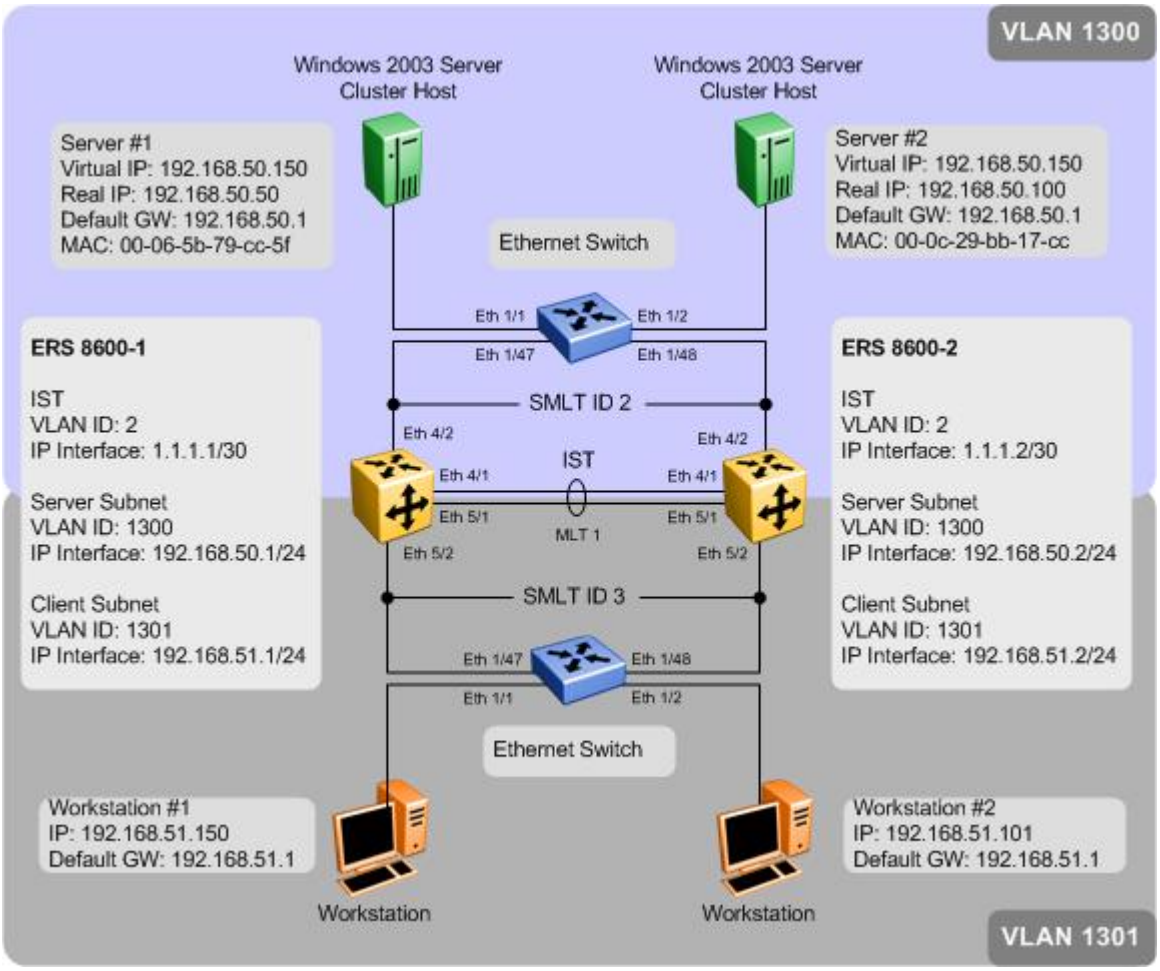

Figure 2.8 – Switch Clustering – Topology 5

2.4.10 Supported Avaya Switching Platforms

The following table provides a list of Avaya Ethernet Routing switching platforms that may be deployed as a SMLT core to support this topology:

Avaya Switch		Microsoft NLB Server Configuration		
Switch	**Configuration**	**Unicast Mode**	**Multicast Mode**	**IGMP-Multicast Mode**
ERS 8600 ERS 8300	NLB Unicast	Yes	No	No
ERS 8600 ERS 8300	Static Multicast entry	No	Yes[1]	No
ERS 8600 ERS 8300	ARP Multicast Flooding	No	Yes	No
ERS 8300	NLB Multicast	No	Yes[2]	No

Table 2.8.1 – Supported Avaya Ethernet Routing switch platforms

Note 1 – Required the system flag enhanced-operational-mode to be enabled and is not supported on legacy I/O modules

Note 2 – Normally, only one of the cluster switches will register the NLB Server ARP and forwarding port entries. Traffic will be forwarded on this cluster switch to the active NLB server either via the local SMLT or via the IST connection. If this node should fail, the peer SMLT cluster switch will forward traffic to the Microsoft NLB server via the SMLT connection.

2.5 Switch Clustering Configuration for Topologies 1 to 5

2.5.1 Enabling Per VLAN NLB Unicast Support

The following commands enables / disables per VLAN NLB unicast assuming VLAN 1300 is used to connect to the Microsoft NLB servers. For this example, we will assume ERS-8600-1 is using CLI and ERS-8600-2 is using PPCLI.

Step 1 – The following command enables / disables per VLAN NLB unicast, multicast or IGMP-multicast support for VLAN 1300
CLI
ERS-8600(config)#***interface vlan 1300*** ERS-8600(config-if)#***nlb-mode <igmp-mcast/multicast/unicast>***
PPCLI
ERS-8600# ***config vlan 1300 nlb-mode <disable/igmp-mcast/multicast/unicast>***
8600-1: Step 2 – The following CLI command enables per VLAN NLB unicast for VLAN 1300 assuming 8600-1 is using CLI
ERS-8600-1(config)#***interface vlan 1300*** ERS-8600-1(config-if)#***nlb-mode unicast*** ERS-8600-1(config-if)#***exit***
8600-2: Step 2 – The following PPCLI command enables per VLAN NLB unicast for VLAN 1300 assuming 8600-2 is using PPCLI
ERS-8600-2# ***config vlan 1300 nlb-mode unicast***

The per VLAN NLB unicast feature needs to be enabled for the VLAN that the cluster hosts are connected to on both Ethernet Routing Switch 8600s in the SMLT core.

2.5.1.1 Verify Operations – Per VLAN NLB Unicast Support

Use the following commands to verify operations assuming ERS-8600-1 is configured with CLI and ERS-8600-2 is configured with NNCLI

Step 1 – The following displays the status on the SMLT cluster when Network Load Balancing unicast support is enabled for VLAN 1300. Note this table may be different on both ERS 8600s depending on MLT port members and VLAN port assignment.

```
ERS-8600-1#show interfaces vlan nlb-mode 1300
ERS-8600-2# show vlan info nlb-mode
```

Result:

```
Response from 8600-1:

================================================================================
                                    Vlan Nlb
================================================================================
VLAN_ID  NLB_ADMIN_MODE NLB_OPER_MODE  PORT_LIST                    MLT_GROUPS
--------------------------------------------------------------------------------
1300     unicast        unicast        4/1-4/2,5/1

Total Entries: 1


Response from 8600-2:

================================================================================
                                    Vlan Nlb
================================================================================
VLAN_ID  NLB_ADMIN_MODE NLB_OPER_MODE  PORT_LIST                    MLT_GROUPS
--------------------------------------------------------------------------------
1300     unicast        unicast        4/1-4/2,5/1

Total Entries: 1
```

Step 2 – The following command displays the ARP entry for the NLB unicast IP address used in this example. The clusters virtual MAC address is a locally administered MAC address and starts with a 02:bf prefix.

```
ERS-8600-1#show ip arp 192.168.50.150
ERS-8600-2# show ip arp info 192.168.50.150
```

Result:

```
Response from 8600-1:

================================================================================
                            IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.150  02:bf:c0:a8:32:96  1300     -     DYNAMIC 2160

1 out of 111 ARP entries displayed


Response from 8600-2:
```

```
================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.150  02:bf:c0:a8:32:96  1300     -     DYNAMIC 2160

1 out of 115 ARP entries displayed
```

Step 3 – The following command displays the MAC entries for VLAN 1300. When unicast mode is enabled, Network Load Balancing binds a bogus MAC address on each hosts adapter which starts with 02 and contains the host ID in the second octet.

ERS-8600-1# ***show vlan mac-address-entry 1300***

ERS-8600-2# ***show vlan info fdb-entry 1300***

Result: If NLB Servers are connected to an SMLT Edge switch where the NLB MAC addresses are learned via the SMLT interface

```
Response from 8600-1:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 learned    00:01:81:29:1e:1c  IST        false   1      true
1300 self       00:80:2d:be:22:0e  Port-cpp   false   1      false
1300 learned    02:01:c0:a8:32:96  MLT-2      false   1      true
1300 learned    02:03:c0:a8:32:96  MLT-2      false   1      true

5 out of 92 entries in all fdb(s) displayed.


Response from 8600-2:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                    QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 self       00:01:81:29:1e:1c  Port-cpp   false   1      false
1300 learned    00:80:2d:be:22:0e  IST        false   1      true
1300 learned    02:01:c0:a8:32:96  MLT-2      false   1      false
1300 learned    02:03:c0:a8:32:96  MLT-2      false   1      false

5 out of 90 entries in all fdb(s) displayed.
```

Result: If NLB Servers are connected are directly connected to the SMLT cluster switches, only the locally attached server MAC address will be learned via the local port whereas the remote NLB server MAC will be learned via the IST.

```
Response from 8600-1:

================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                    QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 learned    00:01:81:29:1e:1c  IST        false   1      true
1300 learned    00:1b:25:e8:b4:00  4505       false   1      true
1300 learned    00:1b:25:e8:b4:32  4505       false   1      true
1300 learned    00:1d:42:36:10:1a  4505       false   1      false
1300 self       00:80:2d:be:22:0e  Port-cpp   false   1      false
1300 learned    02:01:c0:a8:32:96  Port-1/1   false   1      false
1300 learned    02:03:c0:a8:32:96  IST        false   1      true

8 out of 95 entries in all fdb(s) displayed.


Response from 8600-2:

================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                    QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 self       00:01:81:29:1e:1c  Port-cpp   false   1      false
1300 learned    00:1b:25:e8:b4:00  4505       false   1      false
1300 learned    00:1b:25:e8:b4:32  4505       false   1      false
1300 learned    00:1d:42:36:10:1a  4505       false   1      true
1300 learned    00:80:2d:be:22:0e  IST        false   1      true
1300 learned    02:01:c0:a8:32:96  IST        false   1      true
1300 learned    02:03:c0:a8:32:96  Port-1/1   false   1      false

8 out of 93 entries in all fdb(s) displayed.
```

2.5.2 Enabling Per VLAN NLB Multicast Support

The following commands enables / disables per VLAN NLB multicast assuming VLAN 1300 is used to connect to the Microsoft NLB servers. For this example, we will assume ERS-8600-1 is using CLI and ERS-8600-2 is using PPCLI.

Step 1 – The following command enables / disables per VLAN NLB unicast, multicast or IGMP-multicast support for VLAN 1300
CLI
ERS-8600(config)#***interface vlan 1300*** ERS-8600(config-if)#***nlb-mode <igmp-mcast/multicast/unicast>***
PPCLI
ERS-8600# ***config vlan 1300 nlb-mode <disable/igmp-mcast/multicast/unicast>***
8600-1: Step 2 – The following CLI command enables per VLAN NLB unicast for VLAN 1300 assuming 8600-1 is using CLI
ERS-8600-1(config)#***interface vlan 1300*** ERS-8600-1(config-if)#***nlb-mode multicast*** ERS-8600-1(config-if)#***exit***
8600-2: Step 2 – The following PPCLI command enables per VLAN NLB unicast for VLAN 1300 assuming 8600-2 is using PPCLI
ERS-8600-2# ***config vlan 1300 nlb-mode multicast***

The per VLAN NLB multicast feature needs to be enabled for the VLAN that the cluster hosts are connected to on both Ethernet Routing Switch 8600s in the SMLT core.

Please note, per VLAN NLB Multicast configuration is only supported for Switch Cluster Topology 2 and 3 where the Microsoft NLB servers are directly connected to the Switch Cluster instead of going through a SMLT/SLT attached edge switch.

2.5.2.1 Verify Operations - Per VLAN NLB Multicast Support

Use the following commands to verify operations assuming ERS-8600-1 is configured with CLI and ERS-8600-2 is configured with NNCLI

Step 1 – The following displays the status on the SMLT cluster when Network Load Balancing multicast support is enabled for VLAN 1300. Note this table may be different on both ERS 8600s depending on MLT port members and VLAN port assignment.

```
ERS-8600-1#show interfaces vlan nlb-mode 1300
ERS-8600-2# show vlan info nlb-mode 1300
```

Result:

```
Response from 8600-1:

================================================================================
                                    Vlan Nlb
================================================================================
VLAN_ID  NLB_ADMIN_MODE NLB_OPER_MODE  PORT_LIST                  MLT_GROUPS
--------------------------------------------------------------------------------
1300     multicast      multicast      1/1                        1

Total Entries: 1


Response from 8600-2:

================================================================================
                                    Vlan Nlb
================================================================================
VLAN_ID  NLB_ADMIN_MODE NLB_OPER_MODE  PORT_LIST                  MLT_GROUPS
--------------------------------------------------------------------------------
1300     unicast        unicast        1/1                        1

Total Entries: 1
```

Step 2 – The following command displays the ARP entrie for the NLB multicast IP address and the Microsoft NLB server real IP addresses. The clusters virtual multicast MAC address is a locally administered MAC address and starts with a 03:bf prefix.

```
ERS-8600-1#show ip arp 192.168.50.0
ERS-8600-2# show ip arp info 192.168.50.0
```

Result:

```
Response from 8600-1:

================================================================================
                            IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.2    00:01:81:28:86:12  1300     -    LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -    LOCAL   2160
192.168.50.3    00:e0:7b:bc:22:30  1300  MLT 1   DYNAMIC 2041
192.168.50.1    00:00:5e:00:01:82  1300     -    LOCAL   2160
192.168.50.50   00:06:5b:79:cc:5f  1300    1/1   DYNAMIC 2045
192.168.50.100  00:0c:29:bb:17:cc  1300  MLT 1   DYNAMIC 2053
```

```
192.168.50.150  03:bf:c0:a8:32:96  1300    -     DYNAMIC 2158

7 out of 135 ARP entries displayed


Response from 8600-2:

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.3    00:e0:7b:bc:22:30  1300    -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300    -     LOCAL   2160
192.168.50.2    00:01:81:28:86:12  1300  MLT 1   DYNAMIC 2041
192.168.50.1    00:00:5e:00:01:82  1300    -     LOCAL   2160
192.168.50.50   00:06:5b:79:cc:5f  1300  MLT 1   DYNAMIC 2046
192.168.50.100  00:0c:29:bb:17:cc  1300    1/1   DYNAMIC 2053
192.168.50.150  03:bf:c0:a8:32:96  1300    -     DYNAMIC 2158

7 out of 152 ARP entries displayed
```

Step 3 – The following command displays the MAC entries for VLAN 1300. When unicast mode is enabled, Network Load Balancing binds a bogus MAC address on each hosts adapter which starts with 02 and contains the host ID in the second octet.

ERS-8600-1# ***show vlan mac-address-entry 1300***

ERS-8600-2# ***show vlan info fdb-entry 1300***

Result: The Microsoft NLB servers real IP addresses should be displayed. Since the ERS 8600, with the per VLAN NLB multicast parameter enabled, only supports local attached servers connected to the SMLT cluster, the Microsoft NLB Server MAC addresses will be learned either via the local port or from the IST.

```
Response from 8600-1:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN                MAC                                    QOS    SMLT
ID   STATUS         ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self           00:00:5e:00:01:82  Port-cpp   false   1      false
1300 self           00:01:81:28:86:12  Port-cpp   false   1      false
1300 learned        00:06:5b:79:cc:5f  Port-1/1   false   1      false
1300 learned        00:0c:29:bb:17:cc  IST        false   1      true
1300 learned        00:e0:7b:bc:22:30  IST        false   1      true

Response from 8600-2:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN                MAC                                    QOS    SMLT
ID   STATUS         ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self           00:00:5e:00:01:82  Port-cpp   false   1      false
1300 learned        00:01:81:28:86:12  IST        false   1      true
1300 learned        00:06:5b:79:cc:5f  IST        false   1      true
1300 learned        00:0c:29:bb:17:cc  Port-1/1   false   1      false
1300 self           00:e0:7b:bc:22:30  Port-cpp   false   1      false
```

2.5.3 Enabling Static Multicast entries

If NLB Multicast is enabled on the Microsoft NLB servers, the multicast MAC address can be statically entered on both ERS 8600 cluster switches.

Step 1 – The following command enables / disables per VLAN NLB unicast, multicast or IGMP-multicast support for VLAN 1300
CLI
ERS-8600(config)#***vlan static-mcastmac <vlan id> <multicast mac address> <port> mlt <mlt id>*** ERS-8600(config)#***ip arp static-mcast <ip address> <multicast mac address> vid <vlan id> port <slot/port> <mlt id>***
PPCLI
ERS-8600# ***config ip arp static-mcastmac add mac <multicast mac address> ip <ip address> vlan <vlan id> port <slot/port> mlt <mlt id>***
8600-1: Step 2 – The following CLI commands adds a static arp entry of the Microsoft Multicast NLB address. Please note that the IST MLT ID is 1 and the SMLT ID is 2 as used in this example
ERS-8600-1(config)#***vlan static-mcastmac 1300 03:bf:c0:a8:32:96 mlt 1,2*** ERS-8600-1(config)#***ip arp static-mcast 192.168.50.150 03:bf:c0:a8:32:96 vid 1300***
8600-2: Step 2 – The following CLI commands adds a static arp entry of the Microsoft Multicast NLB address. Please note that the IST MLT ID is 1 and the SMLT ID is 2 as used in this example
ERS-8600-2# ***config ip arp static-mcastmac add mac 03:bf:c0:a8:32:96 ip 192.168.50.150 vlan 1300 mlt 1,2***

Please note, the ERS 8600 enchanced-operation-mode flag must be enabled to support static multicast entries. This feature can be enabled by using the PPCLI command *config sys set flags enhanced-operational-mode true* or the CLI command *sys flags enhanced-operational-mode*. Also note that only R and RS modules are supported in enchanced-operational-mode. If you have legacy models, ARP multicast flooding can be used.

If SLT is used instead of SMLT, simply enter the port number and IST MLT ID. For example, the PPCLI command will be *config ip arp static-mcastmac add mac 03:bf:c0:a8:32:96 ip 192.168.50.150 vlan 1300 port 4/2 mlt 1* and the CLI command will be *vlan static-mcastmac 1300 03:bf:c0:a8:32:96 4/2 mlt 1*.

2.5.3.1 Verify Operations - Static Multicast entries

Use the following commands to verify operations assuming ERS-8600-1 is configured with CLI and ERS-8600-2 is configured with NNCLI

Step 1 – The following command displays the ARP entry for the NLB unicast IP address used in this example. As per the configuration used in this example, the NLB multicast MAC

```
ERS-8600-1#show ip arp static-mcastmac
ERS-8600-2# show ip arp static-mcastmac
```

Result:

```
Response from 8600-1:

================================================================================
                    IP Static Multicast MAC Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT            MLT ID
--------------------------------------------------------------------------------
192.168.50.150  03:bf:c0:a8:32:96  1300  -               1,2


Response from 8600-2:

================================================================================
                    IP Static Multicast MAC Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT            MLT ID
--------------------------------------------------------------------------------
192.168.50.150  03:bf:c0:a8:32:96  1300  -               1,2
```

Step 3 – The following command displays the ARP entries for VLAN 1300. The actual NIC MAC address should be displayed for both Microsoft NLB servers via MLT 2 under normal operations.

```
ERS-8600-1#show ip arp 192.168.50.0
ERS-8600-2# show ip arp info 192.168.50.0
```

Result: If NLB Servers are connected to an SMLT Edge switch where the NLB MAC addresses are learned via the SMLT interface

```
Response from 8600-1:

================================================================================
                             IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.2    00:80:2d:be:22:09  1300    -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300    -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300    -     LOCAL   2160
192.168.50.3    00:01:81:29:1e:07  1300  MLT 1   DYNAMIC 2041
192.168.50.50   00:06:5b:79:cc:5f  1300  MLT 2   DYNAMIC 2145
192.168.50.100  00:0c:29:bb:17:cc  1300  MLT 2   DYNAMIC 2101

================================================================================
```

```
                              IP Arp Extn - GlobalRouter
================================================================================
   MULTICAST-MAC-FLOODING                AGING        ARP-THRESHOLD
--------------------------------------------------------------------------------
        disable                           360             500
6 out of 78 ARP entries displayed


Response from 8600-2:

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.3    00:01:81:29:1e:07  1300     -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300     -     LOCAL   2160
192.168.50.2    00:80:2d:be:22:09  1300  MLT 1    DYNAMIC 2047
192.168.50.50   00:06:5b:79:cc:5f  1300  MLT 2    DYNAMIC 2055
192.168.50.100  00:0c:29:bb:17:cc  1300  MLT 2    DYNAMIC 2107
```

Result: If NLB Servers are connected are directly connected to the SMLT cluster switches, only the locally attached server MAC address will be learned via the local port whereas the remote NLB server MAC will be learned via the IST.

```
Response from 8600-1:

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.2    00:80:2d:be:22:09  1300     -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300     -     LOCAL   2160
192.168.50.3    00:01:81:29:1e:07  1300  MLT 1    DYNAMIC 2041
192.168.50.50   00:06:5b:79:cc:5f  1300  1/1      DYNAMIC 2145
192.168.50.100  00:0c:29:bb:17:cc  1300  MLT 1    DYNAMIC 2101

================================================================================
                              IP Arp Extn - GlobalRouter
================================================================================
   MULTICAST-MAC-FLOODING                AGING        ARP-THRESHOLD
--------------------------------------------------------------------------------
        disable                           360             500
6 out of 78 ARP entries displayed


Response from 8600-2:

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.3    00:01:81:29:1e:07  1300     -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300     -     LOCAL   2160
192.168.50.2    00:80:2d:be:22:09  1300  MLT 1    DYNAMIC 2047
192.168.50.50   00:06:5b:79:cc:5f  1300  MLT 1    DYNAMIC 2055
192.168.50.100  00:0c:29:bb:17:cc  1300  1/1      DYNAMIC 2107
```

2.5.4 Enabling Global ARP Multicast MAC Flooding

For Multicast mode Network Load Balancing, the IP ARP Multicast Flooding parameter can be enabled on both SMLT cluster switches. The Per VLAN Multicast and IGMP-Multicast NLB modes are not supported.

8600-1: Step 1 – The following CLI command enables the global ARP multicast MAC flooding feature
ERS-8600-1(config)#***ip arp multicast-mac-flooding enable***
8600-2: Step 1 – The following PPCLI command enables the global ARP multicast MAC flooding feature
ERS-8600-2# ***config ip arp multicast-mac-flooding enable***

The global ARP multicast MAC flooding feature needs to be enabled on both Ethernet Routing Switch 8600s in the SMLT cluster.

2.5.4.1 Verify Operations - Global ARP Multicast MAC Flooding

Step 1 – Verify that the NLB mode configured is set to multicast. The NLB operational state should also display multicast if configured correctly with uplink port to the NLB server.
ERS-8600-1#***show ip arp*** ERS-8600-2# ***config ip arp info***
Result:

```
Response from 8600-1:

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
<arp entries>

================================================================================
                           IP Arp Extn - GlobalRouter
================================================================================
   MULTICAST-MAC-FLOODING               AGING        ARP-THRESHOLD
--------------------------------------------------------------------------------
        enable                           360            500

Response from 8600-2:

multicast-mac-flooding :  enable
aging : 360 (min)
arpreqthreshold : 500
delete : N/A
add :
```

Step 2 – The following command displays the ARP entry for the NLB multicast IP address used in this example. The clusters virtual MAC address is multicast MAC address assigned by Microsoft and will start with a 03:bf prefix.

ERS-8600-1#***show ip arp 192.168.50.0***

ERS-8600-2# ***show ip arp info 192.168.50.0***

Result: The results will display the actual port number for the real IP address even for SMLT connections

```
Response from 8600-1:

================================================================================
                            IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.2    00:80:2d:be:22:09  1300     -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300     -     LOCAL   2160
192.168.50.3    00:01:81:29:1e:07  1300  MLT 1   DYNAMIC 2112
192.168.50.50   00:06:5b:79:cc:5f  1300    4/2   DYNAMIC 2138
192.168.50.100  00:0c:29:bb:17:cc  1300    4/2   DYNAMIC 2138
192.168.50.150  03:bf:c0:a8:32:96  1300     -     DYNAMIC 2138

================================================================================
                          IP Arp Extn - GlobalRouter
================================================================================
   MULTICAST-MAC-FLOODING                AGING        ARP-THRESHOLD
--------------------------------------------------------------------------------
         enable                           360             500
7 out of 75 ARP entries displayed

Response from 8600-2:

================================================================================
                            IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN     PORT TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
192.168.50.3    00:01:81:29:1e:07  1300     -     LOCAL   2160
192.168.50.255  ff:ff:ff:ff:ff:ff  1300     -     LOCAL   2160
192.168.50.1    00:00:5e:00:01:82  1300     -     LOCAL   2160
192.168.50.2    00:80:2d:be:22:09  1300  MLT 1   DYNAMIC 2122
192.168.50.150  03:bf:c0:a8:32:96  1300     -     DYNAMIC 2147
192.168.50.50   00:06:5b:79:cc:5f  1300    4/2   DYNAMIC 2150
192.168.50.100  00:0c:29:bb:17:cc  1300    4/2   DYNAMIC 2150
```

Step 3 – The following displays the MAC entries for VLAN 1300. When multicast mode is enabled on the NLB servers, the real MAC address of the NLB server interface will be used.

ERS-8600-1#***show vlan mac-address-entry 1300***

ERS-8600-2# ***show vlan info fdb-entry 1300***

Result: If NLB Servers are connected to an SMLT Edge switch where the NLB MAC addresses are learned via the SMLT interface

```
Response from 8600-1:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                     QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 learned    00:01:81:29:1e:1c  IST        false   1      true
1300 learned    00:06:5b:79:cc:5f  MLT-2      false   1      true
1300 learned    00:0c:29:bb:17:cc  MLT-2      false   1      true
1300 self       00:80:2d:be:22:0e  Port-cpp   false   1      false

Response from 8600-2:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                     QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 self       00:01:81:29:1e:1c  Port-cpp   false   1      false
1300 learned    00:06:5b:79:cc:5f  MLT-2      false   1      true
1300 learned    00:0c:29:bb:17:cc  MLT-2      false   1      true
1300 learned    00:80:2d:be:22:0e  IST        false   1      true
```

Result: If NLB Servers are connected to an SMLT Edge switch where the NLB MAC addresses are learned via the SMLT interface

```
Response from 8600-1:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                     QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
1300 learned    00:01:81:29:1e:1c  IST        false   1      true
1300 learned    00:1b:25:e8:b4:00  4505       false   1      true
1300 learned    00:1b:25:e8:b4:32  4505       false   1      true
1300 learned    00:1d:42:36:10:1a  4505       false   1      false
1300 self       00:80:2d:be:22:0e  Port-cpp   false   1      false
1300 learned    02:01:c0:a8:32:96  Port-1/1   false   1      false
1300 learned    02:03:c0:a8:32:96  IST        false   1      true

8 out of 95 entries in all fdb(s) displayed.


Response from 8600-2:

================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                                     QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE  MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1300 self       00:00:5e:00:01:82  Port-cpp   false   1      false
```

```
        1300 self       00:01:81:29:1e:1c  Port-cpp   false   1      false
        1300 learned    00:1b:25:e8:b4:00  4505       false   1      false
        1300 learned    00:1b:25:e8:b4:32  4505       false   1      false
        1300 learned    00:1d:42:36:10:1a  4505       false   1      true
        1300 learned    00:80:2d:be:22:0e  IST        false   1      true
        1300 learned    02:01:c0:a8:32:96  IST        false   1      true
        1300 learned    02:03:c0:a8:32:96  Port-1/1   false   1      false

        8 out of 93 entries in all fdb(s) displayed.
```

3. Appendix

3.1 Creating a Network Load Balancing Cluster

The following section demonstrates how to create a Network Load Balancing Cluster using two Windows 2003 servers to provide high available HTTP web services.

The Windows 2003 Servers used in the following examples were configured as follows:

- The Windows 2003 servers have been updated with the latest Service Pack 1 and all the current updates applied.
- Although you can use one network adaptor, for best performance it is recommended that you have two 10/100/1000BASE-T Ethernet network adaptors installed. If you use only one adaptor, it is recommended to select Multicast which allows both the NLB and native traffic to be handled by the adapter. In Unicast mode, NLB will take over the network adapter it is bound to and does not allow any addition network traffic through it.
- Internet Information Services (IIS) is installed and operational with a default web site tied to the Clusters Virtual IP Address.

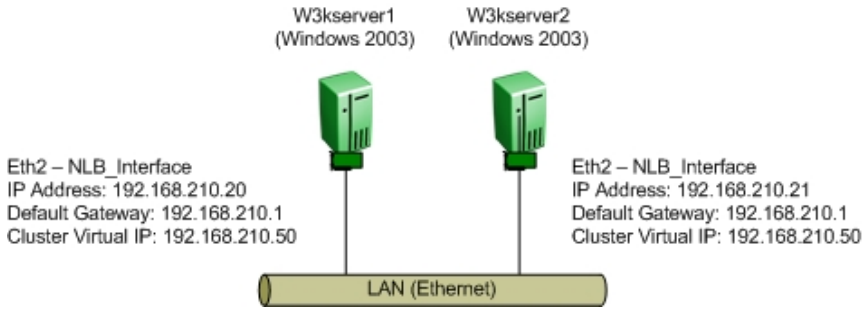

Figure 3.1 – Windows 2003 Server Cluster

Step 1 – Starting the Microsoft Network Load Balancing Manager Application

The Microsoft Network Load Balancing Manager Application may be opened by clicking **Start**, **Run** and typing **nlbmgr** and then clicking OK or go to **Start -> All Programs -> Administrative Tools -> Network Load Balancing Manager.**

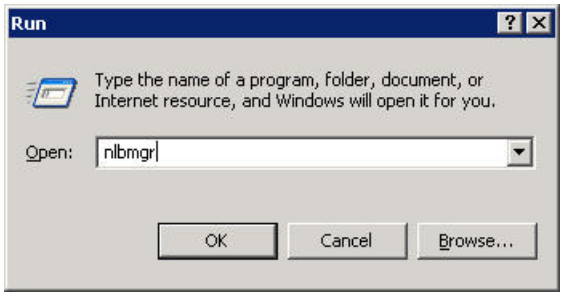

Step 2 – Creating a New Cluster

To create a new cluster, in the Microsoft **Network Load Balancing Manager Application** click **Cluster** then **New**.

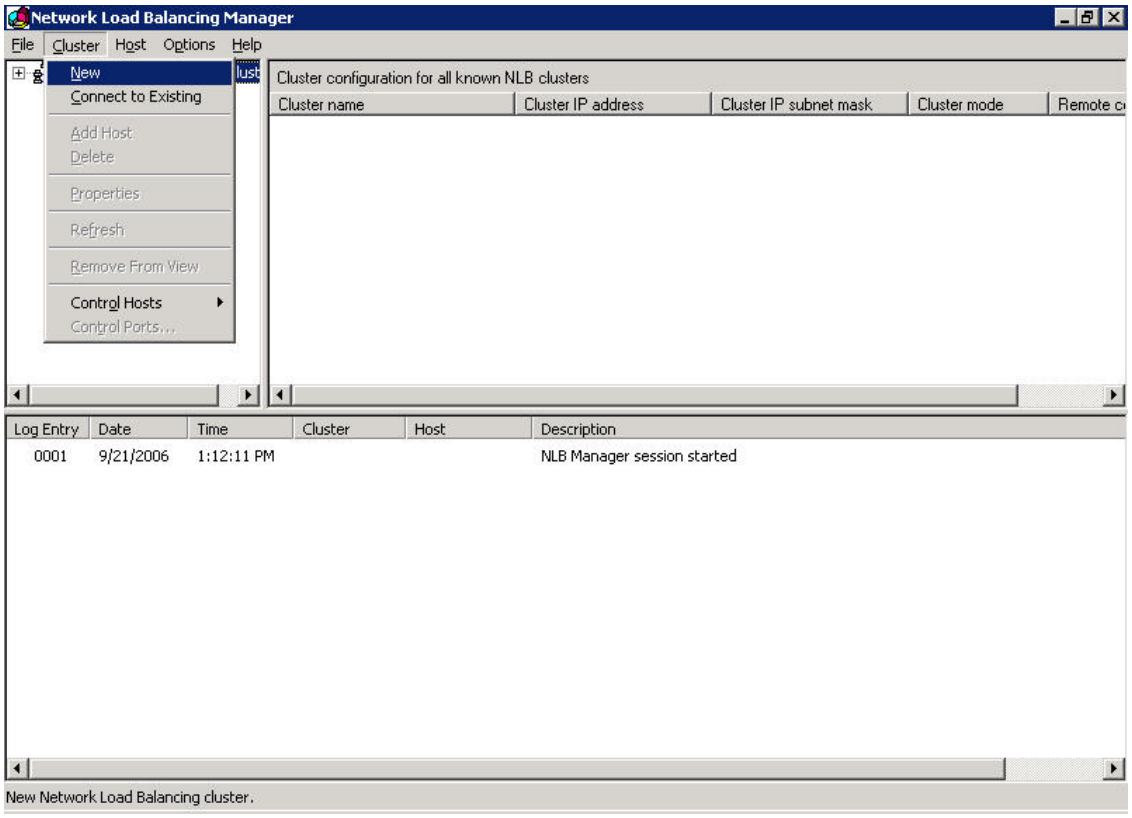

Step 3 – Specify Cluster Parameters and Operational Mode

In the **Cluster Parameters** window, specify the clusters virtual **IP address**, **Subnet Mask** and optionally **Full Internet name** that will used to address this cluster. The Full Internet name is used only for reference. Specify the operational mode for the cluster which can be set to unicast (default), multicast or IGMP-multicast. Note that the **Network address** field will change depending on the cluster operational mode specified.

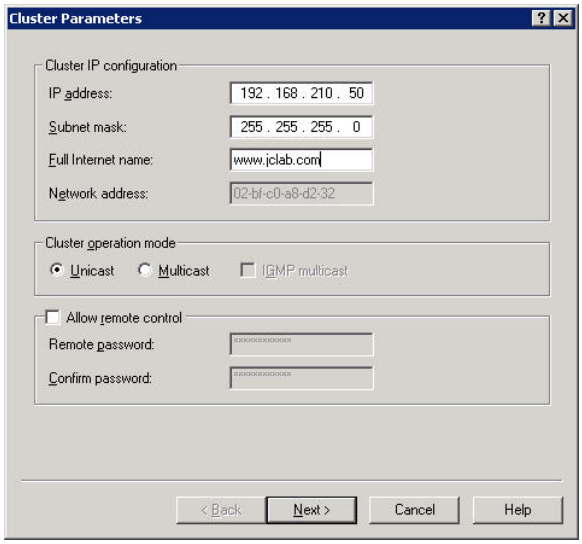

Cluster Parameters window with unicast operational mode is enabled

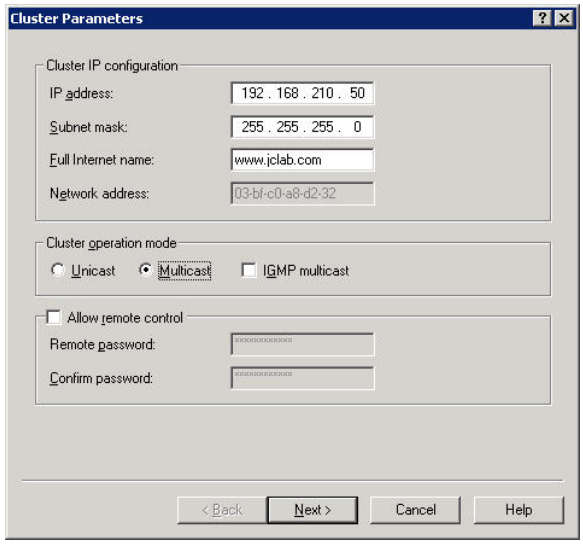

Cluster Parameters window with multicast operational mode is enabled

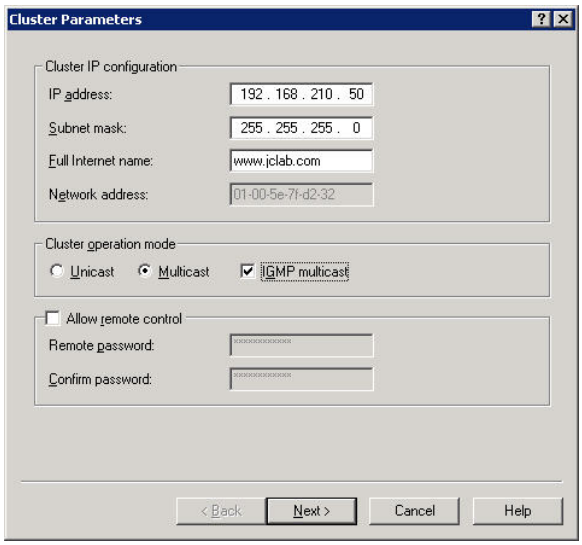

Cluster Parameters window with IGMP-multicast operational mode is enabled

If the cluster operational mode is set to multicast, it is possible to change the operational mode to IGMP-multicast at a later time by simply checking the IGMP multicast checkbox.

Step 4 – Cluster IP Addresses

Click on **Next** to skip adding a Cluster IP address. This is only required if you need additional IP address to be load balanced via multiple sites using different IP addresses.

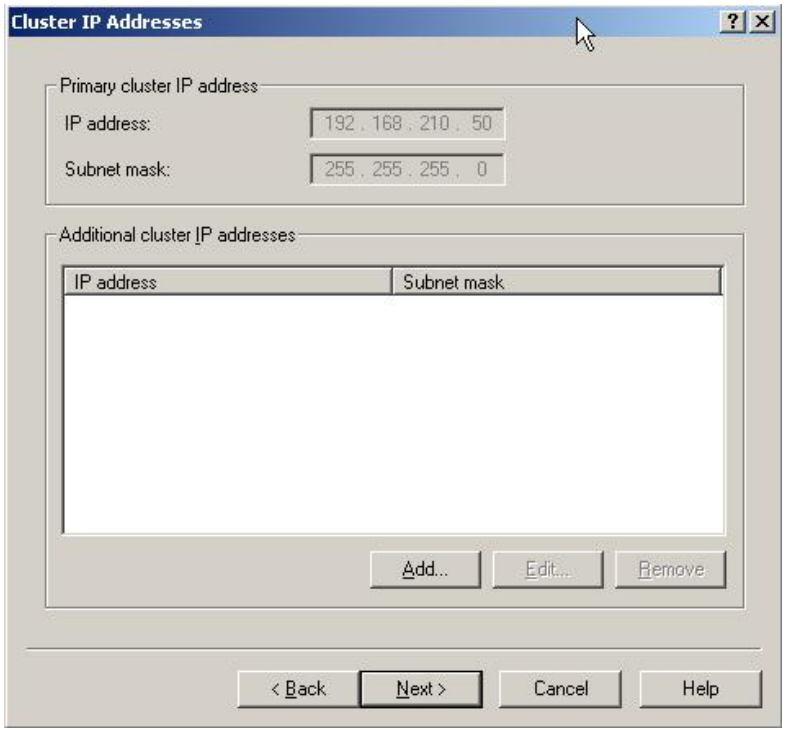

Step 5 – Specify Port Rules

The **Port Rules** window defines the traffic that the load balancing cluster will service as well as how traffic is distributed between hosts. The default port rule will load balance all TCP and UDP traffic using ports 0 through 65535. Administrators may specify a single rule or multiple port rules if the application requires it such as a Web server that requires HTTP and HTTPS. For this example we will modify the default port rule to support HTTP traffic by clicking **Edit**.

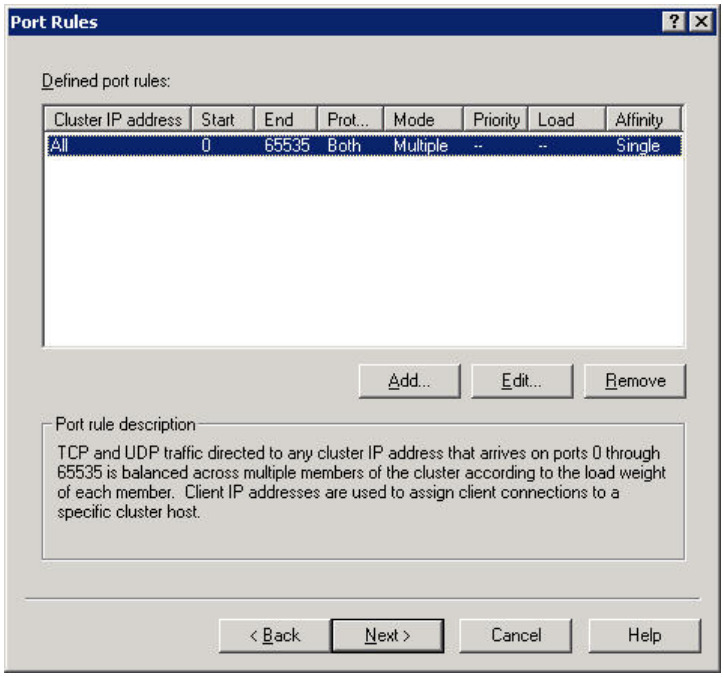

Modify the **Port range** fields so that both the **From** and **To** fields have a value set to **80** (HTTP) with Affinity value of None. Click **OK**. Click on **Add** to another value set to **443** (SSL) with Affinity value of **Single**. Additional port rule parameters are provided in table 3.1.

Step 6 – Specify Port Rules, con't

The default port rule has now been modified so that the Network Load Balancing cluster will load balance HTTP traffic. Click **Next**.

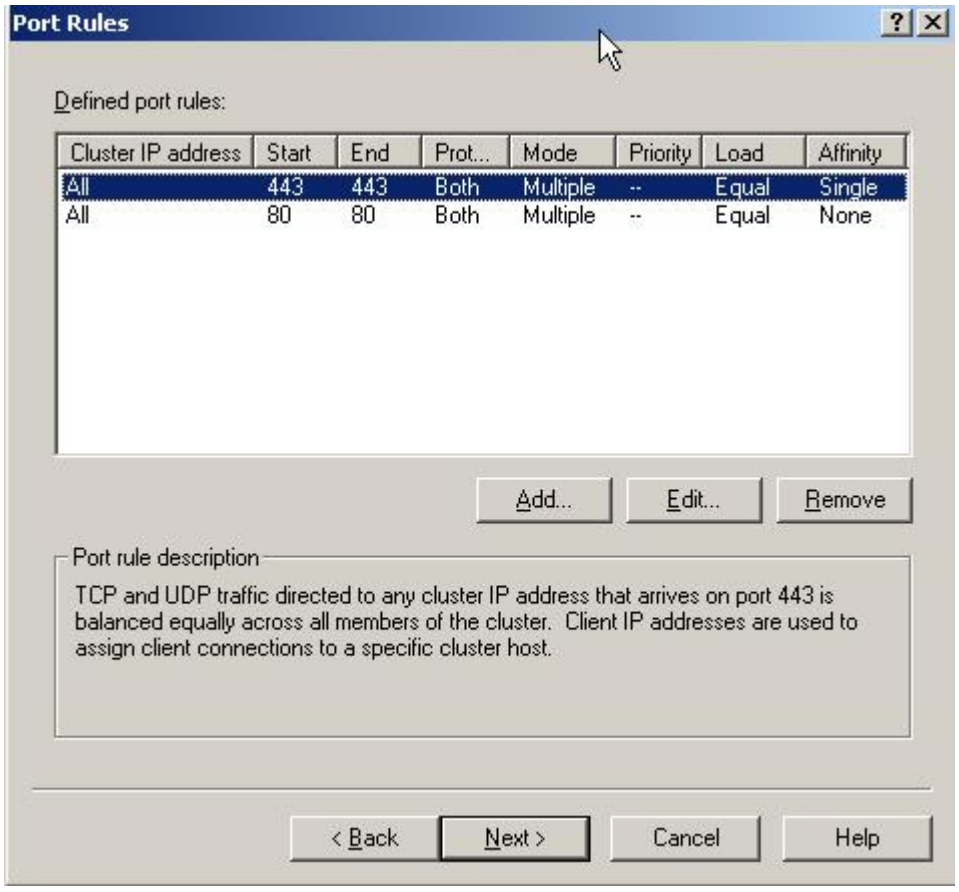

Step 7 – Adding the first Host to the Cluster

In the **Connect** window we will add the first host to the cluster. For this example we have two Windows 2003 servers with the hostnames w3kserver1 and w3kserver2. In the Host field type the hostname for the first server in the cluster and click **Connect**. Once connected a list of interfaces will be displayed. Highlight the **Interface name** where Network Load Balancing will be bound to and click **Next**

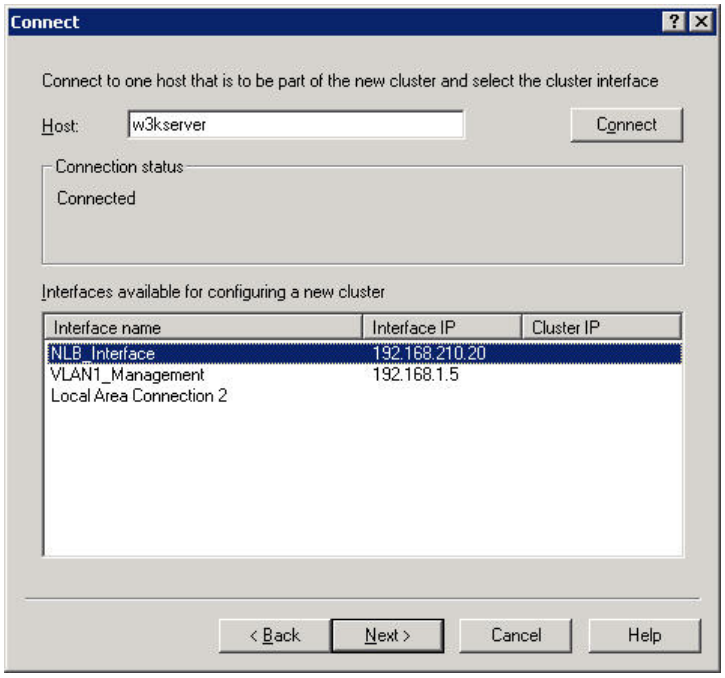

Step 8 – Setting Host Priority

In the **Host Parameters** window set the **Priority** for the host to **1**. This value needs to be unique for each host in the cluster. Optionally modify the **Default state** for the cluster host if you do not wish Network Load Balancing to be immediately started. Click **Finish**

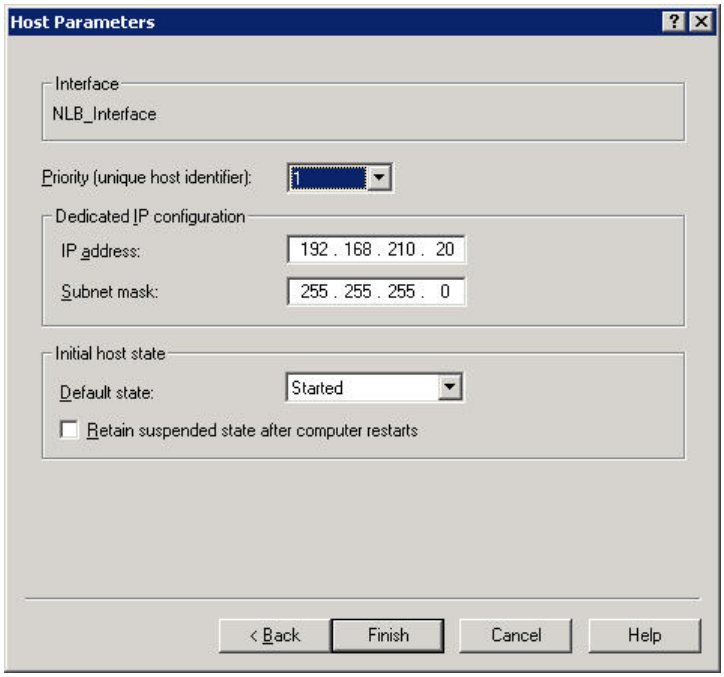

Step 9 – Adding Additional Hosts to the Cluster

To add the second host to the cluster, in the **Network Load Balancing Manager** highlight the **Domain name** of the cluster and then **right click** and click **Add Host To Cluster**.

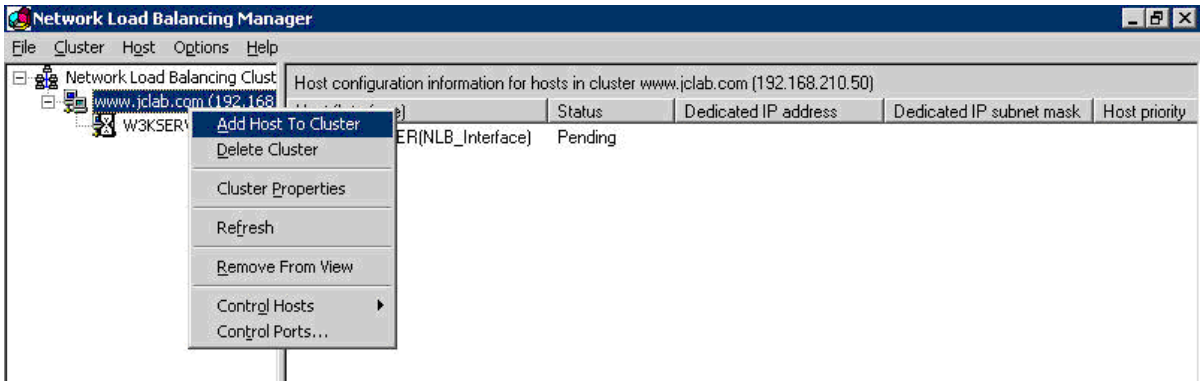

Step 10 – Adding Additional Hosts to the Cluster, con't

In the **Connect** window in the Host field type the hostname for the second server in the cluster and click **Connect**. Once connected a list of interfaces will be displayed. Highlight the **Interface name** where Network Load Balancing will be bound to and click **Next**.

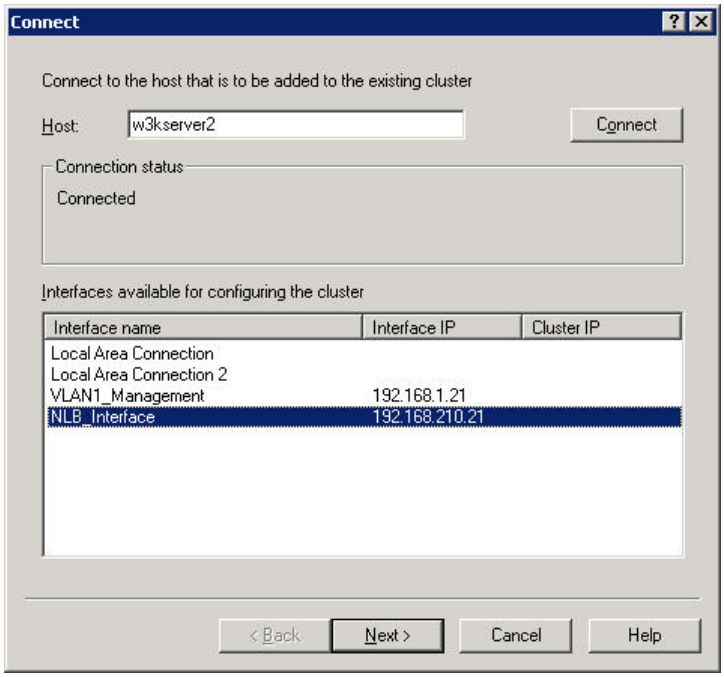

Step 11 – Setting Additional Host Priority

In the **Host Parameters** window set the **Priority** for the host to **2**. This value needs to be unique for each host in the cluster. Optionally modify the **Default state** for the cluster host if you do not wish Network Load Balancing to be immediately started. Click **Finish**.

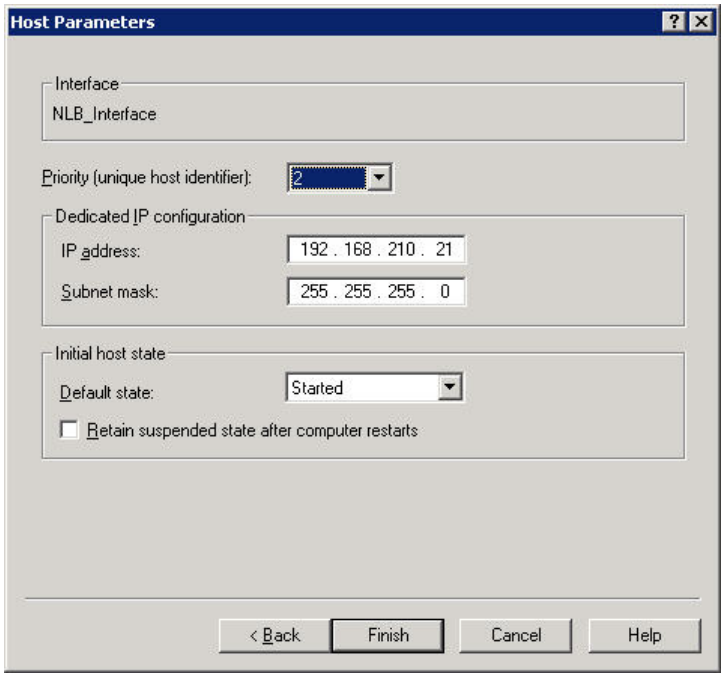

Step 12 – Completed

The cluster is created and once converged all operational hosts will be displayed in **Network Load Balancing Manager** window in a **green** state. Additionally details for all known clusters as well as log entries are displayed in this window.

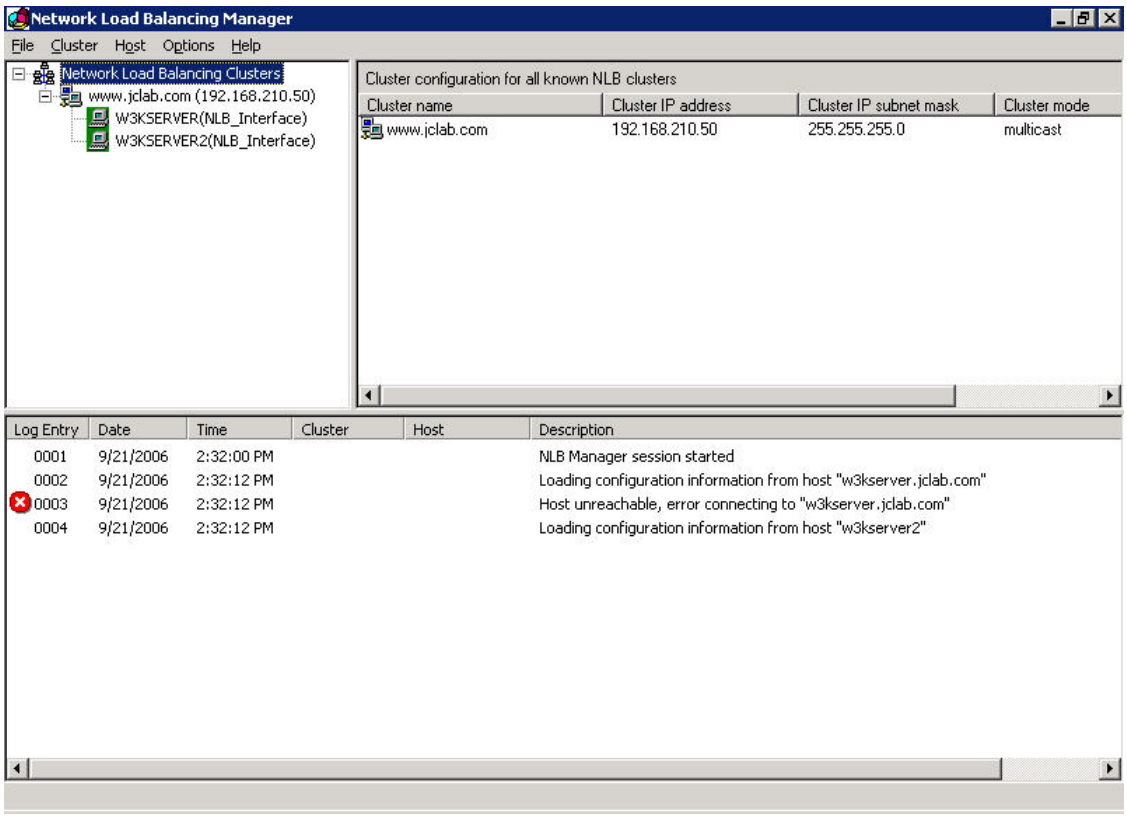

The above configuration assumes you have DNS configured on both NLB servers with the appropriate server names. If DNS is not enabled, you will need to modify the host file (C:\winnt\system32\drivers\etc\hosts) and add appropriate names of each server. If you are using NLB in Unicast mode and you cannot connect to more than one server, please refer to Microsoft article 898867 and 193602.

The following table provides a detailed overview the Port Rule parameters available in the **Add/Edit Port Rule** window.

Parameter	Description
Cluster IP Address	Specifies options regarding which cluster IP addresses that the port rule should cover.
All	Specifies whether the port rule is a global port rule and will cover all cluster IP addresses associated with the particular Network Load Balancing cluster.
Port Range	Specifies the start and end of the port range for the selected port rule. Port numbers in a range of 0 to 65,535 are currently supported. The default port range is 0 to 65,535.
Protocols	Specifies the IP protocol that a port rule should cover: Transmission Control Protocol (TCP), User Datagram Protocol (UDP), or both. Only the network traffic for the specified protocol is affected by the rule. The default host will handle all traffic not covered by a port rule.
Multiple Host	Specifies whether multiple hosts in the cluster handle network traffic for the associated port rule.
Affinity	Specifies how requests are routed to a specific server.
Affinity: None	Specifies whether multiple connections from the same client IP address can be handled by different hosts. Disabling affinity allows for more effective load balancing because it allows multiple connections from the same client to be handled concurrently by different cluster hosts. To maximize scaled performance when client affinity is not needed, disable affinity by selecting None. However, in order to allow Network Load Balancing to properly handle IP fragments, you should avoid using None when selecting UDP or Both for your protocol setting.
Affinity: Single	Specifies that Network Load Balancing direct multiple requests - Transmission Control Protocol (TCP) connections or User Datagram Protocol (UDP) datagram's - from the same client Internet Protocol (IP) address to the same cluster host. Using Single affinity ensures that only one cluster host will handle all connections that are part of the same client session. This is important if the server program running on the cluster host maintains session state (such as "server cookies" or SSL connections for HTTPS) between connections.
Affinity: Class C	Specifies that Network Load Balancing direct multiple requests - Transmission Control Protocol (TCP) connections or User Datagram Protocol (UDP) datagram's - from the same TCP/IP Class C address range to the same cluster host.
Affinity: Single Host	Specifies that network traffic for the associated port rule be handled by a single host in the cluster according to the specified handling priority. This filtering mode provides port specific fault tolerance for the handling of network traffic.
Disable this Port Range	Specifies whether all network traffic for the associated port rule will be blocked.

Table 3.1 – Network Load Balancing Port Rule Options

4. Software Baseline:

Device	Software Release
Windows 2003 Advanced Server	Service Pack 2 and Latest Patches
Ethernet Routing Switch 8600	Release 5.0.5 & 5.1.1.1
Ethernet Routing Switch 8300	Release 4.2.1
Ethernet Routing Switch 5500	Release 6.1.2
Ethernet Routing Switch 1600	Release 2.1.7
Ethernet Routing Switch 4500	Release 5.3.2

Table 4.1 – Software Baseline

5. Reference Documentation:

Ethernet Routing Switch 8600	
Technical Configuration Guide for SMLT	http://www.nortel.com/support
Technical Configuration Guide for VRRP	http://www.nortel.com/support
Network Design Guidelines (per major release)	http://www.nortel.com/support
Configuring IP Routing Operations	http://www.nortel.com/support
Configuring VLANs, Spanning Tree, and Link Aggregation	http://www.nortel.com/support
Ethernet Routing Switch 8300	
Configuring IP Routing and Multicast Operations	http://www.nortel.com/support
Configuring VLANs, Spanning Tree, and Static Link Aggregation	http://www.nortel.com/support
Ethernet Routing Switch 1600	
Configuring IP Routing and Multicast Operations	http://www.nortel.com/support
Configuring VLANs, Spanning Tree, and Static Link Aggregation	http://www.nortel.com/support
Ethernet Switch 4500	
Configuring VLANs, Spanning Tree, and MultiLink Trunking	http://www.nortel.com/support
Configuring IP Multicast Routing Protocols	http://www.nortel.com/support

Table 5.1 – Nortel Reference Documentation

Microsoft TechNet	
Windows Server 2003 Clustering Services	http://technet2.microsoft.com

Table 5.2 – Microsoft Reference Documentation