

# Implementing Microsoft Network Load Balancing in a Virtualized Environment

VMware® Infrastructure 3

---

Network Load Balancing is a feature of recent Microsoft server operating systems, including Windows 2000 Advanced Server, Windows Server 2003, and Windows Server 2008. This clustering technology enables you to improve the scalability and availability of Internet server programs, such as Web servers, proxy servers, DNS servers, FTP servers, virtual private network servers, streaming media servers, and terminal services servers. In addition, it can detect host failures and automatically redistribute traffic to servers that are still operating.

In a VMware® Infrastructure 3 environment, you can create a cluster for Network Load Balancing using virtual machines on the same host or virtual machines on multiple hosts.

This technical note provides instructions for implementing Network Load Balancing in multicast or unicast mode. It covers the following topics:

- [“Network Load Balancing Basics”](#) on page 1
- [“Planning a Network Load Balancing Cluster”](#) on page 1
- [“Configuring Network Load Balancing in Windows”](#) on page 2
- [“References”](#) on page 3

## Network Load Balancing Basics

Network Load Balancing is implemented in a special driver installed on each Windows host in a cluster. The cluster presents a single IP address to clients. When client requests arrive, they go to all hosts in the cluster, and an algorithm implemented in the driver maps each request to a particular host. The other hosts in the cluster drop the request. You can set load partitioning to distribute specified percentages of client connections to particular hosts. You also have the option of routing all requests from a particular client to the host that handled that client’s first request.

Hosts in the cluster exchange heartbeat messages so they can maintain consistent information about what hosts are members of the cluster. If a host fails, client requests are rebalanced across the remaining hosts, with each remaining host handling a percentage of requests proportional to the percentage you specified in the initial configuration.

## Planning a Network Load Balancing Cluster

Network Load Balancing relies on the fact that incoming packets are directed to all cluster hosts and passed to the Network Load Balancing driver for filtering.

You can configure a Network Load Balancing cluster in one of the following modes:

- **Multicast**—Multicast mode allows communication among hosts because it adds a Layer 2 multicast address to the cluster instead of changing the cluster. Communication among hosts is possible because

the hosts retain their original unique media access control (MAC) addresses and already have unique, dedicated IP addresses. However, the address resolution protocol (ARP) reply that is sent by a host in the cluster (in response to an ARP request) maps the cluster's unicast IP address to its multicast MAC address.

Some routers do not support the resolution of unicast IP addresses to multicast MAC addresses, and they discard the ARP reply. As a result, an administrator must add a static ARP entry in the router, mapping the cluster IP address to its MAC address.

To configure a Network Load Balancing cluster in multicast mode correctly, you need to complete the steps in [“Configuring Network Load Balancing in Windows”](#) on page 2.

---

**NOTE** VMware recommends that you use multicast mode, because unicast mode forces the physical switches on the LAN to broadcast all Network Load Balancing traffic to every machine on the LAN.

---

- **Unicast**—Unicast mode works seamlessly with all routers and Layer 2 switches. However, this mode induces switch flooding, a condition in which all switch ports are flooded with Network Load Balancing traffic, even ports to which servers not involved in Network Load Balancing are attached. To communicate among hosts, you must have a second virtual adapter for each host.

Normally, switched environments avoid port flooding when a switch learns the MAC addresses of the hosts that are sending network traffic through it. The Network Load Balancing cluster masks the cluster's MAC address for all outgoing traffic to prevent the switch from learning the MAC address.

On an ESX host, the VMkernel sends a reverse address resolution protocol (RARP) packet each time certain actions occur—for example, when a virtual machine is powered on, when there is a teaming failover, or when certain VMotion operations occur. The RARP packet gives physical switches the MAC address of the virtual machine involved in the action. In a Network Load Balancing cluster environment, after a Network Load Balancing node is powered on, the notification in the RARP packet exposes the MAC address of the cluster NIC. As a result, switches might begin to send all inbound traffic destined for the Network Load Balancing cluster through one switch port to a single node of the cluster.

Because the virtual switch operates with complete data about the underlying MAC addresses of the virtual NICs inside each virtual machine, it always correctly forwards packets containing a MAC address matching that of a running virtual machine. As a result of this behavior, the virtual switch does not forward traffic destined for the Network Load Balancing MAC address outside the virtual environment into the physical network, because it is able to forward it to a local virtual machine.

To correctly configure a Network Load Balancing cluster in unicast mode, you need to complete the steps in [“Configuring Network Load Balancing in Windows”](#) on page 2 and [“Configuring Unicast Mode”](#) on page 3.

## Configuring Network Load Balancing in Windows

- 1 Install a Windows operating system that supports Network Load Balancing in your virtual machines.  
You should install two virtual NICs in each virtual machine that will be part of the Network Load Balancing cluster. One virtual NIC from each virtual machine is used for Network Load Balancing. The other virtual NIC is used for management of the Windows virtual machine.
- 2 Each Network Load Balancing node requires a static IP address to be assigned to the Network Load Balancing-bound adapter. You need one or more additional static IP addresses to be used as the virtual IP addresses of the Network Load Balancing cluster. Use IP addresses that belong to the same subnet.
- 3 Configure Network Load Balancing options using one of the following:
  - **Network Load Balancing Manager**
  - **Network Load Balancing Properties** dialog box accessed through **Network Connections**

VMware recommends that you use the **Network Load Balancing Manager**.

Using both **Network Load Balancing Manager** and **Network Connections** to change **Network Load Balancing** properties can lead to unpredictable results.

You can use **Network Load Balancing Manager** from inside a node or from a remote machine that can communicate with all nodes. By default, **Network Load Balancing Manager** is installed on Windows Server 2003, and you can access it by clicking **Settings > Control Panel > Administrative Tools > Network Load Balancing Manager**.

## Configuring Multicast Mode

You do not need to take any special steps to configure your ESX host when you are using multicast mode.

## Configuring Unicast Mode

This procedure helps prevent RARP packet transmission for the virtual switch as a whole. This setting affects all the port groups that use the switch. You can override this setting for individual port groups using the instructions in [“To Prevent RARP Packet Transmission for a Port Group”](#) on page 3.

### To Prevent RARP Packet Transmission for a Virtual Switch

- 1 Log on to the VI Client and select the ESX host.
- 2 Click the **Configuration** tab.
- 3 Choose **Networking** and, for the virtual switch, select **Properties**.
- 4 On the **Ports** tab, select the virtual switch and click **Edit**.
- 5 Click the **NIC Teaming** tab, set **Notify Switches** to **No**.
- 6 Click **OK** and close the **vSwitch Properties** dialog box.

Complete the following steps to prevent RARP packet transmission only for an individual port group. This setting overrides the setting you make for the virtual switch.

### To Prevent RARP Packet Transmission for a Port Group

- 1 Log on to the VI Client and select the ESX host.
- 2 Click the **Configuration** tab.
- 3 Choose **Networking** and, for the virtual switch, select **Properties**.
- 4 On the **Ports** tab, select the port group and click **Edit**.
- 5 Click the **NIC Teaming** tab, set **Notify Switches** to **No**.
- 6 Click **OK** and close the **vSwitch Properties** dialog box.

For more information, see the VMware knowledge base article “Microsoft NLB Not Working Properly in Unicast Mode.” For a link, see [“References”](#) on page 3.

## References

For additional information, see the following sources:

- “Checklist: Enabling and configuring Network Load Balancing”  
<http://go.microsoft.com/fwlink/?LinkId=18371>
- Clustering Services in Windows 2003  
<http://www.microsoft.com/windowsserver2003/technologies/clustering/default.msp>
- Microsoft Communities Web Site  
<http://go.microsoft.com/fwlink/?LinkId=18374>
- “Microsoft NLB Not Working Properly in Unicast Mode”  
<http://kb.vmware.com/kb/1556>

- “Reasons for Using Network Load Balancing”  
<http://technet2.microsoft.com/windowsserver/en/library/7698646d-510e-47f9-9b09-b31dec12be3a1033.mspx>
- “VMware Virtual Networking Concepts”  
<http://www.vmware.com/resources/techresources/997>
- “Windows Server 2003 Deployment Kit”  
<http://go.microsoft.com/fwlink/?LinkId=18370>

---

If you have comments about this documentation, submit your feedback to: [docfeedback@vmware.com](mailto:docfeedback@vmware.com)

**VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 [www.vmware.com](http://www.vmware.com)**

Copyright © 2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,944,699, 6,961,806, 6,961,941, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,269,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253, and 7,356,679; patents pending. VMware, the VMware “boxes” logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Revision 20080904 Item: EN-000012-00

---