

Network load balancing

Network load balancing is the ability to balance traffic across two or more [WAN](#) links without using complex [routing protocols](#) like [BGP](#).

This capability balances network sessions like Web, email, etc. over multiple connections in order to spread out the amount of [bandwidth](#) used by each [LAN](#) user, thus increasing the total amount of bandwidth available. For example, a user has a single WAN connection to the [Internet](#) operating at 1.5 Mbit/s. They wish to add a second broadband (cable, DSL, wireless, etc.) connection operating at 2.5 Mbit/s. This would provide them with a total of 4 Mbit/s of bandwidth when balancing sessions.

Session balancing does just that, it balances sessions across each WAN link. When Web browsers connect to the Internet, they commonly open multiple sessions, one for the text, another for an image, another for some other image, etc. These sessions can be balanced across the available connections. An [FTP](#) application only uses a single session so it is not balanced; however if a secondary FTP connection is made, then it may be balanced so that the traffic is distributed across two of the various connections and thus provides an overall increase in throughput.

Additionally, network load balancing is commonly used to provide network [redundancy](#) so that in the event of a WAN link outage, access to network resources is still available via the secondary

link(s). Redundancy is a key requirement for [business continuity](#) plans and generally used in conjunction with critical applications like [VPNs](#) and [VoIP](#).

Finally, most network load balancing systems also incorporate the ability to balance both outbound and inbound traffic. Inbound load balancing is generally performed via [dynamic DNS](#) which can either be built into the system, or provided by an external service or system. Having the dynamic DNS service within the system is generally thought to be better from a cost savings and overall control point of view.

Microsoft NLB

Microsoft has also purchased^[1] a technology that it renamed Network Load Balancing (NLB) that allows for efficient utilization of multiple network cards. MS NLB can be configured in [unicast](#) or in [multicast](#) mode^[2] where in multicast mode you can enable [IGMP snooping](#).

MS NLB was introduced for the first time in [Windows NT server](#) to spread traffic over multiple hosts without the need for a hardware based load balancer, e.g. when you host a busy web-server application where a single host wouldn't be able to manage all the traffic. And in more recent applications it would be used in Windows clusters for [Hyper-V](#) or [Microsoft SQL Server](#)

Unicast mode

In [unicast](#) mode MS NLB reassigns the stations [MAC address](#) (which applies to the clusters [IP address](#)) to a virtual MAC address and all [NIC's](#) in the NLB cluster use this same MAC address. This setup will cause all incoming traffic for the cluster to be flooded to all ports of the switch as unknown unicast frames.^[3] even to hosts that are not joining in the cluster. To keep flooding minimal you would need to use a dedicated [VLAN](#) for the cluster.

Multicast mode

Another option is to make NLB in [multicast](#) mode. The unicast [IPv4](#) address of the cluster is linked to a [multicast](#) MAC address. The hosts in the cluster will never send traffic to the [switch](#) using this MAC address with the cluster IPv4 address so one would need to create a static [ARP](#) entry on the router (layer 3) in the attached network. Not all vendors will allow you to create an ARP entry where you use a unicast IP address and a multicast MAC address. Cisco publishes some examples how to set up MS NLB on Catalyst switches running [IOS](#)^[4] and these same examples can be used for switches from many other vendors. As with NLB in unicast mode:

incoming traffic towards the cluster will be flooded to all ports in the switch/VLAN and not all vendors support this setup. To limit the flooding, MS NLB now supports IGMP which should lead to the switches learning which ports are actually using the [multicast address](#),^[5] but it doesn't always lead to the desired result. For example, [Dell PowerConnect](#) multi-layer switches officially don't support MS NLB in multicast. Even though it does work, it will lead to high CPU utilization - affecting (other) traffic in the switch^[6] and on other switches one might have other limitations such that the switch to which the NLB NICs are connected can't be the same switch that does the IP routing.

Server load balancing

When multiple servers are joined to create a [cluster](#). Clusters can use network load balancing whereby simultaneous cluster request are distributed between cluster servers.

[Round-robin DNS](#) records is one form of cluster load balancing. It works by creating multiple host records (usually A and/or AAAA) for one machine. As clients make requests, DNS rotates through its list of records.

In addition to the before mentioned, to configure a terminal server cluster, one needs a load-balancing technology such as Network Load Balancing (NLB) or DNS round robin. A load balancing solution will distribute client connections to each of the terminal servers.

Terminal Server Session Directory is a feature that allows users to easily and automatically reconnect to a disconnected session in a load balanced [Terminal Server](#) farm. The session directory keeps a list of sessions indexed by username and server name. This enables a user, after disconnecting a session, to reconnect to the correct Terminal Server where the disconnected session resides in order to resume working in that session. This reconnection will work even if the user connects from a different client computer.

See also

- [Network Load Balancing Services](#)
- [Load balancing \(computing\)](#)
- [Link aggregation](#)

References

1. "*Microsoft purchases clustering company | Windows Server content from Windows IT Pro*" (<https://web.archive.org/web/20170305215305/http://windowsitpro.com/windows-server/microsoft-purchases-clustering-company>) . Archived from the original (<http://windowsitpro.com/windows-server/microsoft-purchases-clustering-company>) on 2017-03-05. Retrieved 2015-10-18.
2. TechNet *Overview of NLB* (<https://technet.microsoft.com/en-us/library/cc725691.aspx>) , visited 8 June 2013
3. VMware knowledge base: *MS NLB in unicast and multicast mode* (http://kb.vmware.com/self-service/microsites/search.do?language=en_US&cmd=displayKC&externalId=1006580) , visited 8 May 2013
4. Cisco: *Catalyst configuration examples* (http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_example09186a0080a07203.shtml) Archived (https://web.archive.org/web/20140131115541/http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_example09186a0080a07203.shtml) 2014-01-31 at the Wayback Machine, visited 8 June 2013
5. Microsoft Support knowledge base: *KB283028: IGMP Snooping support on NLB* (<http://support.microsoft.com/kb/283028>) , visited 8 June 2013
6. Dell Tech Communities:*Trying to limit switch flooding from a Microsoft Windows Network Load Balancing (NLB)* (<http://en.community.dell.com/support-forums/network-switches/f/866/p/19435793/20091643.aspx>) , February 2012. Visited: 8 June 2013

Retrieved from

["https://en.wikipedia.org/w/index.php?](https://en.wikipedia.org/w/index.php?title=Network_load_balancing&oldid=1110978534)

[title=Network_load_balancing&oldid=1110978534"](https://en.wikipedia.org/w/index.php?title=Network_load_balancing&oldid=1110978534)

Last edited 3 months ago by Polyamorph

WIKIPEDIA
