

Overview: Load Balancing with the MNLB Feature Set for LocalDirector

This chapter provides a conceptual overview of load balancing and introduces Cisco's MultiNode Load Balancing (MNLB) Feature Set for LocalDirector, an IP server load balancing solutions for high-end Web and IBM Parallel Sysplex sites. Terms and acronyms are defined in a Glossary at the end of this chapter.

Challenges in Networking Today

The explosive growth of internet business, along with the proliferation of new IP-based enterprise resource planning applications, is creating a heightened requirement for continuous availability of mission-critical data residing on UNIX, NT, or IBM servers. Enterprises are implementing internet business strategies to increase revenue, raise customer satisfaction levels, streamline their supply chain management, and optimize their work force. Consequently, enterprises are experiencing ever-increasing demands from customers, suppliers, and employees for access to applications and data. These challenges are being met by a number of networking strategies and services.

The Role of Server Clusters

A server cluster is a set of computer systems connected together through multisystem hardware or software to provide services that were traditionally provided by a single system. Clustering of servers has been in practice for many years and has been used to achieve the following benefits:

Scalability—Traditional multiprocessor systems have relatively low scalability limits due to contention for shared memory and I/O. Clustered machines can scale farther in that the workload for any particular user is bound to a particular machine and far less sharing is needed.

Non-disruptive growth—When workloads grow beyond the capacity of a single machine, the traditional approach is to replace it with a larger machine or, if possible, add additional processors within the machine. In either case, this requires downtime for the entire machine. With clustering, machines can be added to the cluster without disrupting work that is executing on the other machines. When the new machine comes on-line, the work can start to migrate to that machine, thus reducing the load on the existing machines.

Load Balancing—Spreading users across multiple independent systems can result in wasted capacity on some systems while others are overloaded. By employing load balancing within a cluster of systems the users are spread to available systems based on the load on each system.

Continuous Availability—Individual application instances or machines can fail (or be taken down for maintenance) without shutting down service to end-users. Users on the failed system reconnect and should not be aware that they are using an alternate image. Users on the other systems are virtually unaffected except for the additional load caused by services for some portion of the users

that were formerly on the failed system. Some high-availability systems go a step further by moving the users transparently from the failed system to the backup system but that requirement diminishes in importance if session reconnect time is small.

Content Aware Networking

Content aware networking, as the term suggests, is a networking strategy that enables content to be dynamically distributed. Because content can be dynamically cached, it can be located at any given place at any given time and distributed between the servers and the location of the Web cache. Cisco has developed the ContentFlow architecture and the Dynamic Feedback Protocol (DFP) to enable networks to provide content aware networking services

Cisco's Family of Load Balancing Products

The MNLB Feature Set for LocalDirector is the newest of Cisco's family of load balancing products.

Distributed Director delivers geographical load balancing on network topology and traffic patterns.

LocalDirector performs IP server load balancing for small to medium-sized web sites.

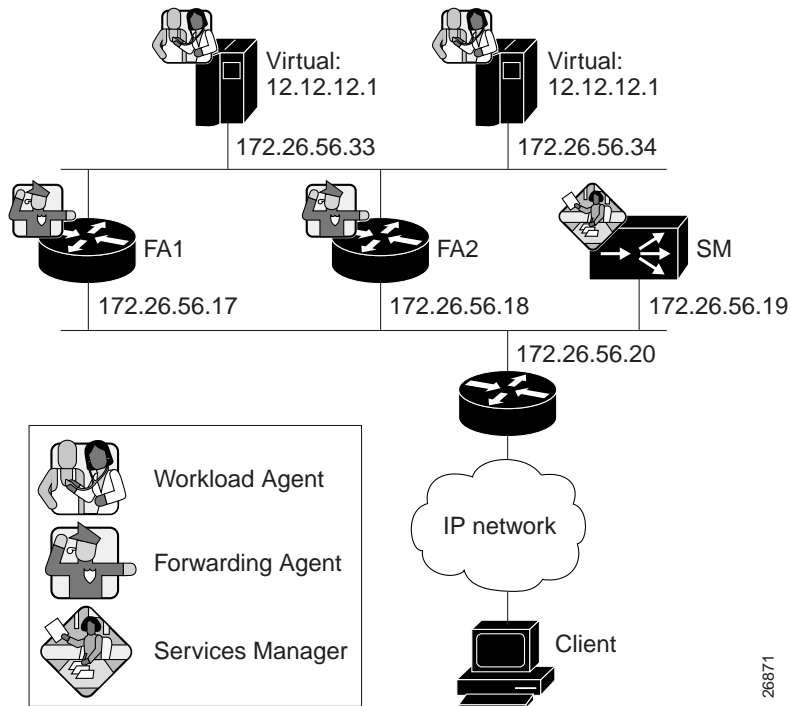
MNLB Feature Set for LocalDirector fits the needs of the large enterprise web site and IBM Parallel Sysplex data center, which require high availability and scalability for new e-Commerce and internet business applications. The MNLB Feature Set for LocalDirector comprises software running on Cisco routers and switches, Cisco's LocalDirector, and application server platforms.

MNLB Feature Set for LocalDirector

The MNLB Feature Set for LocalDirector distributes load balancing capabilities across any number of devices, enabling the highest levels of availability, scalability, and performance for server applications. The MNLB Feature Set for LocalDirector blasts the conventional restriction that all traffic inbound for a server cluster must pass through a single load-balancing engine. It enables the customization of routing decisions in neighboring IP routing engines, called forwarding agents, under the direction of a services manager. The combination of fast forwarding agents (whether routers or IP-aware switches) and load-balancing service managers (to provide synchronization and control) eliminates the scalability limitations of the past. And MNLB Feature Set for LocalDirector employs an IP-based feedback mechanism that enables the continuous adjustment of load-balancing decisions.

Figure 1 shows the functional components of the MNLB Feature Set for LocalDirector.

Figure 1-1 Components of the MNLB Feature Set for LocalDirector



26871

MNLB Services Manager




The MNLB Services Manager makes the load-balancing decisions based on application availability, server capacity, and load distribution algorithms such as round robin or least connections, or the Dynamic Feedback Protocol (DFP).

The Services Manager runs on a Cisco LocalDirector running LocalDirector version 3.1.1 or higher.

In this example, the Services Manager connects to the lower LAN. The Services Manager could also connect to the upper LAN, but not to the upper and lower LANs at the same time.

MNLB Forwarding Agent



The MNLB Forwarding Agent is an IOS-based packet redirector that forwards packets based on instructions received from the Services Manager. Each Forwarding Agent “learns” the destination of specific connection requests and forwards packets between the appropriate client and chosen destination. When a Forwarding Agent receives a connection request, the request is forwarded to the Services Manager. The Services Manager makes the load balancing decision and instructs the Forwarding Agents with the optimal destination. After destination selection, session data is forwarded directly to the destination without further Services Manager participation. There is no limit to the number of Forwarding Agents that can be configured in the MNLB Feature Set for LocalDirector.

The Forwarding Agent runs on an IOS router or RSM.

Note If you are running CASA, you must configure **ip pim dense**. Some Forwarding Agents might be many hops away, so TTL=1 might not work in some cases.

MNLB Workload Agent



Workload Agents are value-added software components developed for specific platforms by third-party developers. Workload Agents run on server platforms or on platforms that manage server farms. Workload Agents deliver server and application information to the Services Manager. This information enables the Services Manager to make optimum server selection.

Supported Platforms

The Cisco MNLB Feature Set for LocalDirector supports the hardware and software platforms listed in Table 1-1

Table 1-1 Cisco MultiNode LocalDirector Hardware and Software Platforms

Component	Hardware Platform	Software Version
Services Manager	LocalDirector	LocalDirector Version 3.1.1 and higher
Forwarding Agent	Cisco 75xx Cisco 72xx Cisco Cat5000/RSM Cisco 47xx Cisco 36xx	Cisco IOS Release 12.0(5)T and higher
	Cisco Cat6500 with MSFC Cisco 76xx	Cisco IOS Release 12.1(1)E and higher

Supported MIBs and RFCs

The MNLB Forwarding Agent supports the following MIB:

- cisco-casa-fa-mib.my

For descriptions of supported MIBs and how to use MIBs, see Cisco's MIB website on CCO at <http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>.

No RFCs are supported by this feature.

Glossary

affinity

The set of information that uniquely describes the association of a client to a particular host within a server cluster. It relates the addresses found in the IP packet (protocol, client IP address, port, local port, cluster address) to the IP address of the local host that has been assigned to handle all packets with that unique set of addresses.

CEF

Cisco Express forwarding.

cluster

A set of computer systems that are connected together through multisystem hardware or software to provide services traditionally provided by a single system. This arrangement provides higher availability and better scalability.

cluster address

The IP address that represents the entire cluster of hosts. It is defined in each server and each Forwarding Agent router. The routers advertise routes to their internal instances of this address. The servers do not advertise the address, but recognize it as one of their local, or loopback, addresses.

ContentFlow architecture

Cisco's content aware networking architecture that describes message flows and actions in a distributed environment.

control address

An IP address assigned to the IP cluster function within each router. It is unique to each router and is used for management flows.

Forwarding Agent

Fulfills routing decisions made by the MNLB Services Manager. The Forwarding Agent filters packets coming into the virtual network and sends to the services manager packets that are without a known server destination.

IGMP

The MultiNode LocalDirector Forwarding Agent uses Internet Group Management Protocol (IGMP) multicast to listen to the MultiNode LocalDirector Services Manager broadcasts. IP hosts use IGMP to report their group membership to directly-connected multicast routers. IGMP uses group addresses, which are Class D IP addresses. The high-order four bits of a Class D address are 1110. This means that host group addresses can be in the range 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is guaranteed not to be assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

load balancing

Spreading user requests among available servers within a cluster of servers, based on a variety of algorithms.

LocalDirector

Cisco LocalDirector is hardware and software that intelligently load balances IP traffic across multiple servers.

MD5

Message Digest Algorithm Version 5, a neighbor router authentication scheme used to ensure reliability and security when routing updates are to be exchanged between neighbor routers.

MVS Sysplex

One or more MVS systems connected and sharing resource information.

MVS Workload Manager

MVS component consisting of software and hardware and sharing resource information for each MVS system.

NetFlow switching

High-performance network-layer switching path that captures as part of its switching function a rich set of traffic statistics including user, protocol, port, and type of service information.

OS/390 Workload Agent

The software that supervises registration and deregistration of TCP/IP stacks and applications. Sends information to the Services Manager for workload balancing across the sysplex.

server farm

Also called a server cluster, a group of real servers that provide various applications and services.

Services Manager

Using load balancing and server/application feedback, determines a real server for the packet flow. Once the optimal destination is decided, all other packets in the packet flow are directed to a Forwarding Agent and real server, increasing packet throughput. Within the MNLB Feature Set for LocalDirector, the Services Manager function is performed by the LocalDirector.

virtual server

Presents a single address that represents an application server farm for clients.

workload manager

An MVS function that tracks the use of the MMS.