



NetWare Link Services Protocol

Building a Link State Database

Laura Chappell

With IntranetWare and NetWare 4.11, Novell changed its default routing protocol from Routing Information Protocol (RIP) to NetWare Link Services Protocol (NLSP). NLSP offers the following advantages over RIP:

- ? Load-balancing capabilities
- ? Improved manageability
- ? Improved reliability
- ? Faster convergence
- ? Less network overhead
- ? Better internetwork support (since NLSP can traverse more hops than RIP can)

This article is the first part of a three-part series that presents a packet-level view of NLSP communications. This article explains how NLSP devices perform the following tasks:

- ? Discover other NLSP devices on the network
- ? Create a link state database and broadcast the database throughout the network
- ? Make routing decisions

The second article in this series will examine how NLSP devices maintain the link state database and revise and purge invalid entries. Finally, the third article in this series will focus on how NLSP devices interoperate with RIP devices. The third article will also explain how NLSP devices receive and exchange RIP information and how these devices make routing decisions in a mixed NLSP and RIP environment. (An NLSP device can be a router, an IntranetWare server, or a NetWare server.)

DISTANCE VECTOR ROUTING VERSUS LINK STATE ROUTING

RIP is a *distance vector* routing protocol. A distance vector router knows only the next router in the data path or the address of the destination device (if the destination device is on a directly attached network).

By contrast, NLSP is a *link state* routing protocol that is similar to Open Shortest Path First (OSPF), the link state routing protocol used in a TCP/IP environment. A link state router maintains a map of the entire network and knows the route from one network to another. In addition, if a link state router is congested, it can assume a less important role in exchanging link information.

FRIENDLY NEIGHBORS

When an NLSP device is booted, it participates in a process called *neighbor notification*. The NLSP device transmits a hello packet onto its local network, notifying other network devices that the NLSP device is on the network. (All NLSP communications use socket number 0x9001.) The NLSP device then waits to receive a hello packet from other NLSP devices on the same network. If an NLSP device receives a hello packet, this device knows that it has a neighboring NLSP device.



For example, suppose that you attached the MPR1 router to a network on which the CORP-FS router was already running. (See [Figure 1](#).) The MPR1 router would initiate the neighbor notification process by broadcasting a hello packet. Each hello packet includes a Neighboring Routers field. Since the MPR1 router would not yet know if it had any neighboring NLSP devices, this field would be empty.

When the CORP-FS router received the hello packet, CORP-FS would record the MPR1 router in an adjacency database. (See [Figure 2](#).) The CORP-FS router would also note that MPR1 was in the initializing state. Each NLSP device maintains its own *adjacency database*, which contains information about all neighboring NLSP devices on the local network. The next hello packet the CORP-FS router sent would include the MPR1 router's network interface card (NIC) address. (See [Figure 3](#).)

After receiving this packet, the MPR1 router would record the CORP-FS router in an adjacency database. The MPR1 router would also note that the CORP-FS router was in the up state because CORP-FS had processed information about MPR1. The next hello packet the MPR1 router sent would include the CORP-FS router's NIC address in the packet's Neighboring Routers field.

If you attached additional NLSP devices to the same network, these devices would complete the same neighbor notification process. The MPR1 router and the CORP-FS router would then add the new devices to their adjacency database. Both routers would also list the new devices' NIC address in subsequent hello packets.

The number of neighboring NLSP devices an NLSP device can list in a hello packet depends on the maximum packet size used on the network. Listing each neighboring NLSP device requires 6 bytes for the NIC address, plus 2 bytes to define the field name and length. Thus, a 1,518-byte Ethernet packet could list approximately 189 neighboring NLSP devices.

Identifying the Designated Router

During the neighbor notification process, NLSP devices must determine which device has the highest priority. The NLSP device with the highest priority is the network's *designated router*, which helps each device maintain a map of the entire network. (The role of the designated router is explained in more depth later in this article.)

The first NLSP device that is booted on the network has priority number 44. If the first NLSP device transmits two hello packets but does not receive any hello packets, this device assumes that it is the designated router. The first NLSP device then increases its priority number by 20, resulting in priority number 64. (See [Figure 4](#).)

If another NLSP device is booted on the same network, this device broadcasts a hello packet, using the default priority number 44. As a result, the second NLSP device is not the designated router.

If two NLSP devices have the same priority number, they use their NIC address as a tiebreaker. The NLSP device with the highest NIC address becomes the designated router.

If you want to ensure that a particular NLSP device is always the designated router, you can use the INETCFG utility to specify a high priority number (such as 100) for this device.

The designated router requires additional memory and processing power. If a server is already overloaded, you might want to assign the server a low priority number, thereby ensuring that this server does not become the designated router.



In some cases, an NLSP device will be the designated router on one network but will not be the designated router on another network. In [Figure 2](#), for example, the MPR1 router is the designated router for network 0xAA-BB-CC-DD, and the CORP-FS router is the designated router for network 0xCC-DD-EE-FF.

The designated router transmits a hello packet approximately every 10 seconds. Other NLSP routers transmit a hello packet approximately every 20 seconds.

Hello, WAN Neighbor

If NLSP devices are connected by a WAN link (such as an X.25 or T1 link), the neighbor notification process differs slightly. Before these NLSP devices can exchange hello packets, they must initialize the WAN link and use the IPXWAN protocol to define operational characteristics such as WAN link throughput and delay.

During the neighbor notification process, each NLSP device on the WAN link sends a circuit ID number to other NLSP devices on the WAN link. The circuit ID number is assigned to the NLSP device's circuit when this circuit is created. (A *circuit* is a logical connection between devices.)

NLSP devices on the WAN link also exchange their state information (up, down, or initializing) and the maximum packet size (excluding data-link header) that their circuit supports.

OUTSIDE THE NEIGHBORHOOD

After an NLSP device knows its neighboring NLSP devices, it participates in a process called *Link State Protocol (LSP) information exchange*. The NLSP device transmits LSP packets, which contain information such as the device's links, services, and external routers.

Other NLSP devices on the network use these LSP packets to create a *link state database*, which is a map of the entire network. The sending NLSP device, in turn, uses LSP packets transmitted by other NLSP devices on the network to create its own link state database.

To help all of the NLSP devices synchronize their link state database, the designated router transmits Complete Sequence Number Packets (CSNP), which summarize the information in the link state database. Using this summary, each NLSP device checks its copy of the link state database, ensuring that this copy is up-to-date and accurate. In this way, each NLSP device has an identical copy of the link state database and can use this database to route packets to their destination.

NLSP devices on a local network do not acknowledge the receipt of LSP packets. When an NLSP device receives an LSP packet from another device on the same network, the device checks its link state database to see if it has more current information than the incoming LSP packet. When an NLSP device receives an LSP packet over a WAN link, however, this device replies with a Partial Sequence Number Packet (PSNP), acknowledging the receipt of the LSP packet.

You can view the entries in the network's link state database by typing `LOAD IPXCON` at the console of the NLSP device. If you have a third-party NLSP device, check the documentation to find out how you can view the device's copy of the link state database.

THE VERY REAL PSEUDONODE



Because each NLSP device tracks neighboring NLSP devices in an adjacency database, large networks with many NLSP devices could potentially have high overhead. To eliminate this overhead, NLSP is designed to use a *pseudonode*, which is a fictitious device that represents the entire network. (See [Figure 5](#).) Because the pseudonode is not a physical device, the designated router sends LSP packets on behalf of the pseudonode. These LSP packets announce the network and all of its routing devices.

The pseudonode LSP packet is the most important LSP packet transmitted on the network: An NLSP device could build the link state database using only the pseudonode LSP packet.

NLSP uses two types of pseudonodes: *LAN pseudonodes* and *WAN pseudonodes*. NLSP defines a LAN pseudonode as the network to which neighboring NLSP devices are attached. On the other hand, NLSP defines a WAN pseudonode as a *virtual circuit*. Whereas LAN links can be considered simple point-to-point connections between devices, WANs can include multiple devices connected by multiple virtual circuits. For example, an X.25 WAN link might contain five virtual circuits between the corporate office and a single branch office. In this case, NLSP considers each virtual circuit as a different point-to-point connection and assigns a pseudonode ID number to each virtual circuit.

TRANSMITTING LSP PACKETS

Now that you understand the designated router, the link state database, LSP packets, and pseudonodes, the following example will help you understand how the LSP information exchange works:

The CORP-FS Router

Suppose that the CORP-FS router had just completed the neighbor notification process. This router would then transmit its first LSP packet. (See [Figure 6](#).) The first LSP packet is assigned LSP ID number 0x02-00-00-22-22-22-00-00 and contains information about the local server and the NetWare Core Protocol (NCP) services located at socket number 0x0451. As [Figure 6](#) shows, the network does not contain neighboring NLSP devices, and the LSP packet was transmitted to the NLSP socket, which is socket number 0x9001.

If the CORP-FS router did not receive a hello packet or an LSP packet from another NLSP device, CORP-FS would transmit a second LSP packet. In this packet, the CORP-FS router would increase the value of the Sequence Number field. For example, the first LSP packet was assigned sequence number 2. (See [Figure 6](#).) The second LSP packet would then be assigned sequence number 3. The higher the sequence number is, the more up-to-date the LSP packet is.

The CORP-FS router would then transmit a third LSP packet (sequence number 4), which would include the name of the NDS tree in which CORP-FS resides. This LSP packet would also include the following information:

- ? The CORP-FS router is a file server.
- ? The CORP-FS router offers NCP services.
- ? The CORP-FS router's network ID number is 0x02-00-00-22-22-22-01.
- ? The network is a 10 Mbit/s Ethernet network with 10 million bits/second throughput.
- ? The maximum packet size (excluding the Ethernet header) is 1,497 bytes.
- ? The delay to transmit 1 byte of data to a destination on the network is approximately 200 microseconds.



- ? The metric cost of the route is 20. (This value is assigned by default. The *metric cost* is based on LAN throughput values or WAN-tested values. The higher the throughput is, the lower the metric cost is.)

Because the CORP-FS router transmitted this LSP packet, all listening devices would compare the packet with the information they already had about the CORP-FS router. If the LSP packet was more up-to-date, the listening devices would override their existing information for this router with the information contained in the packet. (The second article in this series will focus on how NLSP devices update their link state database.)

The CORP-FS router would then transmit another LSP packet, this time on behalf of the pseudonode on its network. This LSP packet would be assigned LSP ID number 0x02-00-00-22-22-22-01-00 and would indicate that the network address is 0xCC-DD-EE-FF.

Since metric costs are associated with an NLSP device's connection to the network, the pseudonode LSP packet would not include cost information for crossing the network. Metric costs are assigned only in LSP packets that report network routes.

The information contained in the transmitted LSP packets is entered into the link state database of all receiving devices. The CORP-FS router would use the link state database to create a map of the connection between itself and the pseudonode. (See [Figure 7](#).)

The MPR1 Router

Now suppose that you attached the MPR1 router to the network and that this router transmitted its first LSP packet to announce its services. In this LSP packet, MPR1 would provide information about its connection to network 0x02-00-00-22-22-22-01, including the following:

- ? The MPR1 router is a file server.
- ? The MPR1 router offers NCP services.
- ? The MPR1 router's internal network number is 0x00-11-11-11.

The MPR1 router would then transmit a second LSP packet (sequence number 3). The second LSP packet would contain information about the other network to which the MPR1 router was attached. This network has a network address of 0xAA-BB-CC-DD.

Next, the MPR1 router would transmit a third LSP packet (sequence number 4), which would contain information about TREE1, the NDS tree in which MPR1 resides. This LSP packet would also contain information about the MPR1 router's other neighboring NLSP device, which has an LSP neighbor ID number of 0x02-00-00-11-11-11-01. In addition, the LSP packet would contain the following:

- ? The network is a 10 Mbit/s Ethernet network with 10 million bits/second throughput.
- ? The metric cost of the route is 20.

Since the MPR1 router is the designated router on network 0xAA-BB-CC-DD, MPR1 would then send a fourth LSP packet, this time on behalf of the network's pseudonode.



The CORP-FS Router . . . Again

Finally, the CORP-FS router would transmit its fifth LSP packet, also on behalf of the pseudonode on its network, which has a network address of 0xCC-DD-EE-FF. The CORP-FS router would transmit this LSP packet because a new neighboring NLSP device, MPR1, had been attached to the pseudonode. The LSP packet would indicate that two neighboring NLSP devices are now attached to the pseudonode.

You could use the information contained in the LSP packets transmitted by the CORP-FS and MPR1 routers to create a conceptual diagram of the network. You could also use this information to show the metric cost of transmitting data in each direction. (See [Figure 8](#).)

THE OVERLOADED DESIGNATED ROUTER

Because the designated router transmits LSP packets on behalf of the pseudonode, this router must allocate a lot of memory and other resources for these tasks. If the designated router depletes its resources and cannot perform its duties, this router will do the following:

- ? **Enter LSP Database Overload State.** If the designated router cannot store an LSP packet, this router discards the packet. If the designated router discards an LSP packet, the link state database may be incomplete.
- ? **Lower Its Priority Number.** If the designated router lowers its priority number, this router will not be elected as the designated router for any other network.

BROADCAST OR MULTICAST

By default, an NLSP device uses a broadcast address to transmit LSP packets. However, you can use the INETCFG utility to specify that a particular NLSP device should use a multicast address to send LSP packets.

With a broadcast address, an NLSP device transmits LSP packets to all devices that are attached to the network. With a multicast address, on the other hand, an NLSP device transmits LSP packets only to NLSP devices.

If you enable multicast addressing for NLSP devices, each NLSP device will use the following multicast addresses:

IEEE 802.3 0x09-00-1B-FF-FF-FF

IEEE 802.5 0xC0-00-10-00-00-00

FDDI 0x09-00-1B-FF-FF-FF

MAKING ROUTING DECISIONS

Although NLSP devices use the link state database and the adjacency database to create a map of the entire network, these databases are not ideally suited to help the devices make routing decisions. Instead, an NLSP device takes the relevant information from the link state database to build a *forwarding database*. The NLSP device then consults the forwarding database for all routing processes.



The forwarding database is based on a decision-making process called Dijkstra's algorithm, which was named after Edsger W. Dijkstra, the man who defined the basic methodology. Dijkstra's algorithm treats each NLSP device and pseudonode as a node and each link as a line connecting two nodes. (See [Figure 9](#).) Each link has an associated metric cost.

Dijkstra's algorithm calculates the total metric cost of each route and then uses these costs to rank each route in priority. When a change is made to the link state database, each NLSP device must perform this decision-making process to update the forwarding database.

What if the network expands and NLSP devices have multiple routes from one network to another? For example, suppose that the network in [Figure 9](#) were expanded to include a second route from the MPR1 router (C) to network number 0x00-00-00-99 (M). (See [Figure 10](#).)

The MPR1 router could now use two paths to transmit data to M: One route is through B, E, F, G, H, and N and has a metric cost of 74 (20+14+20+20). The second route is through A, D, I, J, K, and L and has a metric cost of 80 (20+20+20+ 20). Because Dijkstra's algorithm uses the metric cost to identify the best path, the MPR1 router would transmit data to M through B, E, F, G, H, and N.

Load Splitting

If two routes have equal metric costs, NLSP chooses one route to use for all network communications; NLSP uses only the first route learned. However, you can use the INETCFG utility to configure NLSP to divide network communications among two or more routes. This process is called *load splitting*.

However, you should not manually assign metric costs to links that are unequal in throughput (such as 4 Mbit/s Token Ring and 10 Mbit/s Ethernet) to force load splitting. You might create problems with network implementations that are order-dependent, such as networks that use Novell's burst mode technology.

Asymmetrical Routing

You can also change the metric cost of a link to indicate that crossing a link in one direction costs more than crossing a link in another direction. Called *asymmetrical routing*, this process can create two separate routes--one for each direction.

Before you implement asymmetrical routing, however, you should know this process can make troubleshooting network communications difficult. Ideally, you want to see both the request and the corresponding reply on the same path when you analyze network communications.

CONCLUSION

This article explained how NLSP devices build the adjacency database, the link state database, and the forwarding database to create a map of the entire network and to discover the best route to particular destinations. The next article in this series will explain how NLSP devices maintain the link state database and revise and purge invalid entries.

NetWare Connection, September 1997, pp.38-45