

NetWare Link-Services Protocol (NLSP)

Background

The *NetWare Link-Services Protocol (NLSP)* is a link-state routing protocol from Novell designed to overcome some of the limitations associated with the IPX Routing Information Protocol (RIP) and its companion protocol, the Service Advertisement Protocol (SAP). NLSP is based on the OSI Intermediate System-to-Intermediate System (IS-IS) protocol and was designed to replace RIP and SAP, Novell's original routing protocols that were designed when internetworks were local and relatively small. As such, RIP and SAP are not well-suited for today's large, global internetworks. This chapter summarizes the routing processes and protocol components of NLSP.

As compared to RIP and SAP, NLSP provides improved routing, better efficiency, and scalability. In addition, NLSP-based routers are backward-compatible with RIP-based routers. NLSP-based routers use a reliable delivery protocol, so delivery is guaranteed. Furthermore, NLSP facilitates improved routing decisions because NLSP-based routers store a complete map of the network, not just next-hop information such as RIP-based routers use. Routing information is transmitted only when the topology has changed, not every 60 seconds as RIP-based routers do, regardless of whether the topology has changed. Additionally, NLSP-based routers send service-information updates only when services change, not every 60 seconds as SAP does.

NLSP is efficient in several ways. It is particularly useful over a WAN link because its support of IPX header compression makes it possible to reduce the size of packets. NLSP also supports multicast addressing so that routing information is sent only to other NLSP routers, not to all devices as RIP does.

In addition, NLSP supports load balancing across parallel paths and improves link integrity. It periodically checks links for connectivity and for the data integrity of routing information. If a link fails, NLSP switches to an alternate link and updates the network-topology databases stored in each node when connectivity changes occur anywhere in the routing area.

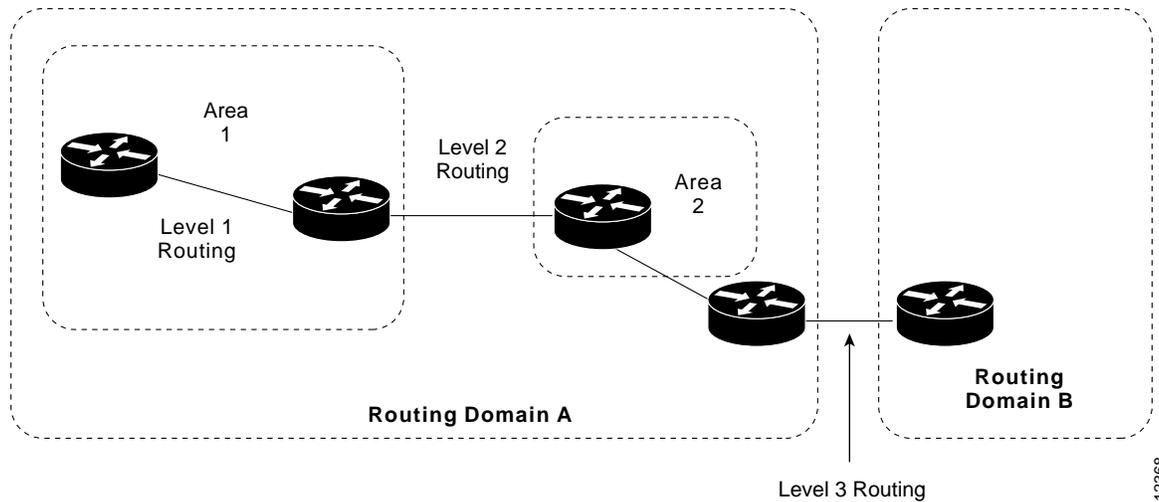
In terms of scalability, NLSP can support up to 127 hops (RIP supports only 15 hops) and permits hierarchical addressing of network nodes, which allows networks to contain thousands of LANs and servers.

NLSP Hierarchical Routing

NLSP supports hierarchical routing with area, domain, and global internetwork components. An area is a collection of connected networks that all have the same area address. A domain is a collection of areas that belong to the same organization. A global internetwork is a collection of domains that usually belong to different organizations but with an arms-length relationship. Areas can be linked to create routing domains, and domains can be linked to create a global internetwork.

NLSP supports three levels of hierarchical routing: *Level 1*, *Level 2*, and *Level 3* routing. A Level 1 router connects network segments within a given routing area. A Level 2 router connects areas and also acts as a Level 1 router within its own area. A Level 3 router connects domains and also acts as a Level 2 router within its own domain. Figure 40-1 illustrates the three routing levels NLSP defines.

Figure 40-1 NLSP defines three routing levels.



12368

Hierarchical Routing Efficiencies

Hierarchical routing simplifies the process of enlarging a network by reducing the amount of information that every router must store and process to route packets within a domain. A Level 1 router is required to keep detailed information only about its own area instead of storing link-state information for every router and network segment in its domain. To exchange traffic with other areas, a Level 1 router must only find the nearest Level 2 router. Between areas, Level 2 routers advertise the area address(es) only for their respective areas, not their entire link-state databases. Level 3 routers perform similarly between domains.

NLSP Adjacencies

By exchanging hello packets, a router determines the reachability of its neighbors and uses this information to establish adjacency. *Adjacency* is a record that a router keeps about the state of its connectivity with a neighbor and the attributes of the neighboring router. The router stores these records in its adjacency database.

Adjacency-establishment procedures vary depending upon whether the router is establishing and maintaining adjacencies over a WAN or a LAN.

Establishing router adjacency over a WAN involves first establishing the underlying data-link connection (details depend upon the medium). The routers then exchange identities by using the IPX WAN Version 2 protocol and determine certain operational characteristics of the link. Hello packets are exchanged and the routers update their adjacency databases. The routers then exchange both link-state packets (LSPs) describing the state of their links and IPX data packets over the link. To maintain a WAN link, the router maintains a state variable indicating whether the link is up, down, or initializing for each adjacency. If the router does not hear from a neighbor within the time specified in a holding timer, the router generates a message indicating that the link is down and deletes the adjacency.

WAN Hello packets enable routers to discover each other's identity, decide whether they are in the same routing area, and determine whether other routers and links are operational. A router sends Hello packets when the circuit is first established, when a timer expires, or when the contents of the next Hello to be transmitted are different than the contents of the previous Hello transmitted by this system (and one or more seconds have elapsed since the previous Hello). Hello packets are sent as long as the circuit exists.

Establishing a New WAN Adjacency

A typical startup procedure between two routers (A and B) on a WAN link begins with the link in the down state. Router A sends a WAN Hello indicating the down state to Router B, which changes its state for the link to Initializing. Router B sends a WAN Hello with a field indicating Initializing to Router A. Router A then changes its state for the link to initializing and sends a WAN Hello with a field indicating this to Router B. Router B changes its state for the link to the up state and sends a WAN Hello with a field indicating its new state. Finally, Router A changes its state for the link to Up.

Maintaining Adjacencies over LANs

When a broadcast circuit, such as an 802.3 Ethernet and 802.5 Token Ring, is enabled on a router, the router begins sending and accepting Hello packets from other routers on the LAN and starts the Designated Router election process.

The Designated Router represents the LAN as a whole in the link-state database, makes routing decisions on behalf of the whole, and originates LSPs on behalf of the LAN. This ensures that the size of the link-state databases that each router must construct and to manage to stay within reasonable limits.

Periodically, every router sends a multicast Hello packet on the LAN. The router with the highest priority (a configurable parameter) becomes the Level 1 Designated Router on the LAN. In case of a tie, the router with the higher MAC address wins.

Sending LAN Hello Packets

Hello packets enable routers on the broadcast circuit to discover the identity of the other Level 1 routers in the same routing area on that circuit. The packets are sent immediately when any circuit has been enabled to a special multicast destination address. Routers listen on this address for arriving Hello packets.

NLSP Operation

An NLSP router extracts certain information from the adjacency database and adds locally derived information. Using this information, the router constructs a link-state packet (LSP) that describes its immediate neighbors. All LSPs constructed by all routers in the routing area make up the link-state database for the area.

The NLSP specification intends that each router maintain a copy of the link-state database and keep these copies synchronized with each other. The link-state database is synchronized by reliably propagating LSPs throughout the routing area when a router observes a topology change. Two methods ensure that accurate topology-change information is propagated: *flooding* and *receipt confirmation*.

Flooding is instigated when a router detects a topology change. When such a change is detected, the router constructs a new LSP and transmits it to each of its neighbors. Such LSPs are directed packets on a WAN and multicast packets on a LAN. Upon receiving an LSP, the router uses the sequence

number in the packet to decide whether the packet is newer than the current copy stored in its database. If it is a newer LSP, the router retransmits it to all its neighbors (except on the circuit over which the LSP was received).

The receipt-confirmation process is different for LANs and WANs. On WANs, a router receiving an LSP replies with an acknowledgment. On LANs, no explicit acknowledgment occurs, but the Designated Router periodically multicasts a packet called a complete sequence number packet (CSNP) that contains all the LSP identifiers and sequence numbers it has in its database for the entire area. This ensures that other routers can detect whether they are out of synchronization with the Designated Router.

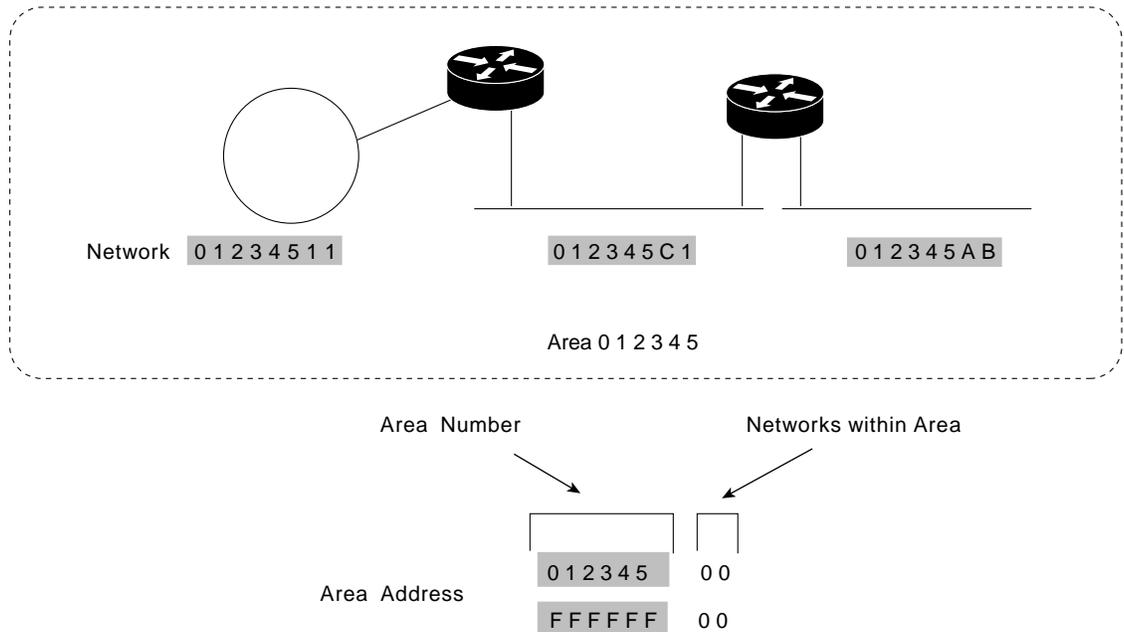
NLSP Hierarchical Addressing

NLSP supports a hierarchical addressing scheme. Each routing area is identified by two 32-bit quantities: a *network address* and a *mask*. This pair of numbers is called an *area address*. Expressed in hexadecimal, an example of an area address follows:

- 01234500—This number is the network address for this routing area. Every network number within that area starts with the identification code 012345.
- FFFFFFF0—This number is the mask that identifies how much of the network address refers to the area itself and how much refers to individual networks within the area.

In the example area address above, the first 24 bits (012345) identify the routing area. The remaining 8 bits are used to identify individual network numbers within the routing area (for example, 012345AB, 012345C1, 01234511). Figure 40-2 highlights the above addressing concepts with three different networks in a single area.

Figure 40-2 NLSP addresses consist of a network address and a mask.



A routing area can have as many as three different area addresses, each with a different mask. Having more than one area address allows the routing area to be reorganized without interrupting operations. Any combination of area addresses can be used within a domain.

NLSP Hello Packets

Two types of NLSP Hello packets exist: WAN Hello and Level 1 LAN Hello packets.

WAN Hello Packet

Figure 40-3 illustrates the fields of a WAN Hello packet.

Figure 40-3 Fourteen fields make up a WAN Hello packet.

WAN Hello			Number of Bytes
Protocol ID			1
Length Indicator			1
Minor Version			1
Reserved			1
Reserved	Packet Type		1
Major Version			1
Reserved			2
Reserved	State	Cct Type	1
Source ID			6
Holding Time			2
Packet Length			2
Local WAN Circuit ID			1
Variable Length Fields			Variable

inf3803

WAN Hello Packet Fields

The following WAN hello packet field descriptions summarize each field illustrated in Figure 40-3.

- *Protocol ID*—Identifies the NLSP routing layer with the 0x83 hex number.
- *Length Indicator*—Determines the number of bytes in the fixed portion of the header.
- *Minor Version*—Contains 1 possible decimal value and is ignored on receipt.
- *Reserved*—Contains no decimal values and is ignored on receipt.
- *Packet Type (5 bits)*—Contains 17 possible decimal values.
- *Major Version*—Contains 1 possible decimal value.
- *Reserved*—Contains no decimal values and is ignored on receipt.

- *State (2 bits)*—Sends the router's state associated with the link (0 = Up, 1 = Initializing, 2 = Down).
- *Circuit Type (Cct type)*—Consists of 2 bits. This field can have one of the following values:
 - 0 = Reserved value, ignore entire packet
 - 1 = Level 1 routing only
 - 2 = Level 2 routing only (sender uses this link for Level 2 routing)
 - 3 = Both Level 1 and Level 2 (sender is a Level 2 router and uses this link for Level 1 and Level 2 traffic)
- *Source ID*—Serves as the system identifier of the sending router.
- *Holding Time*—Contains the holding timer, in seconds, to be used for the sending router.
- *Packet Length*—Determines the entire length of the packet, in bytes, including the NLSP header.
- *Local WAN Circuit ID*—Acts as a unique identifier assigned to this circuit when it is created by the router.
- *Variable Length Field*—Consists of a series of optional fields.

NLSP LAN Hello Packets

Figure 40-4 illustrates the fields of a LAN Level 1 Hello packet.

Figure 40-4 A LAN Level 1 Hello packet consists of 16 fields.

LAN Level 1 Hello				Number of Bytes
Protocol ID				1
Length Indicator				1
Minor Version				1
Reserved				1
Reserved		Packet Type		1
Major Version				1
Reserved				2
Reserved	NM	Res	Cct Type	1
Source ID				6
Holding Time				2
Packet Length				2
R	Priority			1
LAN ID				7
Variable Length Fields				Variable

ifh3804

Level 1 LAN Hello Packet Fields

The following Level 1 LAN Hello packet field descriptions summarize each field illustrated in Figure 40-4:

- *Protocol ID*—Identifies the NLSP routing layer with the 0x83 hex number.
- *Length Indicator*— Determines the number of bytes in the fixed portion of the header (up to and including the LAN ID field).
- *Minor Version*—Contains 1 possible decimal value and is ignored on receipt.
- *Reserved*—Contains no possible decimal values and is ignored on receipt.
- *Packet Type (5 bits)*—Contains 15 possible decimal values.
- *Major Version*—Contains 1 possible decimal value.
- *Reserved*—Contains no possible decimal values and is ignored on receipt.
- *No Multicast (NM) (1 bit)*—Indicates, when set to 1, that the packet sender cannot receive traffic addressed to a multicast address. (Future packets on this LAN must be sent to the broadcast address.)

- *Circuit Type (Cct Type) (2 bits)*—Can have one of the following values:
 - 0 = Reserved value, ignore entire packet
 - 1 = Level 1 routing only
 - 2 = Level 2 routing only (sender uses this link for Level 2 routing)
 - 3 = Both Level 1 and Level 2 (sender is a Level 2 router and uses this link for Level 1 and Level 2 traffic)
- *Source ID*—Contains the system ID of the sending router.
- *Holding Time*—Contains the holding timer, in seconds, to be used for the sending router.
- *Packet Length*—Determines the entire length of the packet, in bytes, including the NLSP header.
- *R*—Contains no possible decimal values and is ignored on receipt.
- *Priority (7 bits)*—Serves as the priority associated with being the LAN Level 1 Designated Router. (Higher numbers have higher priority.)
- *LAN ID*—Contains the system ID (6 bytes) of the LAN Level 1 Designated Router, followed by a field assigned by that Designated Router.
- *Variable Length Fields*—Consists of a series of optional fields.