

Novell NetWare® 6

www.novell.com

INTERNETWORK PACKET EXCHANGE

N



Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1994-1995, 2001 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

U.S. Patent No. 5,157,663; 5,349,642; 5,455,932; 5,553,139; 5,553,143; 5,572,528; 5,594,863; 5,608,903; 5,633,931; 5,652,859; 5,671,414; 5,677,851; 5,692,129; 5,701,459; 5,717,912; 5,758,069; 5,758,344; 5,781,724; 5,781,724; 5,781,733; 5,784,560; 5,787,439; 5,818,936; 5,828,882; 5,832,274; 5,832,275; 5,832,483; 5,832,487; 5,850,565; 5,859,978; 5,870,561; 5,870,739; 5,873,079; 5,878,415; 5,878,434; 5,884,304; 5,893,116; 5,893,118; 5,903,650; 5,903,720; 5,905,860; 5,910,803; 5,913,025; 5,913,209; 5,915,253; 5,925,108; 5,933,503; 5,933,826; 5,946,002; 5,946,467; 5,950,198; 5,956,718; 5,956,745; 5,964,872; 5,974,474; 5,983,223; 5,983,234; 5,987,471; 5,991,771; 5,991,810; 6,002,398; 6,014,667; 6,015,132; 6,016,499; 6,029,247; 6,047,289; 6,052,724; 6,061,743; 6,065,017; 6,094,672; 6,098,090; 6,105,062; 6,105,132; 6,115,039; 6,119,122; 6,144,959; 6,151,688; 6,157,925; 6,167,393; 6,173,289; 6,192,365; 6,216,123; 6,219,652; 6,229,809. Patents Pending.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.

www.novell.com

Internetwork Packet Exchange
October 2001
103-000176-001

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

ConsoleOne is a trademark of Novell, Inc.

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

IPX/SPX is a trademark of Novell, Inc.

IPXWAN is a trademark of Novell, Inc.

MacIPX is a registered trademark of Novell, Inc., in the United States and other countries.

NetWare is a registered trademark of Novell, Inc., in the United States and other countries.

NetWare/IP is a trademark of Novell, Inc.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

NetWare Mobile IPX is a trademark of Novell, Inc.

NetWare MultiProtocol Router is a trademark of Novell, Inc.

Novell is a registered trademark of Novell, Inc., in the United States and other countries.

Novell Client is a trademark of Novell, Inc.

Novell Directory Services and NDS are registered trademarks of Novell, Inc., in the United States and other countries.

Packet Burst is a trademark of Novell, Inc.

Sequenced Packet Exchange and SPX are trademarks of Novell, Inc.

System Fault Tolerance, SFT, and SFT III are trademarks of Novell, Inc.

Third-Party Trademarks

All third-party trademarks are the property of their respective owners.

Contents

About This Guide	9
Documentation Conventions	9
1 Understanding	11
The IPX Protocol	12
IPX Packet Structure	13
IPX Addressing	17
Network Number	17
Node Number	19
Socket Number	20
How IPX Routing Works	21
When a Workstation Sends an IPX Packet.	21
When a Router Receives an IPX Packet	22
When a Router Forwards an IPX Packet.	23
IPX Operation over WAN Links	24
Call Types	25
Permanent Calls	25
On-Demand Calls.	25
Routed On-Demand Calls	26
Routing Types	27
Static Routes and Services	28
Watchdog Packet Spoofing	30
Header Compression.	31
Compression Slots	32
Compression Packet Types	33
IPX Route Aggregation.	34
Introducing Aggregated Routes into NLSP.	35
Consistent Use of Routers that Support Route Aggregation.	36
Interaction with SAP	36
Metrics Used with Aggregated Routes.	36
IPX Address Mapping Gateway	37
2 Planning	41
IPX Configuration Decisions	41
NetWare Mobile IPX Configuration Decisions	45
Mobile Client Driver Selection	45
Planning for Efficient Use of Your Mobile Client	45
Deciding Where to Locate a Home Router	45

3	Setting Up	51
	Turning Off IPX Packet Forwarding	51
	How to Turn Off IPX Packet Forwarding	52
	Configuring Static Routes and Services	53
	Configuring Static Routes and Services with NIASCFG	54
	Configuring Static Routes and Services with STATICON	56
	Configuring Watchdog Spoofing	65
	How to Configure Watchdog Spoofing on an Interface	66
	How to Configure Watchdog Spoofing for Call Destinations	66
	Configuring Routed or Static On-Demand Calls	67
	How to Configure Routed or Static On-Demand Calls	68
	Configuring IPX and NCP Header Compression	69
	How to Configure IPX and NCP Header Compression on an Interface	70
	How to Configure IPX and NCP Header Compression per Call Destination	71
	Configuring NLSP	72
	How to Configure NLSP	74
	How to Change the LSP Size	76
	Configuring RIP and SAP	77
	How to Configure RIP	78
	How to Configure SAP	80
	Accepting and Advertising Services from a Network Not Listed in the Routing Information Table	
	82	
	Proxying a NetWare File Server	84
	How to Proxy a NetWare File Server	84
	How to Check the Proxy Configuration	85
	Configuring the IPX Address Mapping Gateway	86
	Configuring IPX Route Aggregation	88
	Controlling the Propagation of Type 20 Packets	89
	How to Control Propagation of Type 20 Packets	90
	Changing the Hop Count Limit for IPX Packets	90
	How to Change the Hop Count Limit	91
	Balancing Traffic Loads over Equal-Cost Routes	92
	How to Balance Traffic Loads over Equal-Cost Routes	93
	Configuring SPX Connection Parameters	94
	How to Configure SPX Connection Parameters	95
	Setting Delay and Throughput for a Slow Link	95
	How to Set Delay and Throughput for a Slow Link	97
	Configuring IPX for Wireless Connectivity	97
	Configuring a Home Router	98
	Configuring a Mobile Client	99
	Configuring the MacIPX Gateway	105
	Configuring and Binding the Gateway Driver	107
	Restricting Gateway Service to Selected Networks	109

4	Managing	111
	Using the IPXCON Utility	111
	Using the IPXPING Utility on the Server	112
	Using the IPXPING Utility on the Workstation	113
	Syntax	113
	Parameters	113
	Example	114
	Using the SPFCON Utility	114
	Main Window	115
	Interfaces Window	115
	Connections Window	115
	Spoofing Statistics Window	115
	Viewing NetWare IPX Configuration Information	116
	Determining Whether a Remote IPX Router Is Reachable	116
	Determining Which IPX Services Are Reachable	117
	Checking an IPX Network for Inactive Routers	117
	Checking the IPX Routing Table	117
	Checking an IPX Network for Duplicate Network Numbers	118
	Checking an IPX Network for Duplicate System IDs	118
	Determining Where NLSP Is Running in Your Network	119
	Finding NLSP Routers with Insufficient Memory	119
	Finding the Designated Router on a LAN	120
	Monitoring Error Counters	121
	Viewing the MacIPX Gateway Configuration	122
	Viewing MacIPX Gateway Statistics	122
5	Troubleshooting	125
	Troubleshooting Tools	125
	IPXCON	125
	System Console Commands	126
	Configuration Tips	127
	Troubleshooting Checkpoints	127
	IPX Checkpoints	127
	IPX Connectivity Problems (Duplicate ID or Network Number)	129
	NLSP Checkpoints	131
	Common Problems	131
	Login Times Out	132
	Load Balancing over IPX Is Not Working	133
	Only One IPX Packet Is Sent and Received	133
	IPXCON Counters Are Increasing (Duplicate ID or Network Number)	133
	Error Messages Are Displayed (Duplicate ID or Network Number)	134
	NLSP Decision Process Is Running Frequently (Duplicate System ID)	135
	Other Router Names Are Not Displayed	135
	System Frequently Appears and Disappears on the LAN	135

Multiple Systems on a LAN Become Unreachable Intermittently	138
Connectivity Across a Point-to-Point Link Has Been Lost	139
An NLSP Server on a LAN Cannot Be Accessed	140
LAN Is Partitioned	141
No Communication Occurs between Two Networks.	142
Services Are Inaccessible in the Area	143
Number of Routes and Services on a System Shows Local Connectivity Only	144
Services or Routes are Fluctuating Excessively.	145
Heavy Network-Layer Traffic Occurs on a Point-to-Point Link.	146
Applications Perform Poorly	147
CALLMGR Shows an IPX Circuit but IPXCON Does Not	149
Many Systems Are Entering an Overloaded State	149
Connectivity Is Lost on Only One LAN.	150
NetWare Mobile IPX Client Loses Connectivity to the Server	151

About This Guide

This guide provides the information you need to configure and manage the Novell[®] Internet Access Server 4.1 IPX[™] routing software. In addition to planning information, this guide provides troubleshooting tips, techniques, and tools, as well as the symptoms of and solutions to commonly occurring problems for the IPX components of Novell Internet Access Server 4.1.

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

Also, a trademark symbol ([®], [™], etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

1

Understanding

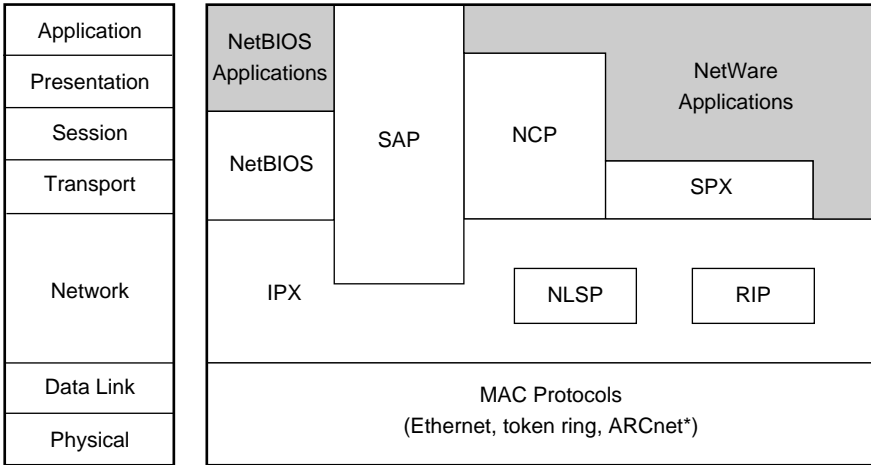
This section describes the processes and protocols that govern Internetwork Packet Exchange™ (IPX™) routing in the NetWare® networking environment. In particular, it examines the mechanics of IPX packet routing and the administration of routing and service information on an IPX internetwork.

Each NetWare protocol plays a different role in enabling a NetWare router to perform its tasks. Media access control (MAC) protocols and IPX provide the addressing mechanism that delivers packets to their destination. The Routing Information Protocol (RIP), Service Advertising Protocol (SAP), and NetWare Link Services Protocol™ (NLSP™) protocols provide the means by which routers gather routing and service information and share it with other routers on an internetwork.

Although the NetWare Core Protocol™ (NCP™) software does not play a direct role in routing, it does provide session control and packet-level error checking between NetWare workstations and routers. Similarly, the Sequenced Packet Exchange™ (SPX™) protocol neither routes packets nor advertises service information, but guarantees delivery of each packet to its destination.

Figure 1 on page 12 shows how the NetWare protocols correspond to the Open Systems Interconnection (OSI) reference model. Because this model represents only a basic framework for networking functionality, not all NetWare protocols fit neatly into a single functional layer.

Figure 1 How NetWare Protocols Correspond to the OSI Reference Model



The higher-level protocols (NetBIOS, SAP, NCP, SPX, NLSP, and RIP) rely on the MAC protocols and IPX to handle lower-level communications, such as node addressing. With the exception of NetBIOS, NCP, and SPX, each of these protocols plays a role in the operation of IPX routing.

The IPX Protocol

Novell adapted IPX from the Xerox* Network System (XNS*) Internet Datagram Protocol (IDP). IPX is a connectionless datagram protocol. *Connectionless* means that when a process running on a particular node uses IPX to communicate with a process on another node, no connection between the two nodes is established. Thus, IPX packets are addressed and sent to their destinations, but there is no guarantee or verification of successful delivery. Any packet acknowledgment or connection control is provided by protocols above IPX, such as SPX. *Datagram* means that each packet is treated as an individual entity, having no logical or sequential relation to any other packet.

As shown in [Figure 1](#), IPX operates at the OSI Network layer. As a Network-layer protocol, IPX addresses and routes packets from one location to another on an IPX internetwork. IPX bases its routing decisions on the address fields in its header and on the information it receives from RIP or NLSP. IPX uses this information to forward packets to their destination node or to the next router providing a path to the destination node.

For more information about the IPX protocol, refer to

- ♦ “IPX Packet Structure” on page 13
- ♦ “IPX Addressing” on page 17
- ♦ “How IPX Routing Works” on page 21

IPX Packet Structure

The IPX packet is similar to an XNS IDP packet and comprises two parts:

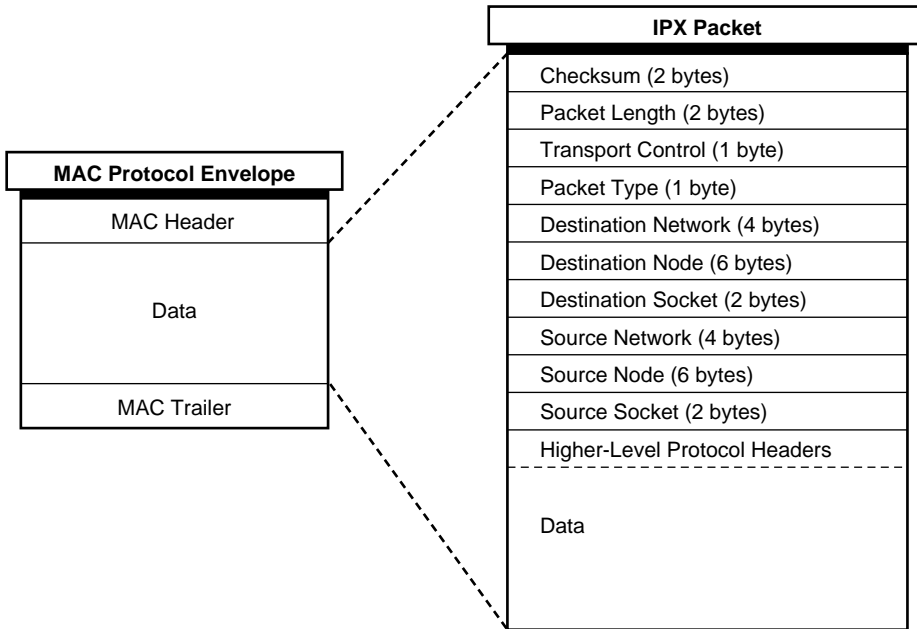
- ♦ A *30-byte IPX header*, which includes the network, node, and socket addresses for both the destination and the source
- ♦ A *data* section, which often includes the header of a higher-level protocol, such as SPX

The minimum IPX packet size—excluding the MAC header—is 30 bytes (IPX header only). Historically, the maximum size of routed IPX packets has been 576 bytes (IPX header and data). Until recently, all routed IPX packets were between 30 and 576 bytes. However, the IPX protocol has always allowed packet sizes up to 65,535 bytes.

NOTE: Media constraints typically limit the actual maximum packet size allowed to something less than 65,535 bytes. Ethernet II packets, for example, are limited to a data size of 1,500 bytes, not including the MAC header.

The IPX header is placed after the MAC header and before the data. [Figure 2 on page 14](#) shows the structure of an IPX packet.

Figure 2 IPX Packet Structure



The following describes the IPX packet fields:

- ♦ *Checksum* —Packet integrity check.
The checksum is used by the NetWare SFT™ III™ software and NetWare 4 software. Older versions of NetWare did not use the IPX checksum and required that this field be set to 0xFFFF.
- ♦ *Packet Length* —Length, in bytes, of the complete packet, which is the length of the IPX header plus the length of the data.
The packet length is at least 30 bytes (for the IPX header).
- ♦ *Transport Control* —Number of routers a packet has traversed on the way to its destination.
IMPORTANT: On a traditional, RIP-based IPX router, IPX packets whose Transport Control field reaches a value of 16 are discarded. With NLSP, an IPX packet can travel up to 127 hops to reach its destination. You make this possible by setting the Hop Count Limit parameter from the Novell Internet Access Server Configuration utility (NIASCFG). This enables you to limit the number of routers (hops) an IPX packet traverses before it is discarded.

Sending nodes always set the Transport Control field to zero when building an IPX packet. When a router receives a packet that requires further routing, it increments this field by one and forwards the packet.

- ◆ *Packet Type*—Type of service offered or required by the packet.

Novell currently uses the packet types listed in [Table 1](#).

Table 1 Packet Types

Packet Type	Field Value (Hex)	Purpose
NLSP	0x00	NLSP packets
Routing information	0x01	RIP packets
Service advertising	0x04	SAP packets
Sequenced	0x05	SPX packets
NCP	0x11	NCP packets
Propagated	0x14	NetBIOS and other propagated packets

- ◆ *Destination Network*—Number of the network to which the destination node is attached.

When a sending node sets this field to 0x0 (that is, 0x00000000), the destination node is assumed to be on the same network segment as the sending (or source) node.

A special case exists when a workstation sends SAP Get Nearest Server and RIP Get Local Target (or Route Request) broadcast requests at initialization time. Because the workstation does not yet know which network it belongs to, it sets both the Source Network and Destination Network fields to 0 for these requests. When a router receives one of these requests, it sends a reply directly to the sending workstation, filling in the Source Network and Destination Network fields with the appropriate network numbers.

NOTE: IPX does not have a broadcast network number (such as 0xFFFFFFFF).

In addition to network number 0, the numbers 0xFFFFFFFF and 0xFFFFFFFFE are reserved for specific purposes. For this reason, they should not be assigned to any IPX network. For more information about

reserved network numbers, refer to “Reserved Network Numbers” on page 18.

- ◆ *Destination Node* —Physical address of the destination node.

Not all LAN topologies use the same size address field. A node on an Ethernet network requires all 6 bytes to define its address; a node on an Ammonite network requires only 1 byte.

A node address of 0xFFFFFFFF (that is, 6 bytes of 0xFF) broadcasts the packet to all nodes on the destination network.

- ◆ *Destination Socket* —Socket address of the packet destination process.

Sockets route packets to different processes within a single node. Novell reserves several sockets for use in the NetWare environment. Refer to Table 2 on page 20 for a partial list of NetWare socket numbers.

NOTE: IPX does not have a broadcast socket number (such as 0xFFFF).

- ◆ *Source Network* —Number of the network to which the source node is attached. If a sending node sets this field to zero, the local network to which the source is connected is unknown. For routers, the rules that apply to the Destination Network field also apply to the Source Network field, except that routers can propagate packets that were received with this field set to zero.

- ◆ *Source Node* —Physical address of the source node.

Broadcast addresses are not allowed.

- ◆ *Source Socket* —Socket address of the process that transmits the packet.

Processes communicating in a peer-to-peer fashion do not need to send and receive on the same socket number.

On a network of workstations and servers, the server usually listens on a specific socket for service requests. In such a case, the source socket is not necessarily the same or even significant. All that matters is that the server reply to the source socket. For example, all NetWare file servers have the same socket address, but requests to them can originate from any socket number.

Source socket numbers follow the same conventions as those for destination sockets.

- ◆ *Higher-Level Protocol Headers* —Headers of higher-level NetWare protocols, such as NCP or SPX. These headers occupy the data portion of the IPX packet.

IPX Addressing

IPX defines its own internetwork and intranode addressing. For intranetwork (node) addressing, IPX uses the physical address assigned to the network interface board.

The *IPX network address* uniquely identifies an IPX server on an IPX network and individual processes within the server. A complete IPX network address is a 12-byte hexadecimal number comprising the following components:

- ◆ A 4-byte network number (server)
- ◆ A 6-byte node number (server)
- ◆ A 2-byte socket number (server process)

The following is an example of a complete IPX network address:

FEDCBA98 1A2B3C5D7E9F 0453

Each number in an IPX address is contained in a field in the IPX header and represents a source or destination network, node, or socket. The network number is used only for Network-layer operations, namely routing. The node number is used for local, or same-segment, packet transmission. The socket number directs a packet to a process operating within a node.

For more information about each address component, refer to

- ◆ [“Network Number” on page 17](#)
- ◆ [“Node Number” on page 19](#)
- ◆ [“Socket Number” on page 20](#)

Network Number

The IPX network number is the 4-byte hexadecimal address that serves as the basis for IPX packet routing. Each network segment on an internetwork is assigned a unique network number. NetWare routers use this number to forward packets to their final destination network.

An IPX network number can contain up to eight digits, including zeros. (Leading zeros are usually not displayed.) For example, 0xFEDCBA98, 0x1234567D, and 0xC7 are all valid network numbers.

With the fast setup feature available from NIASCFG, the routing software can automatically detect the network number and data-link *frame type* used on an

IPX network. After you configure a board and select a driver during the initial router configuration, the router sends a RIP *all routes request* packet to the network. From the responses it receives, the router determines the network number and frame type it needs to use.

Reserved Network Numbers

The destination network of an IPX packet is typically an IPX network to which a unique network number has been assigned. However, three network numbers—0x0, 0xFFFFFFFF, and 0xFFFFFFF0—are reserved and cannot be used to identify a specific network. These numbers have the following meanings:

- ♦ **0x0** —Represents the local network segment. If a router receives a packet whose destination network number is 0, the packet's source and destination nodes are attached to the same segment.
- ♦ **0xFFFFFFFF** —Represents an *all routes request* between NetWare routers. If a router receives a packet whose destination network number is FFFFFFFF, it sends all the routes it knows about to the requesting router.
- ♦ **0xFFFFFFF0** —Represents the *default route*. This is an advertised destination to which IPX packets with unknown destination networks are forwarded.

With NetWare routing software, a router that receives an IPX packet with an unknown destination network can do one of two things: If another router on the network is advertising 0xFFFFFFF0, the router forwards the packet to that router. If 0xFFFFFFF0 is not advertised on the network, the packet's destination remains unknown and the router discards the packet.

Both RIP and NLSP have been modified to recognize 0xFFFFFFF0 as the default route. On a RIP network, the default route is typically advertised by a RIP router that connects the LAN to a larger network infrastructure, such as a corporate backbone or *transit LAN*.

The routing software cannot advertise the default route dynamically, but you can configure the router to advertise it as a *static route*. To read about static routes, refer to [“Static Routes and Services.”](#)

Internal Network Number

NetWare 3 servers and NetWare 4 servers have an additional identifier called an *internal network number*. This is a unique hexadecimal number between one and eight digits that is assigned to the server at installation. The internal network is a *logical* network that NetWare uses to advertise services and route IPX packets to the physical networks attached to the server.

The internal network number overcomes some routing and connectivity limitations inherent in NetWare 2. These are summarized in the following paragraphs.

A NetWare 2 server selects a primary interface and advertises its services as reachable through that interface. On a network with more than one server, packets might travel an extra hop to reach their destination.

A NetWare 2 server loses network connectivity if its primary network interface board fails, even if the server has Network-layer connectivity through another interface. Consider a NetWare 2 server with connections to two networks. The server advertises its services through the primary interface attached to one of the networks. If that interface fails, workstations attached to the server through the second network might not be able to log in to the server.

Node Number

The node number is the 6-byte hexadecimal address that identifies a device on an IPX network. This device can be a file server, router, workstation, or printer. The node number is identical to the physical address assigned to the interface board that connects the device to the network.

The IPX header contains a Destination Node field and a Source Node field. These fields contain the same destination and source node addresses found in the MAC header. A NetWare workstation, for example, uses the destination node address to locate and forward packets to another workstation on the same network segment.

IPX requires the node number to be unique only within the same IPX network. For example, a node on network FEDCBA98 can use the number 1A2B3C5D7E9F, and a node on network 1234567D can also use the number 1A2B3C5D7E9F. Because each node has a different network number, IPX recognizes each node as having a legitimate, unique address.

Socket Number

The socket number is the 2-byte hexadecimal number that identifies the ultimate destination of an IPX packet within the node. This destination is actually a *process*—such as routing (RIP) or advertising (SAP)—that operates within the node. Because several processes are typically operating at any given time, socket numbers provide a type of mail slot by which each process can identify itself to IPX.

A process that must communicate on the network requests that a socket number be assigned to it. Any packets that IPX receives that are addressed to that socket are passed to the process. Socket numbers provide a quick method of routing packets within a node.

Table 2 lists some socket numbers and processes used in the NetWare environment.

Table 2 NetWare Socket Numbers and Processes

Socket Number	Process
0x451	NCP
0x452	SAP
0x453	RIP
0x455	Novell NetBIOS
0x456	Diagnostics
0x9001	NLSP
0x9004	IPXWAN™ protocol

(For information about IPXWAN, refer to “**IPX Operation over WAN Links**” on page 24.)

Socket numbers between 0x4000 and 0x7FFF are *dynamic* sockets; these are used by workstations to communicate with file servers and other network devices. Socket numbers between 0x8000 and 0xFFFF are *well-known* sockets; these are assigned by Novell to specific processes. For example, 0x9001 is the socket number that identifies NLSP. Software developers writing NetWare applications can contact Novell to reserve well-known sockets.

How IPX Routing Works

NetWare routers interconnect different IPX network segments and receive instructions for addressing and routing packets between these segments from the IPX protocol. IPX accomplishes these and other Network-layer tasks with the help of RIP, SAP, and NLSP.

For more information about how IPX routing works, refer to:

- ◆ [When a Workstation Sends an IPX Packet](#)
- ◆ [When a Router Receives an IPX Packet](#)
- ◆ [When a Router Forwards an IPX Packet](#)

When a Workstation Sends an IPX Packet

Consider a NetWare workstation that wants to send data to another workstation. If both workstations share the same network number (both are on the same segment), the sending workstation addresses and sends packets directly to the destination workstation's physical address. If the two workstations have different network numbers (each is on a different segment), the sending workstation must first find a router on its own segment that can forward packets to the segment on which the destination workstation resides.

To find this router, the sending workstation broadcasts a RIP packet requesting the fastest route to the destination segment. The router on the sending segment with the shortest path to the destination segment responds to the request. In its response, the router includes its own network and node address in the IPX header.

NOTE: If the sending node is a router instead of a workstation, the router does not need to broadcast a RIP request to obtain this information; the router obtains the information from its internal routing table.

When the sending workstation knows the router node address, it addresses and sends packets to the destination workstation as follows:

1. The sending workstation places the destination node IPX network address—network, node, and socket numbers—in the corresponding destination fields of the IPX header.
2. The sending workstation places its own IPX network address—network, node, and socket numbers—in the corresponding source fields of the IPX header. The sending workstation also fills out all other fields in the header.

3. The sending workstation places the node address of the router that responded to the RIP request in the Destination Address field of the MAC header.
4. The sending workstation places its own node address in the Source Address field of the MAC header.
5. The sending workstation sends the packet.

When a Router Receives an IPX Packet

When a router receives an IPX packet, it performs the following tasks:

1. The router checks the Transport Control field of the IPX packet header.
A RIP router discards the packet if the value in this field is greater than 16.
An NLSP router discards the packet if the value in this field is greater than the value of the Hop Count Limit parameter.
2. The router checks the IPX header Packet Type field.
Note that if the packet type is 20 (0x14, NetBIOS), the packet is handled as follows:
 - a. The router examines the Transport Control field of the IPX header. If this value is 8 or greater, the router discards the packet. (Propagation of a NetBIOS packet is limited to eight networks.)
 - b. The router compares each network number entry in the packet to the network number of the segment on which the router received the packet.
If the router finds a match, it discards the packet to prevent multiple traversals of the same network segment. If the router finds no match, it performs the next step.
 - c. The router places the address of the network segment from which the packet arrived in the next available Network Number field.
 - d. The router increments the Transport Control field of the IPX header and broadcasts the packet to all directly connected network segments that are not represented in the Network Number fields.
3. The router checks the IPX header Destination Address fields—network, node, and socket numbers—to determine how to route the packet.

If the packet is addressed to the router, the appropriate socket process handles it internally; otherwise, the router forwards the packet.

NOTE: Packets that must be handled internally are those addressed directly to the router and those broadcast (destination node = 0xFFFFFFFFFFFF) to any network segment to which the router is directly connected. Usually, only RIP, SAP, and diagnostic packets fall into this category.

When a Router Forwards an IPX Packet

When forwarding packets, the router can take one of two possible actions. If the packet is destined for a network number to which the router is directly connected, the router performs the following tasks:

1. The router places the destination node address from the IPX header in the Destination Address field of the MAC header.
2. The router places its own node address in the Source Address field of the MAC header.
3. The router increments the Transport Control field of the IPX header and forwards the packet to the destination node segment.

IMPORTANT: If the Transport Control field equals the maximum allowed hop count before the field is incremented, the router discards the packet. For RIP routers, the hop count limit is 16; for NLSP routers, this limit can be set to any number between 8 and 127.

Note also that broadcast packets are never rebroadcast onto the network segment from which they are received.

If the router is not directly connected to the segment on which the final destination node resides, it sends the packet to the next router in the path to the destination node, as follows:

1. The router places the node address of the next router in the Destination Address field of the MAC header.
The router gets this information from its Routing Information Table.
2. The router places its own node address in the Source Address field of the MAC header.
3. The router increments the Transport Control field in the IPX header and forwards the packet to the next router.

IPX Operation over WAN Links

For packets to travel between two IPX network segments separated by a WAN, there must be a connection between the two routers representing each segment. This connection is represented by the *WAN call destination*, a unique name that identifies the router on the other end of the connection.

A WAN connection can be initiated by any of the following methods:

- ♦ **Automatic** —Established between the local router and the call destination when IPX is bound to a WAN port. Automatic connections are established typically at router startup.
- ♦ **Manual** —Established between the local router and the call destination by a user from the Call Manager utility (CALLMGR).
- ♦ **Data-initiated** —Established when the local router needs to send data to the (remote) router represented by the WAN call destination. This connection method is characteristic of *on-demand calls*, which are described in the following section.

After a WAN connection is established, the routers use the IPXWAN protocol to negotiate the values or states of various connection characteristics, such as speed, throughput, routing type, and IPX header compression. These and other characteristics are negotiated before the routers exchange any routing information or data.

For more information about IPX operation over WAN links, refer to

- ♦ [“Call Types” on page 25](#)
- ♦ [“Routing Types” on page 27](#)
- ♦ [“Static Routes and Services” on page 28](#)
- ♦ [“Watchdog Packet Spoofing” on page 30](#)
- ♦ [“Header Compression” on page 31](#)

Call Types

Associated with each WAN call destination is the *call type*, which characterizes the behavior of the call after it is established. Calls can be *permanent* or *on-demand*.

For more information about call types, refer to:

- ♦ “Permanent Calls” on page 25
- ♦ “On-Demand Calls” on page 25
- ♦ “Routed On-Demand Calls” on page 26

Permanent Calls

A *permanent call* is a connection that remains active between the local router and the remote router identified by the call destination. A permanent call can be established automatically from configured protocol-to-board bindings, or manually from CALLMGR. The call remains active until IPX is unbound from the interface, or until the connection is disconnected manually from CALLMGR. If a permanent call fails, IPXRTR tries to reestablish the connection.

IPX routing and service information crosses permanent calls as required by the operative routing/service protocol, which can be RIP/SAP or NLSP. If you do not want routing and service traffic to cross a permanent-call link, the routing software enables you to configure *static routes and services* on each router. This is typically how IPX routers are made aware of remote routes and services over links that use on-demand calls.

For information about static routes and services, refer to “[Static Routes and Services](#)” on page 28.

For more information about permanent calls, refer to “Understanding” in the routing documentation for NetWare/Link PPP.

On-Demand Calls

An *on-demand call* is a dedicated, point-to-point connection between two routers that becomes active only when one router must send user data to the other. Because the on-demand call relies on configured static routes and services, no routing or service information crosses the link while the call is active.

With an on-demand call, the link remains inactive until user data needs to cross it. Workstations needing to reach remote destinations send packets to their local IPX router advertising the routes, assuming the packets can reach their destination. The local router stores the packets and tries to establish a connection to the remote router. After the local router completes the call and negotiates on-demand service, it forwards the stored packets to the remote router, which then forwards them to their destination.

NOTE: To avoid activating potentially expensive connections, IPX routers do not forward type 20 (NetBIOS) packets over on-demand calls.

For more information about on-demand calls, refer to “Understanding” in the routing documentation for NetWare/Link PPP.

Routed On-Demand Calls

NetWare routing also enables you to configure a *routed on-demand call*. Unlike the standard on-demand call, which relies on statically configured routes and services, a routed on-demand call runs a routing protocol while the link is active. When the link goes down, the routes and services made known by the routing protocol become unavailable.

If no data crosses the link after some period of time, a Data-Link layer timer triggers the termination of an on-demand call. However, the routing protocol running over a routed on-demand call resets this timer each time it transmits a packet. This keeps the link continuously active. To solve this problem, the routing software uses a similar timer that operates at the Network layer. This timer is reset only when data packets—not protocol packets—cross the link. In this way, the routing updates do not keep the link active when no data is being transmitted.

Routed on-demand calls are well-suited for large corporate networks that have many branch offices. In this type of internetwork, most of the traffic is unidirectional: from the branch office to the corporate network. Configuring each branch office with a single (default) route to the corporate network is sufficient. When a branch office router establishes a link to the router serving the corporate network, the routing protocol floods the branch office routes into the corporate network. This is necessary so that responses to branch office service requests know how to reach their destination in the branch office network. As long as the branch office forwards information to the corporate network, the link remains active. If the link is idle for some predetermined period of time, it goes down.

You configure a routed on-demand call the same way you configure a standard on-demand call with one exception: you must configure a routing protocol to operate over the link.

Routing Types

IPXWAN negotiates the *WAN routing type*, which determines which IPX routing protocol—if any—runs over the connection. NetWare routing supports the following routing types for IPX exchanges over a WAN connection:

NOTE: The first four routing types operate only between routers; the fifth, *WAN Workstation*, operates between a router and a NetWare workstation.

- ◆ **WAN NLSP** —Uses NLSP to exchange routing and service information over the connection. This is the most efficient—and preferred—routing type for WAN connections requiring an active routing protocol. A WAN NLSP connection does not require an IPX network number.
- ◆ **Unnumbered RIP** —Uses RIP and SAP to exchange routing and service information over the connection but requires no IPX network number.
- ◆ **Numbered RIP** —Also uses RIP and SAP to exchange routing and service information over the connection but does require an IPX network number.
- ◆ **On-Demand** —Uses no active routing or service advertising protocol, but rather a set of static routes and services on each router.
- ◆ **WAN Workstation** —Enables a NetWare workstation to connect to an IPX internetwork through a remote router. No routing protocol runs over the connection, except when the workstation sends a route or service request to the router.

To choose the most suitable routing type, IPXWAN considers the following criteria during its negotiation process:

- ◆ Which versions of the routing software are running on the routers
Earlier versions of the routing software, such as NetWare Multiprocessor Router Plus™ 2.1x software and NetWare WAN Links™ 2.0 software, support only Numbered RIP connections.
- ◆ Which routing protocol—RIP or NLSP—is enabled on the WAN interfaces at each end of the connection

For example, two routers running NLSP at their respective WAN interfaces automatically use the WAN NLSP routing type over the connection.

- ◆ Whether a third-party router is running at the other end of the connection

Some third-party routers might support only Numbered RIP connections for IPX routing over WANs.

- ◆ Whether the WAN call destination is configured as a permanent call or an on-demand call

WAN NLSP, Unnumbered RIP, and Numbered RIP operate only over permanent calls; the On-Demand routing type operates only over on-demand calls.

Static Routes and Services

A *static route* is a RIP route that is added to the Routing Information Table by a network administrator, rather than by the active routing protocol—in this case, RIP—operating over a network link. For a WAN connection, a static route comprises a WAN call destination, the destination IPX network number, and the route metrics (hops and ticks) to reach the destination. A *static service* is a SAP service that is also added manually rather than dynamically by SAP. A static service comprises a WAN call destination; the service name and type; the service address network, node, and socket; and the service metrics (hops and ticks) to reach the destination advertising the service. With the routing software, you can configure static routes and services for both permanent and on-demand calls.

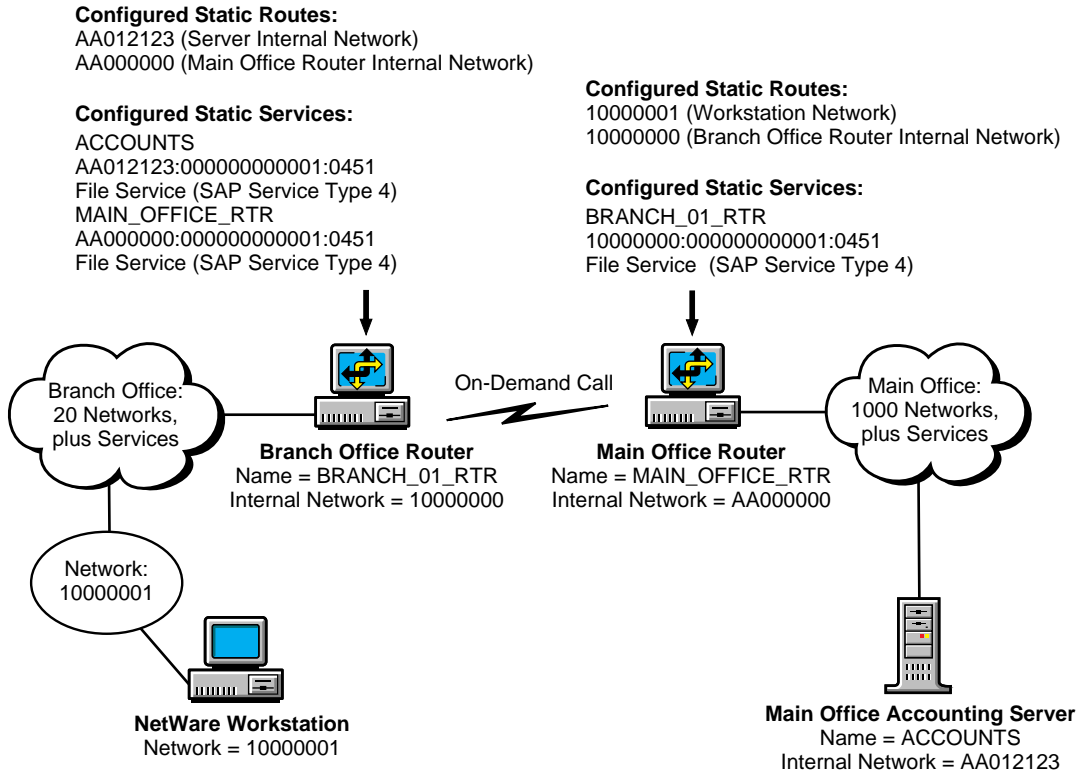
When used with permanent calls, static routes and services are useful for redirecting traffic to a particular network, perhaps for security reasons, and for conserving bandwidth on slow or low-capacity links. A single static route is also useful as a default route. In this way, the only routing information crossing the link is that required by users to access a specified set of services.

When used with on-demand calls, static routes and services are useful for connections that use expensive telecommunications carriers and for slow links over which it is undesirable to exchange routing and service information. Consider an internetwork that connects tens to hundreds of branch offices to a single main office. Typically, each branch office requires periodic access to information at the main office. However, it is most likely that the main office periodically polls the branch offices to get up-to-date information, such as the day's sales figures. Because a permanent call to each branch office is not

necessary, connections to the main office need only be low-speed, dial-up lines. For this reason, the first several minutes of the call should not be taken up by a flood of routing and service information into a branch office. Nor should there be a relatively smaller flood of (mostly irrelevant) routing and service information from a branch office into the main office.

Figure 3 shows a typical configuration for static routes and services over an on-demand call.

Figure 3 On-Demand Call Between a Branch Office and the Main Office



In this configuration, the branch office router, BRANCH_01_RTR, must know only the addresses and names of a few servers and services. This small number of extra routes and services is of minimal burden to the branch office network. The main office router, MAIN_OFFICE_RTR, must keep track of only a few networks and services from each branch office. This is significantly better than being flooded with potentially thousands of extra routes and services that are of no use to the main office network.

To configure static routes and services for permanent and on-demand calls, you can use either of the following utilities:

- ♦ NIASCFG, which you also use for configuring boards and network interfaces, and for enabling protocols and bindings on the router. If you use NIASCFG, you must configure *all* routes and services manually.
- ♦ STATICON, the static routes and services configuration utility for IPX. STATICON discovers which routes and services are available through a remote router and then adds them automatically to the static routing table on a local router. Because STATICON gathers and exchanges the information automatically, it is essential for configuring large numbers of static routes and services.

Watchdog Packet Spoofing

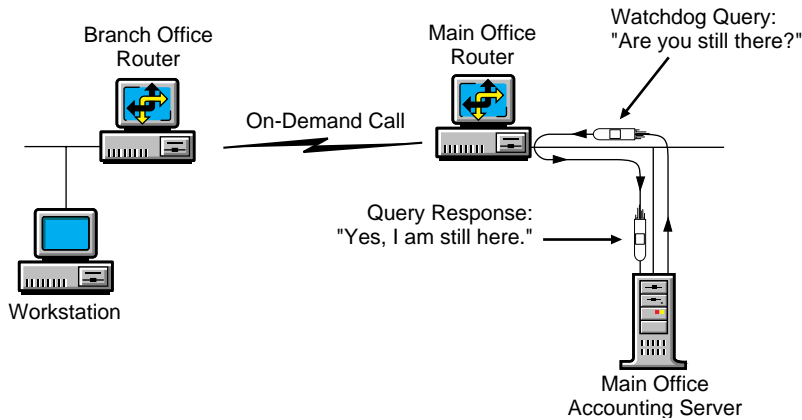
NetWare servers use the *Watchdog protocol* to validate workstation connections periodically. When a workstation is logged in to a server but has not transmitted a packet for some period of time (the default is 5 minutes), the server sends a watchdog query packet to the workstation. If the workstation does not reply with a watchdog response packet after 5 minutes, the server sends additional queries at specified intervals until 15 minutes have elapsed. If the workstation still has not replied, the server terminates the connection.

With several workstations operating over an on-demand call, the exchange of watchdog packets can keep the connection active most of the time. Depending on the telecommunications carrier you use for the connection, this can become expensive.

You can avoid this problem by configuring your router to perform *watchdog spoofing*. This means that the router captures watchdog query packets on their way to a workstation and responds on behalf of the workstation without activating the on-demand call. Because of the spoofing, however, the workstation's server connection remains occupied unless it logs out. A way to avoid this is for the remote server to execute a forced logout of all workstations at a predetermined time (midnight, for example), so that all server connections are freed for the next day.

Figure 4 on page 31 shows how watchdog spoofing works over an on-demand call.

Figure 4 Watchdog Spoofing Enabled over an On-Demand Call



When watchdog spoofing is enabled on an on-demand call, watchdog packets, going from a server to a client, cause the router to reply that the workstation is active without initiating the call. If watchdog spoofing is disabled, an on-demand call is initiated for each watchdog packet that crosses the connection.

Header Compression

NOTE: NCP header compression is not used for NCP packets using the Packet Burst™ protocol. Because IPX headers are the standard, IPX header compression is used.

Header compression increases the throughput of IPX and NCP packets over low-speed serial lines (except for NCP packets using the Packet Burst protocol). An IPX packet header is 30 bytes and is typically followed by an upper-layer protocol header, such as an SPX header. Header compression reduces the size of this combined packet header to just a few bytes.

Header compression is negotiated by the IPXWAN protocol when a call is established over any WAN connection type. Header compression is not used on the connection if IPXWAN detects that one of the routers does not support it. The routers at each end of the connection must have header compression enabled and must allocate the same number of *compression slots*.

For more information about header compression, refer to

- ♦ “Compression Slots” on page 32
- ♦ “Compression Packet Types” on page 33

Compression Slots

When you enable header compression, you can also specify the number of compression slots. A compression slot is a location in router memory that stores packet header information. The compression algorithm uses this information to compress outgoing—and decompress incoming—packet headers.

IMPORTANT: You must allocate the same number of compression slots on each router. If the values are different, the IPXWAN protocol chooses the lesser of the two.

If too few compression slots are allocated for the number of different-style packets crossing the connection, the values in the following IPXCON counters become large:

- ◆ Initialization Packets Sent
- ◆ Initialization Packets Received
- ◆ Uncompressed Packets Sent
- ◆ Uncompressed Packets Received

The compression algorithm is running efficiently if the number of compressed packets sent is significantly higher than the values in these counters.

A router sends an uncompressed packet when it is considered beneficial not to reuse a compression slot.

Allocating too many compression slots has its own consequences. More memory is required to store all the headers, and the compression algorithm must scan through more stored headers to find a match for each transmitted packet. This results in a higher processing load and slower performance.

Compression Packet Types

Five packet types are used to exchange compression-state information about packets sent over a connection on which header compression is enabled. Three of these packet types—slot initialization, reject, and acknowledgment packets—manage the flow of compressed and uncompressed packets over the connection; these are the *compression protocol* packets. The packet type, along with other information, is indicated in the first byte of a compressed packet. Compression packet types are defined as follows:

- ◆ *Uncompressed packet* —A standard, uncompressed IPX or NCP packet. A router sends an uncompressed packet when the packet cannot be compressed or a decision was made not to compress it. When the remote router receives an uncompressed packet, it simply removes the 1-byte compression header and passes the packet to IPX.

IPXCON tracks uncompressed packets exchanged on a connection in the Uncompressed Packets Sent and Uncompressed Packets Received counters.

- ◆ *Compressed packet* —A compressed IPX or NCP packet. Compressed packets do not contain the standard packet header. Instead, they contain the number of a compression slot on the receiving router. This slot contains the information necessary for the compression algorithm to decompress the packet header before passing the packet to IPX.

IPXCON tracks compressed packets in the Packets Sent and Packets Received counters.

- ◆ **Slot initialization packet** —A compression protocol packet that a router sends to prepare a compression slot on the receiving router for use. Initialization packets can prepare new slots or previously used slots for reuse. The routing software uses two different initialization packets: one for IPX packets and one for NCP packets.

IPXCON tracks initialization packets in the Initialization Packets Sent and Initialization Packets Received counters.

- ◆ **Reject packet** — A compression protocol packet that a router sends when it receives another compression protocol packet that it does not understand.

IPXCON tracks reject packets in the Reject Packets Sent and Reject Packets Received counters.

- ◆ **Acknowledgment packet** — A compression protocol packet that acknowledges receipt of an IPX slot initialization packet. The transmitting node continues to send slot initialization packets until it receives an acknowledgment packet. After receiving the acknowledgment packet, the transmitting node begins to send compressed packets.

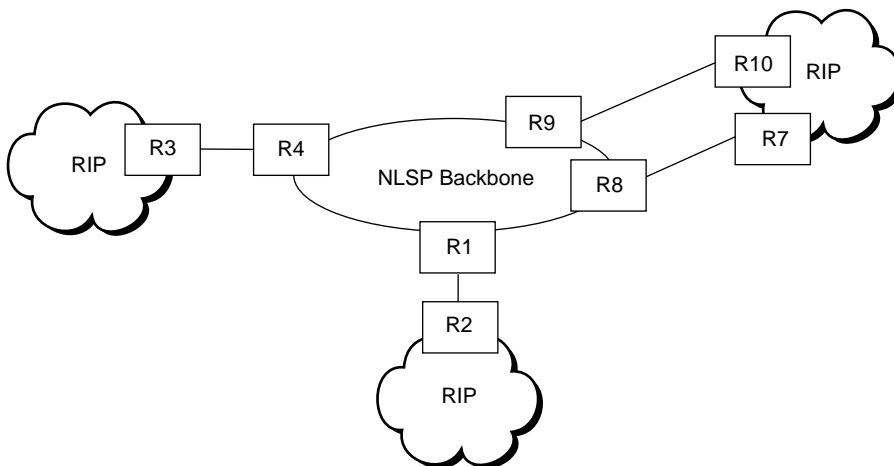
IPXCON does not track acknowledgment packets.

IPX Route Aggregation

IPX Route Aggregation enables you to introduce routes learned through RIP into an NLSP backbone in a summarized form. Route aggregation compactly describes many IPX network numbers simultaneously by using an address and mask pair. For example, all addresses from C9000000 to C9FFFFFFF can be represented using the address C9000000 and the mask FF000000.

Figure 5 on page 34 shows a typical topology for using route aggregation, a backbone network with several RIP areas attached. In Figure 5 on page 34, the information within the backbone is minimized by having routers R1, R4, R8, and R9 report address summaries for the attached areas. In this scenario, the default route is the only information about reachable external addresses that must be transmitted within an attached area. Therefore, information within the attached areas is minimized.

Figure 5 Aggregated Routes Topology



For more information about IPX route aggregation, refer to

- ♦ “Introducing Aggregated Routes into NLSP” on page 35
- ♦ “Consistent Use of Routers that Support Route Aggregation” on page 36
- ♦ “Interaction with SAP” on page 36
- ♦ “Metrics Used with Aggregated Routes” on page 36

Introducing Aggregated Routes into NLSP

Aggregated routes are introduced into NLSP in the same way that external RIP routes are introduced. There are two methods of introducing aggregated routes into NLSP:

- ♦ Aggregated routes are introduced through static configurations. If an aggregated route is configured for a static routing link, the configured aggregated route is reported into the NLSP area as soon as IPXRTR binds to the interface.
- ♦ Aggregated routes are learned from RIP. An address summary can be configured for a link, but unless at least one address matches the address summary, as learned through RIP, the summary is not reported.

For example, if a router is configured with address summaries 572* and 5729* on a link running RIP and learns from that link that the destination 57285489 is reachable and that no other matching destinations are learned, then the router reports the aggregated route 572* to the NLSP area. The asterisk represents a wildcard character. If the router learns that destination 57298381 is reachable, then only 5729* is reported to the NLSP area. If both 57298381 and 57212376 are reachable, then the router reports both 572* and 5729* to the NLSP area.

Routers always report aggregated routes with the longest match. For example, if a router is configured with address summaries C9* and C91* and learns that the destination C9123456 is reachable, then the router reports only the aggregated route C91*.

Consistent Use of Routers that Support Route Aggregation

Route aggregation into an NLSP area is possible only if all routers in that area support address summaries. Routers that do not support route aggregation do not recognize destination addresses for aggregated routes; they forward packets to the default route or drop packets if no default route is configured. Because routers that support route aggregation route packets differently from routers that do not support route aggregation, routing loops can occur. Therefore, if a router that supports route aggregation detects that the next-hop router on the path to an aggregated destination is a router that does not support route aggregation, it will drop the packet.

Interaction with SAP

Without route aggregation, if a service is announced with an address that is not explicitly reachable, RIP assumes the service is unreachable. When NLSP and RIP are used with route aggregation, SAP is accepted, provided a packet can be forwarded to the corresponding network number. A packet can be forwarded when there is a default route and address summary or when there is an explicit advertisement of that network number.

This feature is disabled by default. It can be enabled by entering the following command after IPXRTR.NLM has been loaded:

```
SET REQUIRED NETWORK FOR SERVICES=ON
```

Metrics Used with Aggregated Routes

If a router is configured to introduce an address summary into an NLSP area, it does so, with the number of ticks equal to 1 (as the default), an area count equal to 6 (as the default), and the ability to configure both parameters with a different value. The area count is the number of areas through which the route information is allowed to propagate. As the route information is passed through each of these areas, the area count is decremented by one. This enables the initial sending router to control the spread of information through all NLSP areas that are connected to each other.

Address summaries are not exported from an NLSP area into a RIP cloud because the capability to express a summarized route within RIP does not exist. In addition, updating RIP with this capability is not cost-effective because the default route is sufficient for RIP.


IPX Address Mapping Gateway

Using the IPX Address Mapping Gateway (IAMG) offers the following three advantages:

- ◆ Your hosts can connect to a backbone network even when your local network numbers are not compatible with the backbone addressing scheme.
- ◆ If the routing protocol in the backbone does not support route aggregation (such as earlier implementations of NLSP), the routing protocol probably cannot manage the number of network addresses from every customer. Even if the routing protocol could handle route aggregation, network numbers might be assigned in a way that does not lend itself to aggregation. IAMG enables the summarization of routes in a manner transparent to the routing protocol by mapping many network numbers to a single number outside the local network. This capability greatly reduces the number of networks that must be advertised throughout the global internetwork.
- ◆ The security of the local intranet is enhanced because local networks are not advertised throughout the global internetwork.

In [Figure 6 on page 38](#), any packets generated from the client on the customer network for the backbone have their source IPX network number converted to the 01014000 gateway address. In addition to mapping a client network address to an address compatible with the backbone, IAMG converts the IPX node address to a unique value based on characteristics of the original node number. Services on the customer network, such as file servers with Novell[®] Directory Services[®] (NDS[®]) that advertise their network addresses through SAP and that must be visible in the backbone, are not translated. To be visible in the backbone network, the services must use a registered backbone address. All registered source addresses are left untranslated.

Figure 6 IAMG Implementation



The IAMG allows multiple gateways to be connected between the customer network and backbone network in parallel. All such gateways on a customer's network share a single gateway address. If a gateway receives a packet from the backbone network with an unknown mapping on the gateway network, the packet is forwarded to all other IAMGs advertising the same gateway address. If necessary, the packet is split into multiple segments using the IPX fragmentation module so that the packet can be forwarded with a new header attached.

If not all gateways support the IPX fragmentation specification, the IAMG is designed to allow the gateways to run in parallel. In this automatic mode of operation, some packets might be lost during the learn cycle. When the gateway that created the mapping receives the forwarded data packets from the backbone, an update reporting the mapping is returned to the gateway that initially received the packet. Because more packets probably will follow, the gateway learns the new mapping.

The IAMG discards any packet destined for the gateway that has the broadcast node address so that outside clients cannot cause an excessive number of broadcasts within a network. NetBIOS broadcasts (packet type 20) cannot be used on a network attached to an interface that has the IAMG enabled. Broadcasts to networks other than the gateway address are allowed.

When the IAMG is configured, care must be taken to avoid address conflicts. We recommend using the following configuration guidelines:

- ◆ Conflicts might occur when token ring and Ethernet interfaces are used on the same network. For instance, when the NetWare Mobile IPX™ software and NetWare remote access software assign node addresses, they set the IEEE Local bit and clear the IEEE Multicast bit. Because token ring networks use a different bit order (canonical instead of noncanonical) within the node address, incorrect interpretations can be made, and locally assigned addresses might conflict with IEEE token ring addresses. It is also possible for IEEE Ethernet node addresses to conflict with IEEE token ring addresses, even though the IEEE assigned addresses are different. You can avoid conflicts by taking the following precautions:
 - ◆ Use command line switches to load the LAN drivers on the conflicting segments with the opposite canonical order for the MAC addresses.
 - ◆ Use a different LAN card.
 - ◆ Because detected conflicts are echoed to the system console, resolve conflicts by manually reconfiguring the address of the offending workstation.
- ◆ If a router supporting IPX WAN client dial-in connections is configured to use node addresses that use the upper 34 bits of the node address, node address conflicts can occur. To avoid conflicts, the IPX WAN client router must have a registered network address for the WAN clients that is not translated, or it must use a low node address (less than 14 bits) on that network.

2 Planning

This section explains what decisions must be made before you can configure IPX™ beyond its most basic configuration.

IPX Configuration Decisions

How you configure IPX beyond its most basic configuration depends on the following decisions:

- ◆ **Whether to turn off IPX packet forwarding or to use the system as an IPX router**

The routing software enables you to turn off IPX packet forwarding on a NetWare® file server. This is useful when you have more than one NetWare system connecting two or more LANs and you want to free one of the systems from the task of forwarding—that is, *routing*—IPX packets between the LANs.

To turn off IPX packet forwarding, refer to [“Turning Off IPX Packet Forwarding” on page 51](#).

- ◆ **If you have WAN connections, whether to configure any of the following:**

- ◆ Static routes and services
- ◆ Watchdog packet spoofing
- ◆ Routed or static on-demand calls
- ◆ IPX and NetWare Core Protocol™ (NCP™) header compression

How you configure a WAN connection for IPX depends on how you want to use the connection and whether you use IPX with PPP, X.25, frame relay, or asynchronous transfer mode (ATM).

NOTE: Before you can configure IPX to run over a WAN connection, you must configure the WAN board, a network interface, and at least one WAN call destination.

To configure static routes and services for a permanent or on-demand call, refer to [“Configuring Static Routes and Services” on page 53](#).

To configure watchdog packet spoofing, refer to [“Configuring Watchdog Spoofing” on page 65](#).

To configure routed or static on-demand calls, refer to [“Configuring Routed or Static On-Demand Calls” on page 67](#).

To configure header compression, refer to [“Configuring IPX and NCP Header Compression” on page 69](#).

◆ **Whether to use NLSP, RIP/SAP, or both**

NetWare Link Services Protocol™ (NLSP™) software is the Novell link state routing protocol for IPX internetworks; Routing Information Protocol (RIP) and Service Advertising Protocol (SAP) are the traditional NetWare routing and service advertising protocols.

To configure NLSP on your router, refer to [“Configuring NLSP” on page 72](#).

To configure RIP and SAP on your router, refer to [“Configuring RIP and SAP” on page 77](#).

◆ **Whether to configure file server proxying on a dedicated router**

File server proxying is useful when you have a dedicated router—a PC running Novell Internet Access Server 4.1 over a two-user version of NetWare 4.11—and several NetWare workstations operating on the same network. Proxying enables the dedicated router to reply to workstations' Get Nearest Server requests with the name of a NetWare file server instead of its own. This enables the server, which has multiple connection slots, to handle simultaneous NCP connection requests from the workstations.

To configure file server proxying, refer to [“Proxying a NetWare File Server” on page 84](#).

- ◆ **Whether to use the IPX Address Mapping Gateway**

The IPX Address Mapping Gateway provides the following advantages:

- ◆ You can connect to a backbone network even when your local network numbers are not compatible with the backbone addressing scheme.
- ◆ If the routing protocol in the backbone does not support route aggregation, like most implementations of NLSP, the routing protocol probably cannot manage the number of network addresses from every user. Even if the routing protocol could handle route aggregation, network numbers might be assigned in a way that does not lend itself to aggregation. The IPX Address Mapping Gateway enables the summarization of routes in a manner transparent to the routing protocol by mapping many network numbers to a single number outside the local network.

To use the IPX Address Mapping Gateway, refer to [“Configuring the IPX Address Mapping Gateway” on page 86](#).

- ◆ **Whether to use IPX Route Aggregation**

IPX Route Aggregation allows your router to compactly report many IPX networks to a connecting backbone network. IPX Route Aggregation is most useful when several RIP networks are attached to an NLSP backbone network. Information in the backbone network is minimized by having the routers that connect to RIP networks report address summaries for these networks.

To use IPX Route Aggregation, refer to [“Configuring IPX Route Aggregation” on page 88](#).

- ◆ **Whether to change how your router propagates type 20 packets**

Type 20 is an IPX packet type that refers to any propagated packet. NetBIOS packets, for example, are type 20 packets. If your router processes a large number of type 20 packets, you can control how it propagates these packets through its interfaces. This can reduce the amount of traffic on a LAN.

To change how your router propagates type 20 packets, refer to [“Controlling the Propagation of Type 20 Packets” on page 89](#).

- ◆ **Whether to change the hop count limit of outbound IPX packets**

This enables you to control the range of outbound IPX packets on your router.

To control the range of IPX packets, refer to [“Changing the Hop Count Limit for IPX Packets” on page 90.](#)

- ◆ **Whether to balance traffic loads over equal-cost routes**

If your router has two or more network interfaces that can reach the same destination, it can distribute outbound traffic among those interfaces for an effective increase in throughput. This is called *load balancing* .

To configure load balancing over equal-cost routes, refer to [“Balancing Traffic Loads over Equal-Cost Routes” on page 92.](#)

- ◆ **Whether to configure SPX connection parameters**

Some NetWare applications have specific requirements for Sequenced Packet Exchange™ (SPX™) connection timeouts, retry counts, and so on. If any of these applications are used on your network, you might need to configure certain SPX parameters to enable these applications to run properly.

To configure any of these parameters, refer to [“Configuring SPX Connection Parameters” on page 94.](#)

- ◆ **Whether to change the delay and throughput values on your router**

This section explains how to set the delay and throughput values on a router to avoid connection timeouts over a slow link. This is often necessary for routers on LANs or bridged network segments that are separated by slow links.

To change the delay and throughput values, refer to [“Setting Delay and Throughput for a Slow Link” on page 95.](#)

NetWare Mobile IPX Configuration Decisions

Each of the following is key to the success of NetWare Mobile IPX™ configuration:

- ◆ **Mobile Client Driver Selection**
- ◆ **Planning for Efficient Use of Your Mobile Client**
- ◆ **Deciding Where to Locate a Home Router**

Mobile Client Driver Selection

When selecting your driver for the mobile client, verify with the vendor that the driver is written specifically with mobile operations in mind and that it supports the following:

- ◆ PCMCIA card in/card out capability
- ◆ In-range and out-of-range capability
- ◆ The NetWare Event Service Layer (NESL)

Planning for Efficient Use of Your Mobile Client

Planning ahead and knowing the appropriate ways to use NetWare Mobile IPX will help you to use your mobile client efficiently. We recommend that you do the following:

- ◆ Disable background products such as E-mail that poll the network.
- ◆ Use only data from the network. Keep your executable files on the mobile client.
- ◆ Complete operations such as saving files before removing the PCMCIA card.

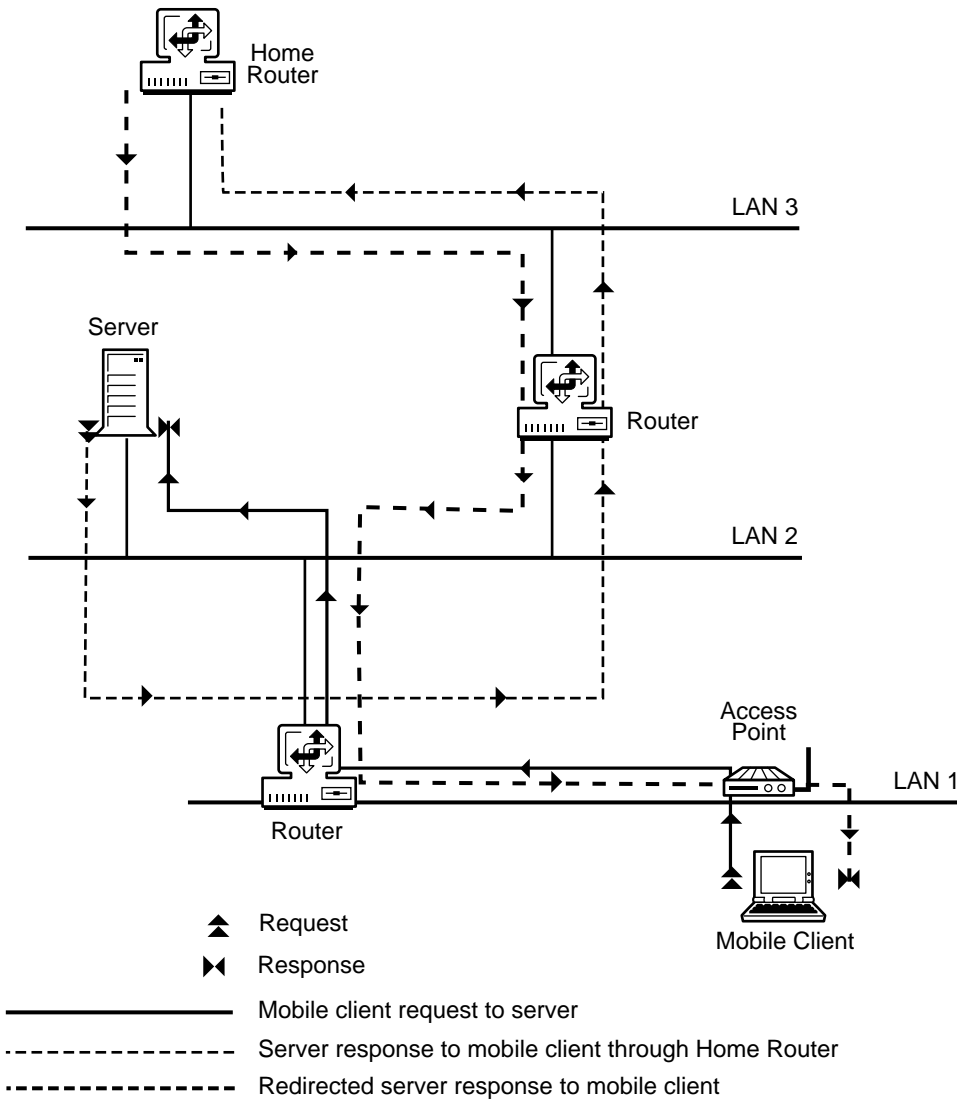
Deciding Where to Locate a Home Router

This section helps you choose the best location on your network to configure the Home Router software. It also explains why more than one Home Router can provide more efficient network operation in certain environments.

A Home Router forwards every packet destined for the mobile clients it serves. If the Home Router is located far from both the file server and the mobile client, and if the mobile client is close to the file server, a packet

destined for the mobile client travels more hops than necessary before arriving at the destination, as shown in **Figure 7**.


Figure 7 Inferior Home Router Placement



The request from the mobile client takes the shortest route to the server. The response from the server is first forwarded to the Home Router, because the destination address is the internal network configured for the server on which the Home Router resides. The router then patches the response with the mobile client's actual address and forwards the packet to the mobile client. In this topology, the response takes an inferior path to the mobile client because of the extra hops taken.


In general, you should install the Home Router in the middle of the network so that most clients are only a few hops away, as shown in [Figure 8](#). The Home Router should be located somewhere on the path between the file server and the mobile clients.

Figure 8 Better Home Router Positioning



The best place to install the NetWare Mobile IPX Home Router software is on the file server that the mobile clients use most, as shown in [Figure 9](#). In this way, when mobile clients access the file server, the responses from the file server are patched with the mobile client's actual location before they ever leave the server. Therefore, the responses do not travel an extra hop before reaching the client.

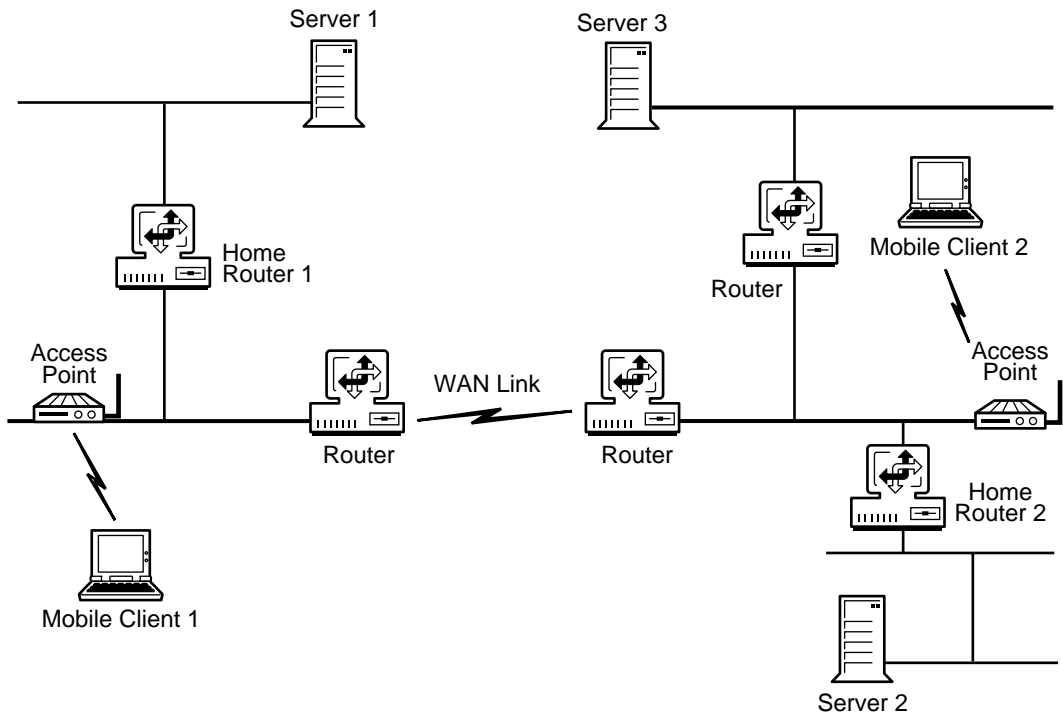
Figure 9 Best Home Router Positioning



HINT: We recommend having a Home Router in each operating area. For example, in a large corporation, you should have a Home Router in marketing, finance, manufacturing, and so on. This enables mobile client users to connect to a preferred Home Router.

If the server and mobile clients are located on one side of a WAN link and the Home Router is on the other side of the link, costly WAN bandwidth is used unnecessarily. If both networks on each side of a WAN link require NetWare Mobile IPX, you should have two Home Routers—one on each side of the WAN link. Mobile clients use the closest Home Router, as shown in [Figure 10](#).

Figure 10 Home Router Positioning over WAN Links



IMPORTANT: When a mobile client is transferred between the two sites over the WAN and not restarted, the client still uses the original Home Router for communications—crossing the WAN if necessary—until the mobile client is restarted.

3

Setting Up

Novell® Internet Access Server 4.1 routing software provides a set of configurable parameters with which you can modify operational characteristics of the Internetwork Packet Exchange™ (IPX™) network protocol. You configure all IPX parameters from the Novell Internet Access Server Configuration utility (NIASCFG).

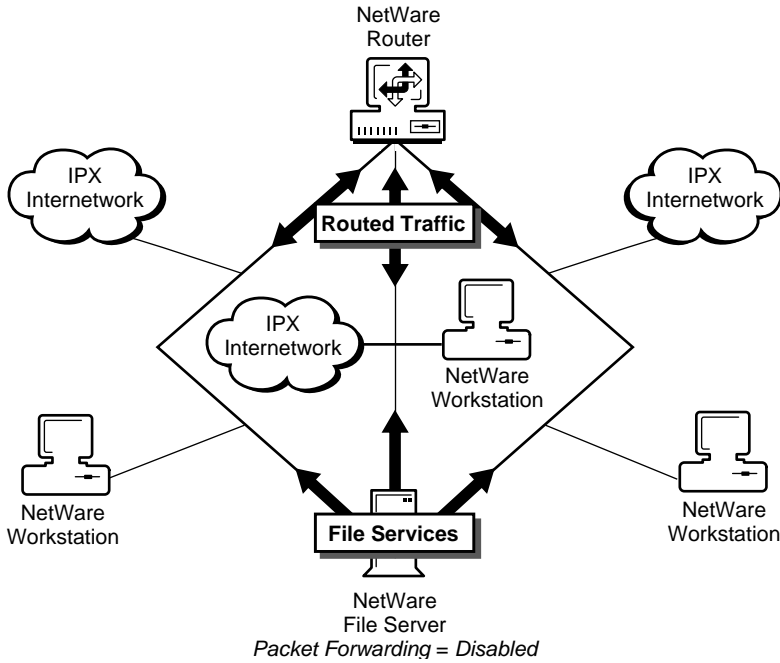
Turning Off IPX Packet Forwarding

As a typical part of its operation, a NetWare file server forwards (routes) IPX packets between its network interfaces. By disabling the Packet Forwarding parameter, you turn off IPX packet forwarding on a NetWare file server. This is useful when you have more than one NetWare system on a network and you want to make more computing resources available for file and print services.

A server with IPX packet forwarding disabled still operates as a file server, but broadcasts only its own services and internal network number—not those associated with its network interfaces. A server operating in this way is sometimes called a *multihomed* server. Although a multihomed server does not reply to routing requests from NetWare workstations, it still accepts incoming RIP and SAP broadcasts to maintain its bindery.

Figure 11 on page 52 shows how the tasks of routing and file service can be divided between a NetWare file server and a dedicated router on the same network. Typically, the task of routing IPX traffic between the two internetworks is shared by the router and the file server. When you disable IPX packet forwarding on the file server, the dedicated router assumes the task of routing all IPX traffic. The file server, now free from the tasks of routing IPX packets and answering route requests, can devote its full processing resources to file service.

Figure 11 Turning Off IPX Packet Forwarding on a NetWare File Server



You should turn off IPX packet forwarding if you do not want the server to forward IPX packets between its interfaces.

For more information about IPX routing and related topics, refer to [Chapter 1, “Understanding,” on page 11](#). For more information about NetWare file and print services, refer to your NetWare documentation.

How to Turn Off IPX Packet Forwarding

Before you begin, you must have at least two NetWare servers, or one NetWare server and a dedicated router on the same IPX network.

To turn off IPX packet forwarding, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Protocols > IPX
- 2** The Packet Forwarding parameter is already highlighted; select it, then select Disabled .

- 3 Press Esc and save your change.
- 4 Press Esc to return to the Internetworking Configuration menu.
- 5 If you want this change to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring Static Routes and Services

A *static route* is a RIP route that is added to a router's Routing Information Table by a network administrator rather than by the active routing protocol—in this case, RIP—operating over a network link. With Novell Internet Access Server 4.1, you can configure static routes and services for both permanent and on-demand calls.

A *permanent call* is a connection that remains active between the local router and the remote router identified by the call destination. An *on-demand call* is a point-to-point connection between two routers that becomes active only when one router must send user data to the router at the other end. On-demand calls are well-suited for connections that use expensive telecommunications carriers and for slow links over which it is undesirable to send routing and service traffic. For more information about permanent and on-demand calls, refer to [“Call Types” on page 25](#).

You can configure static routes and services with the following utilities:

- ♦ **NIASCFG** —With NIASCFG, you must configure *all* routes and services manually. To activate the configuration on both routers, you select Reinitialize System from the Internetworking Configuration menu. Use NIASCFG to set up WAN call destinations at each end of the connection and configure static routes and services.
- ♦ **STATICON** —The static route and service configuration utility for IPX. STATICON uses the Simple Network Management Protocol (SNMP) to discover which routes and services are available through a remote router and adds them to the static routing table on a local router.

Before configuring static routes and services with STATICON, you must use NIASCFG to set up the WAN call destinations and activate the configuration by selecting Reinitialize System . You then load STATICON and configure the static routes and services on the routers. STATICON configures all routes and services on each router

automatically and allows you to try the static configuration before saving it to disk. The STATICON configuration becomes active immediately; you do not need to reinitialize or restart the router.

For more information about static routes and services and related topics, refer to [“Static Routes and Services” on page 28](#).

This section contains the following topics:

- ♦ [“Configuring Static Routes and Services with NIASCFG” on page 54](#)
- ♦ [“Configuring Static Routes and Services with STATICON” on page 56](#)

Configuring Static Routes and Services with NIASCFG

NOTE: If you plan to configure static routing information, we recommend using STATICON to avoid potential routing loops.

Before you configure static routes and services with NIASCFG, determine the addresses of the networks or hosts that you want to access. Then set up a WAN call destination, on *each* router, to the router at the other end of the connection.

WARNING: When setting up a call destination, be sure you set the Remote System ID parameter to the server name of the system initiating the inbound call. The local system checks each inbound call against the remote system ID.

Configuring Static Routes

To configure static routes with NIASCFG, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > WAN Call Destinations.

- 2** Select a WAN call destination from the list, or press **Ins** to choose from a list of available call destinations.
- 3** Select Static Routes.

A new screen displays any configured static routes.

- 4** Press **Ins**, then enter the following information:

- ♦ Network Number —Network number that must be accessed for this on-demand call.
- ♦ Hops to Network —Number of routers crossed to reach the specified network number. If the network number is in the internal network of the remote router, specify 1.

- ◆ Ticks to Network —Number of ticks used to allow a packet to reach the destination network. Add one tick to this value for each LAN hop.

5 Press `Esc` and save your changes.

6 Press `Esc` to return to the Internetworking Configuration menu.

7 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring Static Services

HINT: If you plan to configure any uncommon services, we recommend using `STATICON` to avoid errors.

To configure static services with `NIASCFG`, complete the following steps:

1 Load `NIASCFG`, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > WAN Call Destinations.

2 Select a WAN call destination from the list, or press `Ins` to choose from a list of available call destinations.

3 Select Static Services .

A new screen displays any configured static services.

4 Press `Ins`, then enter the following information:

- ◆ Service Name —Name of the service to be accessed through the on-demand call. This name, which is typically the server name, is added to the local service and routing tables.

NOTE: If you need to obtain service name information, use the `IPXCON` utility at the remote router. If you use `STATICON` to configure static services, this information is obtained automatically.

- ◆ Service Type —SAP service type for this service, expressed as a hexadecimal number. This is typically the file server type (0004). Press `Ins` to display a list of service types.
- ◆ Service Address Network —IPX network number of the service. If you are specifying a file service or a service on a server or router, enter the internal network number of that server or router.

- ◆ Service Address Node —Node address where the service resides. If you are specifying a file service or a service on a server or router, enter **1** . If you are specifying a NetWare 2 file server, specify the NIC address of LAN board A.
- ◆ Service Address Socket —Socket number on which this service listens for service requests. If you are specifying a file service, enter **0x0451** .

NOTE: If you need to obtain service address socket information, use the IPXCON utility at the remote router. If you use STATICON to configure static services, this information is obtained automatically.
- ◆ Hops to Service —Number of routes crossed to reach the service.
- ◆ Ticks to Service —Number of ticks needed for a packet to reach the destination network. Add one tick to this value for each LAN hop.

5 Press Esc and save your changes.

6 Press Esc to return to the Internetworking Configuration menu.

7 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring Static Routes and Services with STATICON

STATICON provides a fast and error-free way to configure static routes and services on routers at each end of a point-to-point WAN connection. However, before loading STATICON, you must use NIASCFG to complete the following preparatory tasks on each router:

- ◆ Set up a WAN call destination to the router at the other end of the connection.
- ◆ Select Reinitialize System to activate the NIASCFG configuration.

IMPORTANT: For STATICON to operate, the router at the other end of the connection must either be running Novell Internet Access Server 4.1 or be a third-party router that supports IPX SNMP and the IPX Management Information Base (MIB) variables.

When you load STATICON, the Select Configuration Method menu appears. [Table 3 on page 57](#) describes each menu option.

Table 3 Select Configuration Method Menu Options

Menu Option	Description
Dynamically Configure Static Routing Tables	Opens an on-demand call to the remote router and allows you to edit the local and remote configurations using dynamically obtained selection lists.
Configure Services from Gatekeepers	Opens an on-demand call to obtain a list of services available from a gatekeeper and allows you to select the services you want. The gatekeeper is normally on the other side of a WAN link.
Configure Local Static Services	Lets you configure static services <i>manually</i> on the local router. This configuration is nearly identical to the one presented in “Configuring Static Services” on page 55 .
Configure Local Static Routes	Lets you configure static routes <i>manually</i> on the local router. This configuration is nearly identical to the one presented in “Configuring Static Routes” on page 54 .
Write Static Routing Tables to Permanent Storage	Sends IPX SNMP requests to the local router to put the routing table information into permanent storage. The router is polled to ensure that the information is updated.
Restore Static Routing Tables from Permanent Storage	Sends IPX SNMP requests to the local router to restore routing table information from permanent storage. The router is polled to ensure that the information is updated.

The most efficient way to configure static routes and services is to select the Dynamically Configure Static Routing Tables option. This enables you to select from the following options:

- ◆ Autoconfigure Local and Remote Routing Tables —Exchanges all routing and service table information automatically with the remote router. Select this option if you want an on-demand call to obtain full routing and service information in the static routing tables. A status screen shows the progress of the exchange. This exchange might take significant time to complete if you are working over a slow link or on a large network.
- ◆ Configure Local Routing Tables —Selectively configures routing and service tables for the local router from information learned from the remote router through IPX SNMP requests.
- ◆ Configure Remote Routing Tables —Selectively configures the routing and service tables for the remote router from information learned from the local router through IPX SNMP requests.

- ◆ Write Connection Routing Tables to Permanent Storage —Sends IPX SNMP requests to the local and remote routers to save the current local and remote routing tables for this connection to permanent storage. Each router is polled to make sure the operation is completed.
- ◆ Restore Connection Routing Tables from Permanent Storage —Sends IPX SNMP requests to the local and remote routers to restore the local and remote static routing tables for this connection from permanent storage. Each router is polled to make sure the operation is completed.

For more information about configuring static routes and services dynamically, refer to “[Selectively Configuring Static Routes and Services](#)” on page 59 and “[Automatically Configuring Static Routes and Services](#)” on page 62.

Checking Write Access on the Remote Router

For STATICON to configure a remote router's routing and service tables, it must support IPX SNMP and the IPX MIB variables and have SNMP-write access to the router. If the remote router is running Novell Internet Access Server 4.1, specifying a *Control Community* from NIASCFG enables write access.

To check write access, complete the following steps:

- 1** Load NIASCFG on the remote router, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Manage Configuration > Configure SNMP Parameters

The Control State field should read Any Community May Write or Specified Community May Write . If it reads Specified Community May Write , note the name in the Control Community field. Use this name when you must provide the name of the SNMP control community in [Step 8 on page 59](#) of [Selectively Configuring Static Routes and Services](#) and [Step 8 on page 63](#) of [Automatically Configuring Static Routes and Services](#).

- 2** Press Esc to return to the Internetworking Configuration menu.
- 3** Exit NIASCFG.

Selectively Configuring Static Routes and Services

Through selective configuration, you can choose specific routes and services you want to add to a routing table. This feature lets you select from an existing routing table the routes and services that your router does not have in its table. You can copy routes to a remote router from a local router, or copy routes to a local router from a remote router.

To selectively configure static routes and services, complete the following steps:

1 Load STATICON.

2 Select Dynamically Configure Static Routing Tables .

STATICON displays the on-demand calls of which IPX is currently aware. It also shows the connection state of each call.

NOTE: The Auto Static Route listed in the display is the nonconfigured static route to the internal network on the other side of the WAN link. This automatic static route ensures a route across the link in case normal RIP filtering might prevent such a route. It must not be deleted.

3 Use the arrow keys on your keyboard to select the WAN call destination associated with the remote router.

4 If the Status field associated with the call reads Not Connected , press Ins to connect the call.

5 Wait for the Status field to change to Connected .

This can take several seconds if you are using a dial-up line.

6 Press Enter .

7 Do one of the following:

To configure the routing table for the local router, select Configure Local Routing Tables.

To configure the routing table for the remote router, select Configure Remote Routing Tables.

8 Enter the SNMP Control Community name associated with the remote router, or press Enter to accept the default Control Community named public .

The remote router must have write access enabled. If you need to check write access, refer to [“Checking Write Access on the Remote Router” on page 58](#).

- 9** If you selected **Configure Local Routing Tables** , complete the following steps; otherwise, continue with **Step 10 on page 60**.

A progress screen appears while the local system reads the currently configured routes and services.

After the **Locally Configured Routes** screen appears, you can remove items from the routing table by highlighting or marking the entries and pressing **Del** .

- 9a** Press **Ins** to add static routes or services to the local routing table.

A progress screen appears while the local system gathers information from the remote router.

The **Selectable Routes and Services** screen appears. The routes and services listed here are from the remote router's table. The list shows only the routes and services that are not already present in the local router's table.

- 9b** Mark the routes or services you want to add to the local routing table.

Use the following keys to mark your selections:

- ◆ **F5** —Marks the current entry.
- ◆ **Tab** —Marks all entries that have the same network number as the currently highlighted entry.
- ◆ **F6** —Lets you use wildcard characters (***** and **?**) to select entries.

If you use **F6** , the **Select Wild Card Marking Option** screen appears. You can select **Match Service Names** or **Match Network Numbers** . After you make a selection, the **Enter Pattern for Matching** screen appears, enabling you to enter the name or number pattern and wildcard.

- 9c** Press **Enter** .

The **Select Currently Marked Routes and/or Services?** screen appears.

- 9d** Select **Yes**. Proceed to **Step 11 on page 61**.

- 10** If you selected **Configure Remote Routing Tables** , complete the following steps:

A progress screen appears while the local system gathers information from the remote router.

After the Remote Router's Configured Routes and Services screen appears, you can remove items from the routing table by highlighting or marking the entries and pressing Del .

10a Press Ins to add static routes or services to the remote routing table.

A progress screen appears while the local system reads the currently configured routes and services.

The Selectable Routes and Services screen appears. The routes and services listed here are from the local router's table. They show only the routes and services that are not already present on the remote router's table.

10b Mark the routes or services you want to add to the remote routing table.

Use the following keys to mark your selections:

- ◆ F5 —Marks the current entry.
- ◆ Tab —Marks all entries that have the same network number as the currently highlighted entry.
- ◆ F6 —Lets you use wildcard characters (* and ?) to select entries.

If you use F6 , the Select Wild Card Marking Option screen appears. You can select Match Service Names or Match Network Numbers . After you make a selection, the Enter Pattern for Matching screen appears, enabling you to enter the name or number pattern and wildcard.

10c Press Enter .

The Select Currently Marked Routes and/or Services? screen appears.

10d Select Yes .

11 Press Esc twice.

STATICON allows you to choose whether to save the static configuration to disk now or test the configuration first by trying to establish a connection.

If you want to save the configuration to disk now, continue with [Step 12 on page 62](#).

If you want to test the configuration before saving it to disk, select Do Not Save the Routing Tables to Permanent Storage.

In this case, the configuration remains in router memory.

If you decide later to save the configuration to disk, return to the Select Configuration Option For This Call screen and select Write Connection Routing Tables to Permanent Storage .

The configuration is saved to disk if you did not do any of the following while testing the configuration:

- ◆ Restart the router
- ◆ Delete the WAN call destination from NIASCFG
- ◆ Unload IPXRTR

12 Select Save the Routing Tables to Permanent Storage .

NOTE: Changes you make from STATICON take effect immediately; you do not need to reinitialize or restart either router after completing the configuration.

The following message appears after the configuration is saved to disk:

```
Writing static routing tables for this call to permanent
  storage completed successfully in router <router_name
  >.
```

```
<Press ENTER to continue>
```

13 Press Enter , then press Esc until you return to the Select Configuration Method menu.

14 Exit STATICON.

Automatically Configuring Static Routes and Services

Configuring static routes and services automatically lets you copy all the missing routes from your local router to a remote router and from the remote router to your local router at the same time.

To automatically configure static routes and services, complete the following steps:

1 Load STATICON.

2 Select Dynamically Configure Static Routing Tables .

STATICON displays the on-demand calls of which IPX is currently aware. It also shows the connection state of each call.

3 Use the arrow keys on your keyboard to select the WAN call destination associated with the remote router.

4 If the Status field associated with the call reads Not Connected , press Ins to connect the call.

5 Wait for the Status field to change to Connected .

This can take several seconds if you are using a dial-up line.

6 Press Enter .

7 Select Autoconfigure Local and Remote Routing Tables .

8 Enter the SNMP Control Community name associated with the remote router, or press Enter to accept the default Control Community public .

The remote router must have write access enabled. If you need to check write access, refer to [“Checking Write Access on the Remote Router.”](#)

A progress screen appears as STATICON exchanges routes and services with the remote router. This might take several minutes if you are working over a large network or slow link.

The following message appears when the exchange is complete:

```
Autoconfiguration of Routing Tables between local and
remote Routers completed successfully.
```

```
<Press ENTER to continue>
```

9 Press Enter .

STATICON allows you to choose whether to save the static configuration to disk now or try the configuration first.

If you want to save the configuration to disk now, continue with [Step 10 on page 64](#) .

If you want to try the configuration before saving it to disk, select Do Not Save the Routing Tables to Permanent Storage.

In this case, the configuration remains in router memory.

If you decide later to save the configuration to disk, return to the Select Configuration Option For This Call screen and select Write Connection Routing Tables to Permanent Storage .

You can do this as long as you do not do any of the following while trying out the configuration:

- ◆ Restart the router
- ◆ Delete the WAN call destination from NIASCFG
- ◆ Unload IPXRTR

10 Select Save the Routing Tables to Permanent Storage .

NOTE: Changes you make from STATICON take effect immediately; you do not need to reinitialize or restart either router after completing the configuration.

The following message appears after the configuration is saved to disk:

```
Writing static routing tables for this call to permanent
storage completed successfully.
```

```
<Press ENTER to continue>
```

11 Press Esc until you return to the Select Configuration Method menu.

12 Exit STATICON.

Configuring Services for a Gatekeeper

A gatekeeper is a special file server that is located in the hub of an internetwork and is able to see all public services from all connected sites. A gatekeeper stores routing and services information in its bindery, Novell's equivalent of a telephone book. Using STATICON, you can obtain a list of services available from a gatekeeper and select the required services for advertisement by your local router. For more information about gatekeepers and binderies, refer to **“Static Routes and Services” on page 28**.

To configure services for a gatekeeper, complete the following steps:

1 Load STATICON.

2 Select Configure Services for Gatekeepers .

STATICON displays the IPX Calls screen, a list of on-demand calls of which IPX is currently aware. It also shows the connection state of each call.

3 Use the arrow keys on your keyboard to select the WAN call destination associated with the remote router.

4 If the Status field associated with the call reads Not Connected , press Ins to connect the call.

5 Wait for the Status field to change to Connected .

This can take several seconds if you are using a dial-up line.

6 Press Enter .

The Select Configuration Option For This Call screen displays.

7 Select Configure Local Routes and Services .

The Configured Services screen displays.

- 8** Press **Ins** to display a list of gatekeepers.
- 9** Use the arrow keys to select a gatekeeper, then press **Enter** .

STATICON displays messages that it is attaching to the selected gatekeeper and that it is scanning the gatekeeper for services. Service names are displayed as STATICON discovers them. When the scanning is completed, a list is displayed of all the services available from the gatekeeper.
- 10** Select all the services you want by using the arrow keys to highlight each desired service, then pressing **F5** .
- 11** Press **Enter** to confirm your selection of all the marked services.
- 12** Optionally, press **Esc** to return to the IPX Calls screen and then repeat **Step 3 on page 64** through **Step 11** for another WAN call.
- 13** Save the information as follows:

If you have collected gatekeeper information for a single WAN call, press **Esc to return to the Select Configuration Option For This Call screen, then select Write Static Routing Tables to Permanent Storage.**

If you have collected gatekeeper information for multiple WAN calls, press **Esc to return to the STATICON main menu, then select Write Static Routing Tables to Permanent Storage.**
- 14** If necessary, press **Esc** to return to the Select Configuration Method menu. Exit STATICON.

Configuring Watchdog Spoofing

When several workstations are operating over an on-demand call, the frequent exchange of watchdog packets can keep the connection active most of the time. Depending on the telecommunications carrier you use for the connection, this can become expensive.

You can avoid this by configuring your router to spoof the watchdog packets. This means that the router captures watchdog query packets on their way to a workstation and responds on the workstation's behalf without activating the on-demand call.

Note, however, that because of spoofing, the workstation's server connection remains occupied unless the workstation logs out. To avoid this, have the remote server execute a forced logout of all workstations at a predetermined time so that all server connections are free for the next day.

For more information about watchdog spoofing and related topics, refer to [“Watchdog Packet Spoofing” on page 30](#).

This section contains the following topics:

- ◆ [How to Configure Watchdog Spoofing on an Interface](#)
- ◆ [How to Configure Watchdog Spoofing for Call Destinations](#)

How to Configure Watchdog Spoofing on an Interface

By default, watchdog spoofing is enabled for all on-demand WAN connections. If you want to *disable* watchdog spoofing on a WAN interface, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > Expert Bind Options
- 2** Select On Demand Spoofing , press Enter , then select Disabled .
- 3** Press Esc and save your change.
- 4** Press Esc to return to the Internetworking Configuration menu.
- 5** If you want this change to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

How to Configure Watchdog Spoofing for Call Destinations

By default, watchdog spoofing is enabled for all on-demand WAN connections. To configure watchdog spoofing for a particular on-demand WAN call destination, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > WAN Call Destinations
- 2** Select a call destination.

If you are modifying an on-demand call that has already been configured, select one from the list.

If you are configuring a new on-demand call, press Ins and choose a call from the list of available calls.

3 Select Expert Options .

4 Select On Demand Spoofing .

The default state is Use Default . This means the call uses the spoofing state to which the interface is currently set.

If spoofing is enabled on the interface but you want to disable it only for this call, select Disabled.

If spoofing is disabled on the interface but you want to enable it only for this call, select Enabled.

5 Press Esc and save your changes.

6 Press Esc to return to the Internetworking Configuration menu.

7 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring Routed or Static On-Demand Calls

Novell Internet Access Server 4.1 enables you to configure a *routed on-demand call* or *static on-demand call* for each WAN call destination.

Unlike the standard on-demand call, which relies on statically configured routes and services at each end of a point-to-point connection, a routed on-demand call runs a routing protocol while the link is active. When the link goes down, the routes and services made known by the routing protocol become unavailable.

Typically, a Data-Link layer timer triggers the termination of an on-demand call after no data has crossed the link for some period of time. Because a routing protocol running over a routed on-demand call would reset this timer each time a protocol packet is sent or received, it would keep a link active because of the protocol data flowing through. To solve this problem, Novell Internet Access Server 4.1 uses a timer that operates at the Network layer. This timer is reset only when data packets—not protocol packets—cross the link. In this way, the routing updates do not keep an on-demand link active when no data is being transmitted.

NOTE: A minimal (seed) set of static routes and services must be associated with a routed on-demand call for key server access. Unless the call is known to get

somewhere, the link will not come up. After the link comes up, other services and routes can be accessed.

For more information about routed and static on-demand calls and related topics, refer to [“Call Types” on page 25](#).

How to Configure Routed or Static On-Demand Calls

Before you begin, you must complete the following tasks:

- ◆ Configure at least one on-demand WAN call destination.
- ◆ Configure a minimal set of routes and services for key server access.

To configure a routed or static on-demand call, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings

2 Select an on-demand call.

If you are modifying an on-demand call that has already been configured, select one from the list.

If you are configuring a new on-demand call, press *Ins* **and choose a call from the list of available calls.**

The WAN Call Destination Entry screen is displayed.

3 Select WAN Call Destinations .

4 Select a WAN call destination from the list of configured calls, then do one of the following:

To configure a routed on-demand call, select Routed On Demand *from the pop-up menu, then proceed to* [Step 5](#).

To configure a static on-demand call, select Static On Demand *from the pop-up menu, then proceed to* [Step 6](#).

By default, on-demand calls are static, and routing traffic over an on-demand call is disabled.

5 For a routed on-demand call only, do the following:

5a Select RIP Bind Options.

5b Configure the routing protocol you want to run over the call.

If you want to run RIP/SAP:

- ◆ Select RIP Options .
- ◆ Set RIP State Override to On .
- ◆ Configure the other RIP override parameters as necessary.
- ◆ Press Esc .
- ◆ Select SAP Options .
- ◆ Set SAP State Override to On .
- ◆ Configure the other SAP override parameters as necessary.

If you want to run NLSP:

- ◆ Select NLSP Options .
- ◆ Set NLSP State Override to On .
- ◆ Configure the other NLSP override parameters as necessary.

6 Press Esc and save your changes.

7 Press Esc to return to the Internetworking Configuration menu.

8 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring IPX and NCP Header Compression

Header compression increases the throughput of IPX and NCP packets over low-speed serial lines. An IPX packet header is 30 bytes and is typically followed by an upper-layer protocol header, such as an SPX header. Header compression reduces the size of this combined packet header to just a few bytes.

Header compression is negotiated by the IPXWAN™ protocol when a call is established over any WAN connection. Header compression is not used on the connection if IPXWAN detects that one of the end nodes does not support it.

When you enable header compression, you can also specify the number of *compression slots* . A compression slot is a location in router memory that stores packet header information. The compression algorithm uses this information to compress outgoing—and decompress incoming—packet headers.

By default, the number of allocated compression slots is 16. In general, a session between two end points uses one slot; routing information uses one or two. Each slot can contain an IPX or an NCP header. When no more slots are available, packet headers are sent uncompressed, or old slots are reused.

For more information about IPX and NCP header compression and related topics, refer to [“Header Compression” on page 31](#).

IMPORTANT: To use header compression, the routers at each end of the connection must have header compression enabled and must allocate the same number of *header compression slots* . If the number of compression slots is different on each router, IPXWAN selects the lesser of the two.

This topic contains the following sections:

- ♦ [“How to Configure IPX and NCP Header Compression on an Interface” on page 70](#)
- ♦ [“How to Configure IPX and NCP Header Compression per Call Destination” on page 71](#)

How to Configure IPX and NCP Header Compression on an Interface

To configure IPX and NCP header compression on a WAN interface, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > Expert Bind Options

2 Select Header Compression .

This parameter enables or disables header compression for all IPX packets sent through this interface. By default, header compression is enabled on all WAN interfaces; if you want to disable it on the interface, select Disabled .

3 Select Compression Slots , then enter the number of slots you want to allocate to this interface.

The more concurrent IPX sessions you use over the interface, the more compression slots you should allocate.

IMPORTANT: Be careful not to allocate too many compression slots. Memory is required to store the headers, and the compression algorithm must scan through stored headers to find a match for each transmitted packet. An excessive number of slots results in a higher processing load and slower performance.

- 4** Press `Esc` and save your changes.
- 5** Press `Esc` to return to the Internetworking Configuration menu.
- 6** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

How to Configure IPX and NCP Header Compression per Call Destination

By default, header compression is enabled for all WAN connections. To configure header compression for a particular WAN call destination, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > WAN Call Destinations
- 2** Select a call destination.
If you are modifying a call that has already been configured, select one from the list.
If you are configuring a new call, press `Ins` and choose a call from the list of available calls.
- 3** Select Expert Options .
- 4** Select Header Compression .

The default state is Use Default . This means the call uses the compression state to which the interface is currently set.

If compression is enabled on the interface but you want to disable it only for this call, select Disabled.

If compression is disabled on the interface but you want to enable it only for this call, select Enabled.

- 5 Select Compression Slots , then enter the number of slots you want to allocate to this call.

IMPORTANT: Be careful not to allocate too many compression slots. Memory is required to store the headers, and the compression algorithm must scan through stored headers to find a match for each transmitted packet. An excessive number of slots results in a higher processing load and slower performance.

- 6 Press Esc and save your changes.
- 7 Press Esc to return to the Internetworking Configuration menu.
- 8 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring NLSP

Novell developed NLSP to meet the demands of large IPX internetworks. As a *link state* routing protocol, NLSP offers better performance, reliability, and scalability than the IPX RIP routing traditionally employed by NetWare servers.

Unlike RIP and SAP, which periodically broadcast routing and service information respectively, NLSP transmits routing information only when a change occurs in a route or service somewhere in the network, *or every two hours*—whichever occurs first. Because NLSP generates fewer routing updates than RIP and SAP, it uses less network bandwidth to maintain its routing database.

To transmit information about its directly connected routers and the links to those routers, an NLSP router uses *Link State Packets* (LSPs). By default, LSPs are 512 bytes, a nominal value that is sufficient for most IPX networks. If you have a large network—on the order of 4,000 routes and 2,000 services or more—you should increase the value of the LSP Size parameter to 1024. To configure this parameter, refer to “[How to Change the LSP Size](#)” on page 76.

By default, NLSP *broadcasts* its packets because some LAN drivers do not properly support *multicast* , a transmission mode that enables only those devices listening for a specific multicast packet address to accept the packet. You can, however, change the NLSP packet transmission mode to multicast with the MAC Channel parameter. An advantage of using multicast

transmission is that NLSP packets sent by multicast do not clutter nonrouting nodes with unnecessary traffic.

NOTE: All NetWare systems on the same LAN must use the same NLSP packet transmission mode.

NLSP makes large IPX internetworks more manageable by allowing you to partition them into administrative domains called *routing areas*. Each routing area can be identified by up to three *area addresses*, a unique, 4-byte hexadecimal number that identifies each NLSP router as being part of a routing area. Although area addresses are not required, they are available chiefly for compatibility with future versions of NLSP and do provide some benefit for large IPX internetworks.

WARNING: Do not configure area addresses unless you must partition a large IPX internetwork. If you make an error in the address assignments, you can partition your network inadvertently and lose connectivity between some routers.

Each NLSP router is identified by a unique, 6-byte hexadecimal number called the *system ID*. The default system ID comprises a 2-byte constant, 0x0200, followed by the router's own internal network number. You should not change the system ID unless you have another numbering scheme in place with which you can reliably track and manage the NLSP routers on your network. If you must change a router's system ID, use another unique number, such as the physical address of one of the router's network interface boards.

Using a default *cost* value based on media throughput, NLSP chooses the best route by which to forward IPX packets. Novell Internet Access Server 4.1 enables you to override this value on an interface. By overriding the default cost, you can establish preferred routes, balance traffic loads among interfaces, and set up specific traffic flows between routers. For more information, refer to [“Balancing Traffic Loads over Equal-Cost Routes” on page 92](#).

All NLSP routers have a configurable parameter called the *Priority*. The router with the highest priority becomes the *Designated Router*, which assumes the responsibility for exchanges of link state information on behalf of all other NLSP routers on the LAN. You do not typically need to change a router's *Priority* value; the NLSP routers automatically elect one from themselves. However, if you want to force a router to become the Designated Router for its LAN, increase its *Priority* value to at least 85. The router you choose should be typically up and should have enough memory to process NLSP routing information and generate the *pseudonode* LSP for its LAN. If you want to prevent a particular router from becoming the Designated Router, decrease its *Priority* value.

IMPORTANT: Novell Internet Access Server 4.1 provides a set of *convergence parameters* that enable you to customize the operation of NLSP on your router. The default values for these parameters are sufficient for most IPX networks and should be changed only on the advice of your technical support representative.

When configuring NLSP on an interface, you can set the NLSP State parameter to one of two states: On or Off. On enables the router to exchange NLSP packets freely with other NLSP routers on the attached network. Off disables NLSP routing on the interface.

This topic contains the following sections:

- ◆ “How to Configure NLSP” on page 74
- ◆ “How to Change the LSP Size” on page 76

How to Configure NLSP

If you implement NLSP throughout a large IPX internetwork, or even if you want to configure NLSP on just a few routers or servers, refer to NLSP Migration for information about planning your implementation of NLSP.

To configure NLSP on the router, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Protocols > IPX
- 2** Select the Routing Protocol parameter, then select NLSP with RIP/SAP Compatibility .
- 3** Press Esc to return to the Internetworking Configuration menu, then select the following path:
Select Bindings > a network interface > Expert Bind Options > NLSP Bind Options
- 4** Select NLSP State .

If you want to run NLSP over the interface, select On.

This enables NLSP routing on the interface.

If the interface is on an area boundary, or if you want to filter incoming RIP or SAP packets at the interface, select Off.

If the router supports two or more interfaces and you want to filter routes and services to a remote site through this interface:

- ◆ Select Off .
- ◆ Enable RIP and SAP on the interface.

To enable RIP and SAP, refer to [“Configuring RIP and SAP” on page 77](#).

5 Select MAC Channel , then select the NLSP packet transmission mode.

If you select Multicast , NLSP automatically determines the multicast address. All systems on a network must be set to Multicast ; otherwise, the systems default to Broadcast , the default state for this parameter.

IMPORTANT: Make sure the driver you are using supports multicast transmission; drivers that do not support multicast can cause systems to become unaware of each other.

6 If you want to customize the interface further, configure one or more of the following parameters:

IMPORTANT: Because the default settings for these parameters are suitable for most NLSP networks, you should change them only for a specific purpose. Misconfiguring these parameters can increase routing traffic or cause loss of connectivity on your network.

- ◆ MTU Override —Overrides the Maximum Transmission Unit (MTU) of the network medium to which this interface is connected. All outbound packets on this interface use the value you enter. The default value is 0, which means use the MTU of the network medium. For example, the Ethernet MTU is 1,500 bytes.

Configure this parameter if you have a bridge or other device on your network, or if you want to transmit smaller packets over a WAN.

- ◆ Priority —Sets the *priority* of the NLSP router on the network segment to which this interface is connected. The default priority is 64; increase this value to at least 85 if you want the router to become the Designated Router for its LAN; decrease it if you want to prevent the router from becoming the Designated Router.
- ◆ Cost Override —Overrides the default cost of the network medium to which this interface is connected. To configure this parameter, refer to [“Balancing Traffic Loads over Equal-Cost Routes” on page 92](#).
- ◆ Pace Override —Specifies the maximum number of NLSP packets that can be sent each second through this interface. On a LAN, the

default—and maximum—pace for NLSP packets is 30 pps; on a WAN, the NLSP pace is derived from the throughput of the link.

The default value for this parameter is 0, which means do not override the current pace.

- 7** Press **Esc** and save your changes.
- 8** Press **Esc** to return to the Internetworking Configuration menu.
- 9** If you want these changes to take effect immediately, select **Reinitialize System** .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

How to Change the LSP Size

The larger the packet a network can carry, the fewer LSPs are required to propagate an NLSP router's link state information on that network. However, the LSP used by the router must be no larger than the largest frame size supported by the network, less 30 bytes for the IPX header. For example, an ARCnet* LAN can transmit no more than 576 bytes at a time. If you leave **LSP Size** at the default value of 1024, the LSP cannot be transmitted across the LAN because it is too large. The result is that the network is prevented from converging.

Before you begin, you should know the maximum frame size supported by the network to which the NLSP router is connected.

To change the size of the LSP that a router transmits, complete the following steps:

- 1** Load **NIASCFG**, then select the following parameter path:
Select **Configure NIAS > Protocols and Routing > Protocols > IPX > Expert Configuration Options**
- 2** Select **LSP Size** , enter a value between 128 and 4096, then press **Enter** .
- 3** Press **Esc** and save your changes.
- 4** Press **Esc** to return to the Internetworking Configuration menu.
- 5** If you want this change to take effect immediately, select **Reinitialize System** .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring RIP and SAP

RIP and SAP are the routing and service advertising protocols traditionally used by NetWare systems to exchange route and service information on an IPX network. RIP is currently the most common routing protocol used on IPX networks.

RIP and SAP perform well in small networks that have simple architectures and few routers. These protocols, however, begin to reveal their limitations in the large, complex internetworks that are becoming increasingly common throughout the installed base of NetWare systems.

By default, RIP and SAP packets are broadcast every 60 seconds, even if no change has occurred anywhere in a route or service. Depending on the size and speed of the network, these periodic broadcasts can consume a significant amount of bandwidth and burden NetWare nodes, especially over WAN links.

Novell Internet Access Server 4.1 provides a configurable parameter, Periodic Update Interval, that enables you to control how often a router broadcasts its route and service updates. This, along with other configurable parameters, such as Aging Interval Multiplier, Pace Override, and Packet Size Override, enables you to fine-tune the operation of RIP and SAP on your router. However, the default values for these parameters are sufficient for most IPX networks and should be changed only on the advice of your technical support representative. A misconfiguration can cause the router to lose routes and services or even generate more traffic than usual.

When configuring RIP on an interface, you can set the RIP State parameter to one of three states: Auto, On, or Off. Auto, the default state, enables the router to accept incoming RIP packets and rebroadcast their routes only if RIP-broadcasting devices, such as NetWare 2 servers, are operating on the attached network. If those devices are removed from the attached network, the Auto-state interface responds by automatically disabling RIP and enabling NLSP. On enables the router to exchange RIP packets freely with other RIP routers on the network. Off disables RIP routing on the interface but does not prevent the router from responding to incoming requests for RIP routes from local NetWare workstations.

Similarly, you configure SAP on an interface with the SAP State parameter, which can also assume one of three states: Auto, On, and Off. Auto, the default state, enables the router to accept incoming SAP packets and rebroadcast their services only if SAP-broadcasting devices, such as NetWare 2 servers, are operating on the attached network. If those devices are removed from the attached network, the Auto-state interface responds by automatically

disabling SAP and enabling NLSP. On enables the router to exchange SAP packets freely with other routers on the network. Off disables SAP advertising on the interface but does not prevent the router from responding to incoming requests for services from local NetWare workstations. Additionally, the Off state still allows the router to import locally advertised services.

NOTE: If you want to filter routes or services between routers, use RIP and SAP. NLSP routers cannot filter routes or services.

You should avoid running RIP and SAP over WAN connections because of the cost they incur from periodic transmissions.

This topic contains the following sections:

- ◆ “How to Configure RIP” on page 78
- ◆ “How to Configure SAP” on page 80
- ◆ “Accepting and Advertising Services from a Network Not Listed in the Routing Information Table” on page 82

How to Configure RIP

To configure RIP, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX

- 2** Select Routing Protocol , then select RIP/SAP Only .

Select RIP/SAP Only only if your network has no NLSP routers.

- 3** Press Esc to return to the Internetworking Configuration menu, then select the following path:

Select Bindings > a network interface > Expert Bind Options > RIP Bind Options

- 4** Select RIP State .

If you want to run RIP over the interface, select On.

This state is necessary for some third-party products that require RIP to operate.

If you do not want to run RIP over the interface, select Off.

This state disables backward compatibility with older routers and servers that use and depend on RIP.

If non-NLSP devices, such as NetWare 2 servers, are operating on the attached network and you want the router to accept and broadcast RIP packets received from these devices, select `Auto`.

`Auto` is the default state.

- 5** If you want to customize RIP operation further, configure one or more of the following parameters:

IMPORTANT: Because the default settings for these parameters are suitable for most RIP-based IPX networks, you should change them only for a specific purpose. Misconfiguring these parameters can increase routing traffic or cause loss of connectivity on your network.

- ◆ **Periodic Update Interval**—Measured in 30-second units, determines the interval at which RIP packets are transmitted through this interface. The default value is 2 (60 seconds).

Each router on the network segment to which this router is attached must use the *same* value for the Periodic Update Interval .

- ◆ **Aging Interval Multiplier**—Controls how long the router keeps route information received through this interface. The product of this parameter and the RIP Periodic Update Interval specifies how long the router keeps route information from periodic RIP updates received through an interface.

Increasing the Aging Interval Multiplier slows the rate at which the router ages the routes in its Routing Information Table. This is necessary to keep routes that might otherwise be aged out of the routing table because of dropped RIP updates.

The default value for the Aging Interval Multiplier is 4. For example, if RIP packets are sent every 60 seconds (Periodic Update Interval equals 2), the router keeps the route information for 240 (60 x 4) seconds without refreshing it.

Each router on the network segment to which this router is attached must use the *same* value for the Aging Interval Multiplier .

- ◆ **Pace Override**—Specifies the maximum number of RIP packets that can be sent each second through this interface. The default—and maximum—pace for RIP packets is 9 pps.

The default value for this parameter is 0, which means do not override the current pace.

- ◆ Packet Size Override —Specifies the size, in bytes, of RIP packets sent on this interface. The default value is 0, which means do not override the current value.

Each router on the network segment must use the *same* RIP packet size.

- 6** Press Esc and save your changes.
- 7** Press Esc to return to the Internetworking Configuration menu.
- 8** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

How to Configure SAP

To configure SAP, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Protocols > IPX
- 2** Select Routing Protocol , then select RIP/SAP Only .
- 3** Press Esc to return to the Internetworking Configuration menu, then select the following path:
Select Bindings > a network interface > Expert Bind Options > SAP Bind Options
- 4** Select SAP State .

If you want to run SAP over the interface, select On.

This state is necessary for some third-party products that rely on SAP to advertise their services. Setting SAP State to On also enables RIP routing on the interface.

If you do not want to run SAP over the interface, select Off.

This state disables backward compatibility with older routers and servers that use and depend on SAP. The router responds to incoming SAP requests, such as Get Nearest Server, even if SAP is disabled on this interface.

If non-NLSP devices, such as NetWare 2 servers, are operating on the attached network and you want the router to accept and broadcast SAP packets received from these devices, select Auto.

Auto is the default state.

- 5** If you want to customize SAP operation further, configure one or more of the following parameters:

IMPORTANT: Because the default settings for these parameters are suitable for most RIP-based IPX networks, you should change them only for a specific purpose. Misconfiguring these parameters can increase routing traffic or cause loss of connectivity on your network.

- ◆ Get Nearest Server Requests Override —Determines whether the router accepts or ignores SAP Get Nearest Server requests it receives through this interface. Select one of the following options:

No Override —Do not override the global setting for the router. This is the default state.

Ignore —Ignore Get Nearest Server requests received through this interface.

Accept —Accept Get Nearest Server requests received through this interface.

- ◆ Periodic Update Interval —Measured in 30-second units, determines the interval at which SAP packets are transmitted through this interface. The default value is 2 (60 seconds).

Each router on the network segment to which this router is attached must use the *same* value for the Periodic Update Interval .

- ◆ Aging Interval Multiplier —Controls how long the router keeps service advertisements received through this interface. The product of this parameter and the SAP Periodic Update Interval specifies how long the router keeps service information from periodic SAP updates received through an interface. This parameter is a holding multiplier for the SAP Periodic Update Interval .

The default value for the Aging Interval Multiplier is 4. For example, if SAP packets are sent every 60 seconds (Periodic Update Interval equals 2), the router keeps the service information for 240 (60 x 4) seconds without refreshing it.

Each router on the network segment to which this router is attached must use the *same* value for the Aging Interval Multiplier .

- ◆ Pace Override —Specifies the maximum number of SAP packets that can be sent each second through this interface. The default—and maximum—pace for SAP packets is 9 pps.

The default value for this parameter is 0, which means do not override the current pace.

- ◆ Packet Size Override —Specifies the size, in bytes, of SAP packets sent on this interface. The default value is 0, which means do not override.

Each router on the network segment must use the *same* SAP packet size.

6 Press Esc and save your changes.

7 Press Esc to return to the Internetworking Configuration menu.

8 If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Accepting and Advertising Services from a Network Not Listed in the Routing Information Table


Novell Internet Access Server 4.1 routing software includes a SET command that enables an IPX router to accept and advertise a service from another network, even if the network number associated with the service is not listed in the router's Routing Information Table.

The syntax of this command is as follows:

```
SET REQUIRED NETWORK FOR SERVICES=ON|OFF
```

To understand how this command works, consider [Figure 12 on page 83](#), which shows two directly connected IPX routers, one serving a large backbone network, the other serving a NetWare LAN.

Figure 12 IPX Router Accepting and Advertising Services from a Network Not Listed in the Routing Information Table



Suppose a NetWare server somewhere within the backbone network advertises its file services, which the backbone router receives as a SAP packet (Service Type=0x0004 and Network Number=0x01234567, for example). When the LAN router receives the SAP packet from the backbone router, it checks its Routing Information Table for the network number 0x01234567.

If the LAN router finds the network number, it adds the associated service information to its services table and advertises the service to the LAN during the next SAP broadcast. If the LAN router *does not* find the network number, it discards the packet. This function occurs if the following command has been entered at the LAN router:

```
SET REQUIRED NETWORK FOR SERVICES=ON
```

This function is disabled by default. However, if this function has been enabled by the preceding command, it can be disabled by entering the following command:

```
SET REQUIRED NETWORK FOR SERVICES=OFF
```

If the previous command has been entered, the LAN router will not discard the packet if it does not find the network number in its Routing Information Table. If the router does not find the network number, it first checks for the nearest NLSP level 2 router and then for the RIP default route (0xFFFFFFFF) or an aggregated route. If the LAN router finds one of these, it adds the associated service information to its services table. If the LAN router does not find any one of these, it then discards the SAP packet.

Proxying a NetWare File Server

If you have a dedicated router—a PC running Novell Internet Access Server 4.1 over a two-user version of NetWare 4.11—and several NetWare workstations operating on a network, the router can reply with a NetWare file server's name—instead of its own—when it receives a SAP Get Nearest Server request. This is called *proxying a file server*.

Proxying avoids the situation in which several workstations on a network restart simultaneously, and the only path to a file server is through the dedicated router. When each workstation restarts, it sends a Get Nearest Server request to the router. Because the dedicated router cannot support multiple, simultaneous logins, all but one of the requesting workstations lock up and fail. Because the proxied file server has multiple connection slots, it can handle simultaneous NCP connection requests from the workstations.

A dedicated router always replies to Get Nearest Server requests. As a proxy, the router still replies but gives the proxy name rather than its own. In fact, the router replies with the server name even if the server is not active.

In addition to having a server proxy on the network, each NetWare workstation should be configured with a *preferred server*:

This topic contains the following sections:

- ◆ [“How to Proxy a NetWare File Server” on page 84](#)
- ◆ [“How to Check the Proxy Configuration” on page 85](#)

How to Proxy a NetWare File Server

Before you begin, you must know the name of the NetWare file server you want to proxy.

To proxy a NetWare file server, complete the following steps:

- 1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX > Expert Configuration Options

By default, the Get Nearest Server Requests parameter is set to Accept. The setting you select applies to all interfaces. You can, however, override this setting on a particular interface. For a description of how to configure a network interface, refer to "Setting Up" in the documentation for the type of WAN interface you are using.

- 2** If you want the router to ignore Get Nearest Server requests, set Get Nearest Server Requests to Ignore .
- 3** Select Override Nearest Server , then select Enabled .

This parameter enables the router to respond to a workstation's Get Nearest Server requests with the proxied server name instead of its own name.
- 4** Select Nearest Server , then enter the name of a reliable server (one that is operating most of the time).
- 5** Press Esc and save your changes.
- 6** Press Esc to return to the Internetworking Configuration menu.
- 7** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

How to Check the Proxy Configuration

To make sure the router is proxying for the correct file server, complete the following steps:

- 1** At the router console prompt, enter
TRACK ON
- 2** Display the SAP Tracking screen.
- 3** Restart one of the NetWare workstations on the network.
- 4** Watch the SAP Tracking screen for the workstation's Get Nearest Server request and for the router's reply with the correct server name.

If you supplied the wrong name or the name of a server that the router cannot reach, the SAP Tracking screen displays the following message:

```
No response to GNS sent - no route to configured server
<server_name>
```

- 5** Return to the router console, then enter
TRACK OFF

Configuring the IPX Address Mapping Gateway

The IPX Address Mapping Gateway allows you to connect to a backbone network even when your local network numbers are not compatible with the backbone addressing scheme.

To configure the IPX Address Mapping Gateway, complete the following steps:

- 1** Load NIASCFG and select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX

- 2** Select Address Mapping Gateway and select Enabled .

- 3** Select Address Mapping Gateway Configuration , select Address Mapping Network Number , and enter the number to which your local network will be mapped.

NOTE: You must enter a registered address unique to the backbone.

One number is supported for each router. Additionally, this number is included as part of the SAP name advertised by the IPX Address Mapping Gateway. The SAP name is used by other gateways to locate gateways that use the same address mapping network number when a packet with an unknown reverse mapping is received.

- 4** Configure the following parameters as needed.

To configure the maximum number of address mappings that is allocated during router initialization, select Maximum Address Mapping Entries *and enter the desired number of mapping entries.*

The default is 1,000 mapping entries.

To configure the amount of time an address mapping is remembered after the last mapping entry was used, select Address Mapping Hold Time *and enter the desired amount of time in minutes, hours, and days.*

After the holding time expires, the mapping is dropped and new packets must flow from the customer network to the backbone to renew the mapping. This process allows mapping slots to be reused. The default hold time for mapping is one hour.

To use outbound RIP filters to determine nonmappable networks, in addition to manually configuring nonmappable entries, select Use RIP Filters for Nonmappable Networks *and select* Enabled.

If the source network number in a packet being forwarded to an IPX Address Mapping Gateway circuit passes the outbound RIP filter on that circuit, the packet is not mapped. If the source network number does not pass a RIP filter, the packet is mapped.

The default is Disabled . When disabled, RIP filters are not used to determine nonmappable network addresses.

NOTE: If enabled, the RIP filter module must be loaded from the IPX protocol menu, and the RIP filter must be configured carefully to block the correct network numbers.

To configure a SAP type list that is used to determine networks that are nonmappable, select Nonmappable SAP Types and select one of the predefined SAP types or press Ins to add a new type. To add a new SAP type, enter the desired SAP type or press Ins and select a SAP type from the list of known service types.

The SAP tables are scanned for SAP entries with matching SAP types. After finding all matching SAP types, the IPX Address Mapping Gateway determines the network numbers on which the services are found and applies those network numbers to the list of nonmappable networks.

This option makes configuring nonmappable network numbers easier. For example, all packets originating from the Novell Directory Services™ (NDS™) software or NetWare Mobile IPX™ software should not be translated. Therefore, SAP types for NDS, NetWare Mobile IPX, and Timesync are included in the list by default.

NOTE: To avoid mapping NetWare 3.x servers, add SAP type 4 to the list of nonmappable SAP types.

5 Press Esc and save your changes.

6 Press Esc to return to the Internetworking Configuration menu.

7 Select the following parameter path:

For LAN interfaces, select Bindings > a LAN interface > Expert Bind Options .

For WAN interfaces, select Bindings > a WAN interface > WAN Call Destinations > a call destination > Expert Options .

8 To enable the IPX Address Mapping Gateway on an interface, select Use For Address Mapping Gateway and select Yes .

When this option is enabled, all packets destined for the interface or WAN call destination are subject to the address mapping rules.

Configuring IPX Route Aggregation

IPX Route Aggregation enables you to introduce routes learned through RIP in a summarized form. Route aggregation compactly describes many IPX network numbers simultaneously by using an address and mask pair. For example, all addresses from C9000000 to C9FFFFFF can be represented using the address C9000000 and the mask FF000000.

To configure IPX Route Aggregation, complete the following steps:

- 1** Load NIASCFG and select the following parameter path:

For LAN interfaces, select Configure NIAS > Protocols and Routing > Bindings > a LAN interface > Expert Bind Options > Aggregate Routes .

For WAN interfaces, select Configure NIAS > Protocols and Routing > Bindings > a WAN interface > WAN Call Destinations > a call destination > Aggregate Routes .

- 2** Press `Ins` and configure the following parameters:

- 2a** Select Address Summary and enter the prefix for the network addresses to be aggregated.

For example, any network beginning with C9, such as C9123829 or C9823878, can be aggregated using the address summary C9000000 with a mask of FF000000.

- 2b** Select Mask and enter a number with Fs for the part of the mask that corresponds to the common prefix of all network addresses to be aggregated and 0s for the rest of the mask.

For example, to summarize all addresses that begin with the same three numbers, set the mask to FFF00000.

- 2c** Optionally, to advertise an aggregate route only when at least one of the routes is received, select Type and set it to Dynamic .

For LANs, Type can be set only to Dynamic . If you select Static for a WAN interface, the aggregate route is always advertised.

Aggregate routes associated with static on-demand WAN calls should always be set to Static .

NOTE: If Type is set to Dynamic , only routes learned through RIP will trigger the dynamic aggregate route to be advertised. Routes learned through NLSP will not trigger the dynamic aggregate route to be advertised. Routers that have dynamic aggregate routes configured on some interfaces should have NSLP disabled and RIP enabled on those interfaces.

- 3 Press Esc and save your changes.
- 4 Press Esc to return to the Internetworking Configuration menu.

Controlling the Propagation of Type 20 Packets

Novell Internet Access Server 4.1 enables you to control the propagation of type 20 packets with the Advanced Packet Type 20 Flooding parameter. Type 20 is an IPX packet type that refers to any propagated packet. NetBIOS packets, for example, are type 20 packets.

The Advanced Packet Type 20 Flooding parameter can be set to one of the following options:

- ◆ **0**—Router discards, rather than propagates, any type 20 packet it receives. This option completely disables type 20 packet propagation.
- ◆ **1**—Router receives and propagates type 20 packets through *all* its interfaces, regardless of whether some of the interfaces are *equal-cost routes* to the same source.
- ◆ **2**—Router propagates type 20 packets only through interfaces that do not lead back to the source of the packets. For example, if Router A receives a type 20 packet from Router B, Router A forwards the packet only through interfaces that do not lead back to Router B. This is a packet forwarding mechanism known as *reverse path forwarding*.

The router does not propagate type 20 packets through the same interface from which it receives them. This is known as *split horizon*, a technique used with RIP and other distance vector routing protocols.

- ◆ **3**—Router propagates type 20 packets the same way as option 2, but *does not* forward them across WAN connections.

If you use FILTCFG to configure NetBIOS packet filters, be aware of the following interactions between these filters and the Advanced Packet Type 20 Flooding parameter:

- ◆ Setting Advanced Packet Type 20 Flooding to 1 (disabled) overrides the effect of NetBIOS packet filters operating on any network interface.
- ◆ Conversely, the action of NetBIOS packet filters on any network interface overrides the effect of setting the Advanced Packet Type 20 Flooding parameter to 2 (enabled).

How to Control Propagation of Type 20 Packets

To configure the propagation of type 20 packets from the server console, enter the following console command:

```
SET IPX NETBIOS REPLICATION OPTION = [0|1|2|3]
```

To check the current setting, you can type the command by itself.

To configure the propagation of type 20 packets from NIASCFG, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select Configure NIAS > Protocols and Routing > Protocols > IPX > Expert Configuration Options
- 2** Select Advanced Packet Type 20 Flooding and select one of the options described previously.
- 3** Press Esc and save your changes.
- 4** Press Esc to return to the Internetworking Configuration menu.
- 5** If you want this change to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Changing the Hop Count Limit for IPX Packets

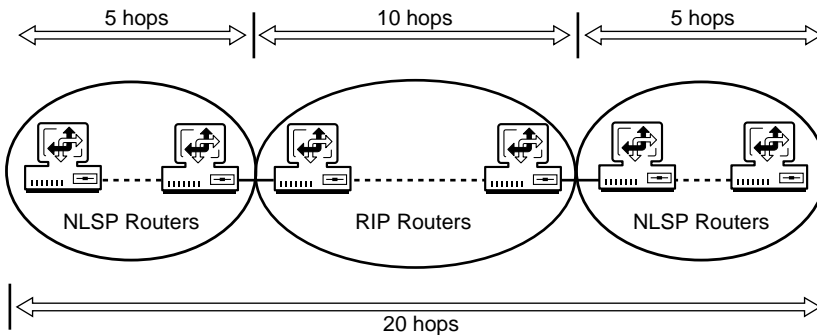
The Novell Internet Access Server 4.1 routing software enables you to increase the range of outbound IPX packets with the Hop Count Limit parameter. The hop count limit is the maximum number of routers (hops) an IPX packet can traverse before it is discarded. You can set the Hop Count Limit parameter to any number between 8 and 127; however, the default value of 64 is sufficient for most IPX networks.

NOTE: The Hop Count Limit parameter applies only to IPX packets. It does not increase the range of RIP and SAP packets, which are limited to 16 hops, or NetBIOS packets, which are limited to eight hops.

Before the release of NetWare MultiProtocol Router 3.0 and NetWare 4.1, the hop count limit for all IPX packets was 16. This limited the size, or *diameter*, of IPX networks.

If the diameter of your IPX network is close to the 16-hop limit, you should run NLSP on the routers at the network boundary to ensure continued connectivity across the network as it grows. **Figure 13** provides a simple, conceptual view of how this can work.

Figure 13 Running NLSP at the Boundary of a Large IPX Network



For more information about hop count, IPX routing, and related topics, refer to **Chapter 1, “Understanding,”** on page 11.

How to Change the Hop Count Limit

To change the hop count limit, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:
Select **Configure NIAS > Protocols and Routing > Protocols > IPX > Expert Configuration Options**
- 2** Select the **Hop Count Limit** parameter, enter a value between 8 and 127, then press **Enter** .
- 3** Press **Esc** and save your changes.
- 4** Press **Esc** to return to the **Internetworking Configuration** menu.
- 5** If you want this change to take effect immediately, select **Reinitialize System** .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Balancing Traffic Loads over Equal-Cost Routes

If a NLSP router has two or more network interfaces with routes to the same destination, it can distribute outbound traffic among those interfaces for an effective increase in throughput. This is called *load balancing* or *load sharing*.

NLSP uses an assigned *path cost* to select the best route by which to forward outbound IPX packets. The higher the throughput of the network medium, the lower the cost of the route.

Table 4 shows the throughput range and default cost of some typical network media.

Table 4 Throughput Range and Default Cost of Typical Network Media

Throughput Range	Default Cost	Typical Network Media
0-16 Kbps	61	9,600-baud line
48-64 Kbps	45	ISDN (U.S.)
64-128 Kbps	45	ISDN (Europe)
1-2 Mbps	27	Corvus Omninet (1 Mbps), T1 (1.5 Mbps)
2-4 Mbps	26	E1 (2 Mbps), ARCnet (2.5 Mbps)
4-8 Mbps	25	Token ring (4 Mbps), Corvus Omninet (4 Mbps)
10-16 Mbps	20	Ethernet (10 Mbps)
16-32 Mbps	19	Token ring (16 Mbps)
64-128 Mbps	14	FDDI (100 Mbps), CDDI (100 Mbps)

You can specify up to eight *equal-cost routes* to a single destination with the `Maximum Number of Path Splits` parameter. Two routes are equal in cost if the cost to the destination is the same for both routes. To equalize the costs of two interfaces, you set their `Cost Override` parameter to the same value. By default, `Cost Override` is set to 0 for all interfaces, which means that NLSP uses the default cost associated with the connected medium and throughput range listed in **Table 4**.

If you configure equal-cost routes on two or more interfaces, make sure the associated media throughputs fall within—or near—the same range, as indicated in [Table 4](#). For example, equal-cost routes between a 10-Mbps Ethernet link and a 16-Mbps token ring link are viable; equal-cost routes between a 4-Mbps token ring link and a 16-Mbps token ring link are not.

WARNING: Do not equalize the cost of routes whose throughputs differ greatly; this can interfere with the operation of IPX applications running over the network.

For more information about load balancing and path cost, refer to [Chapter 1, “Understanding,”](#) on page 11.

How to Balance Traffic Loads over Equal-Cost Routes

Before you begin, make sure of the following:

- ◆ NLSP is enabled on all interfaces you plan to configure.

To enable NLSP globally or on one or more interfaces, refer to [“How to Configure NLSP”](#) on page 74.

- ◆ The media over which you plan to configure equal-cost routes have the same or similar throughput ranges.

For a list of throughput ranges and associated media, refer to [Table 4](#) on page 92.

To configure load balancing over equal-cost routes, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX > Expert Configuration Options

- 2** Select the Maximum Number of Path Splits parameter, enter a value between 2 and 8, then press Enter .

Selecting a value of 2 or above automatically enables local load balancing over the specified number of equal-cost routes.

- 3** Press Esc and save your changes.
- 4** Press Esc to return to the Internetworking Configuration menu.

- 5** Enter a cost for each interface over which you want to balance IPX traffic.
 - 5a** Select the following path:
Select Bindings > a network interface > Expert Bind Options > NLSP Bind Options
 - 5b** Select the Cost Override parameter, enter a value between 1 and 63, then press Enter .
- 6** Press Esc and save your changes.
- 7** Press Esc to return to the Internetworking Configuration menu.
- 8** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring SPX Connection Parameters

Because some Novell and third-party NetWare applications place unique demands on the SPX transport protocol, NIASCFG enables you to adjust the values of the following parameters:

- ♦ Maximum IPX Socket Table Size —Maximum number of concurrent IPX sockets that can be opened by an application.
- ♦ SPX Watchdog Abort Timeout —Time, in ticks (about 1/18 of a second), SPX waits without receiving a packet from the other end of a connection before concluding that the connection is no longer valid.
- ♦ SPX Watchdog Verify Timeout —Time, in ticks, SPX waits without receiving a packet from the other end of a connection before requesting a watchdog, or keep-alive, packet.
- ♦ SPX Ack Wait Timeout —Time, in ticks, SPX waits without receiving an acknowledgment for a data packet it sent, before resending the packet.
- ♦ SPX Default Retry Count —Number of times SPX resends a data packet if it does not receive an acknowledgment.

The product of this parameter and the SPX Ack Wait Timeout is about how long it takes for SPX to conclude that the connection is no longer valid.

- ♦ Maximum Concurrent SPX Sessions —Maximum number of concurrent SPX sessions that can be opened by an application program.

The default values for these parameters are sufficient for most NetWare applications. Any application that requires a change to one or more of these parameters typically tells you so.

How to Configure SPX Connection Parameters

To adjust the value of any SPX connection parameter, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX > IPX/SPX Parameters

The IPX/SPX™ Parameters menu displays the SPX connection parameters.

2 Enter a new value for each parameter you need to change.

3 Press Esc and save your changes.

4 Press Esc to return to the Internetworking Configuration menu.

5 If you want this change to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Setting Delay and Throughput for a Slow Link

Delay is the time, in microseconds, to send a byte of information from one system to another. *Throughput* is the bandwidth of the network medium that connects the systems. Together, these parameters characterize a link between two systems or networks.

On WAN links, delay and throughput are estimated by the IPXWAN protocol. For this reason, you should not need to change these parameters on routers operating over a WAN link. On LAN links, the throughput is reported by the network interface driver; the delay is 200 microseconds, a constant used by all LAN media.

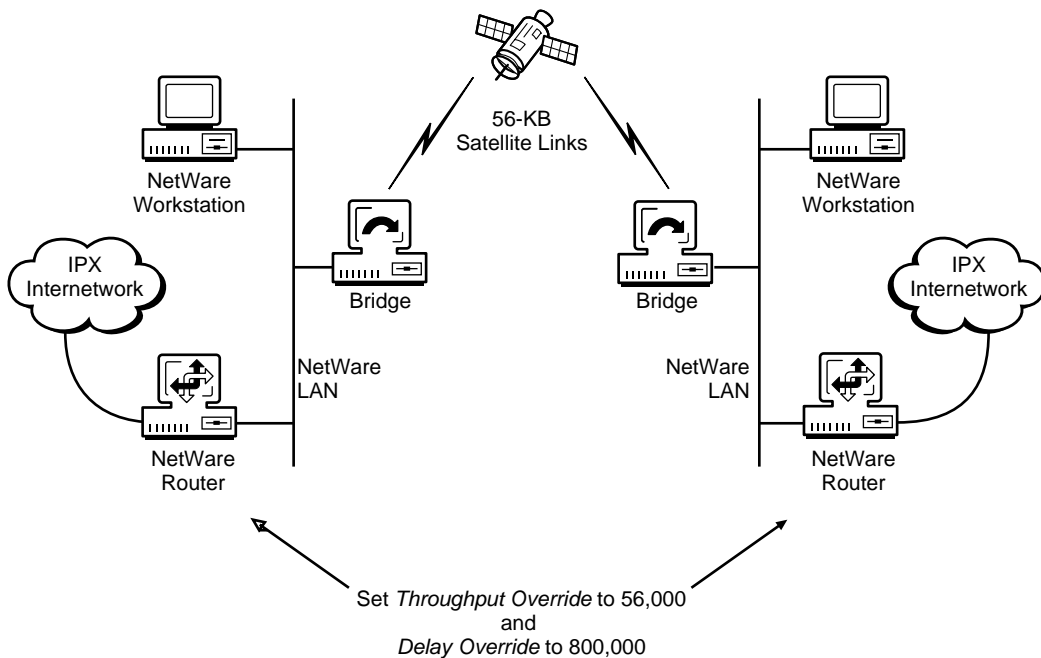
NLSP uses the delay and throughput values to calculate the number of ticks for a route to a destination network. The number of ticks associated with a route is directly proportional to the delay and inversely proportional to the throughput.

Some NetWare protocols, such as SPX, use the ticks value to calculate retransmit timers. If you are configuring LAN routers that must communicate over a bridge, a satellite, or both, you probably need to adjust the delay and throughput values on the routers. Setting the throughput to match the speed of the link and increasing the link delay prevent SPX retransmissions and timeouts between systems separated by a slow link.

Figure 14 shows two NetWare LANs joined by two bridges communicating over a satellite link. To enable the workstations to communicate with the router and the systems in the IPX internetwork on the other end of the link, you set the Throughput Override on each router to 56,000—the throughput of the satellite link—and the Delay Override to 800,000—an arbitrary (but sufficiently high) value to prevent timeouts over the link.

IMPORTANT: Although this configuration enables systems on each end of the link to communicate through the routers, it *does not* enable direct workstation-to-workstation communication between the two LANs.

Figure 14 Setting Delay and Throughput for Systems Communicating over a 56-KB Satellite Link



How to Set Delay and Throughput for a Slow Link

To set delay and throughput on an interface, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > a network interface > Expert Bind Options

- 2** Select Delay Override , then enter a value.

By default, this parameter is set to 0, which means the router uses the default value for LANs or the value estimated by IPXWAN. The valid range is from 1 to 5,000,000 microseconds. One tick equals 55,000 microseconds, or about 1/18 of a second.

The value you enter overrides the default delay for this interface.

- 3** Select Throughput Override , then enter a value.

By default, this parameter is set to 0, which means the router uses the value reported by the LAN driver or estimated by IPXWAN. The valid range is from 300 to 4,294,967,295 bps.

The value you enter overrides the default throughput for this interface.

- 4** Press Esc and save your changes.

- 5** Press Esc to return to the Internetworking Configuration menu.

- 6** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring IPX for Wireless Connectivity

In addition to the standard IPX routing software for wired LANs and WANs, Novell Internet Access Server 4.1 routing software provides wireless connectivity for portable NetWare workstations through NetWare Mobile IPX™ software.

NetWare Mobile IPX consists of router and mobile client software that work in concert to shield network users from the protocol and Network-layer interruptions that occur when a user changes network interfaces or locations during a network session.

This topic contains the following sections:

- ♦ “Configuring a Home Router” on page 98
- ♦ “Configuring a Mobile Client” on page 99

Configuring the NetWare Mobile IPX Home Router and client software is straightforward and simple. The only decision you need to make before you get started is where to locate the Home Router on your network. The next section helps you determine the best location.

Configuring a Home Router

The Home Router serves as the central connection point between mobile clients and NetWare servers. To enable mobile clients to establish and maintain network connections, the Home Router allocates an address from the server's IPX internal network for use by the mobile clients.

How to Configure a Home Router

To configure a Home Router, complete the following steps:

- 1** Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > IPX

- 2** Select Mobile IPX Support , then select Enabled .

NOTE: Selecting Enabled automatically causes client validation on remote access servers to be disabled. Remote access systems check whether packets received from a WAN client have the same source IPX node address that was assigned to the client during the IPXWAN™ negotiation phase. Because NetWare Mobile IPX uses a different filtering method, client validation is turned off to prevent NetWare Mobile IPX packets being discarded by remote access servers.

- 3** Select Mobile IPX Configuration and configure the Home Router parameters.

- 3a** Select Time To Live Override and enter a value, in minutes, from 1 to 10080.

Time To Live Override overrides the mobile client's HR Time To Live parameter, which defines how long the Home Router serves the mobile client without receiving a response from the client. Each time the Home Router receives information from the mobile client, the Time To Live Override counter is reset to the value you enter here. A value of 0, the default, disables the override.

A mobile client cannot obtain a Time To Live value longer than the one you specify here.

- 3b** The Watchdog Spoofing parameter is enabled by default; to disable watchdog spoofing on the Home Router, select Disabled .

Watchdog Spoofing controls whether the Home Router answers NetWare Core Protocol™ (NCP™) watchdog packets on behalf of a mobile client. If Watchdog Spoofing is enabled, users do not lose their connections to file servers as they roam out of wireless range.

- 3c** Configure the Broadcast to Virtual Network parameter.

Broadcast to Virtual Network directs the Home Router to forward or discard broadcast packets destined for the virtual network that the router uses to communicate with its mobile clients.

If a large number of broadcast packets are being directed at mobile clients, or if a mobile client's application does not require broadcast, select Discard . By directing the Home Router to discard broadcast packets, you reduce the amount of bandwidth used on the network.

- 4** Press Esc to return to the Internetworking Configuration menu; save your changes when prompted.

- 5** If you want these changes to take effect immediately, select Reinitialize System .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Configuring a Mobile Client

To configure a mobile client, you modify the client's STARTNET.BAT and NET.CFG files. The changes you make to STARTNET.BAT are required for NetWare Mobile IPX connectivity. Changes to NET.CFG are optional; they are required only if you want to customize the client's NetWare Mobile IPX configuration.

How to Configure a Mobile Client

Before you begin, you must complete the following tasks:

- ◆ Install the standard client software. For instructions, refer to the NetWare client documentation.
- ◆ Install the Novell Internet Access Server 4.1 mobile IPX client.

- ◆ Configure a mobile client board that is mobile aware and supports PCMCIA card in/card out capability, in-range and out-of-range capability, and NESL.

To configure a mobile client, complete the following steps:

- 1** Open the client's STARTNET.BAT file using a text editor.

A typical STARTNET.BAT file looks something like this:

```
@ECHO OFF

SET NWLANGUAGE=ENGLISH

CD NWCLIENT

LSL

driver

IPXODI

VLM /ps=server_name

CD \
```

- 2** Add the following changes, indicated in bold, in the order shown:

```
@ECHO OFF

SET NWLANGUAGE=ENGLISH

CD NWCLIENT

LSL

NESL

driver

IPXODI /M

VLM /ps=server_name

CD \
```

NESL must be loaded for the mobile client to be activated. NetWare Mobile IPX reacts to changes in the system, such as location and the client's adapter board. The MAC driver is the system module that knows of these events (for example, out of range of access point coverage, card insertion or removal, and so on) and notifies IPX of such changes through NESL.

The IPXODI **/M** switch enables the NetWare Mobile IPX client software.

A STARTNET.BAT file configured for NetWare Mobile IPX operation looks something like this:

```
@ECHO OFF
CD C:\NWCLIENT
SET NWLANGUAGE=ENGLISH
LH C:\NWCLIENT\LSL.COM
LH C:\NWCLIENT\NESL.COM
LH C:\NWCLIENT\NE2000.COM
LH C:\NWCLIENT\IPXODI.COM /M
C:\NWCLIENT\VLM.EXE /ps=MY_SERVER
```

- 3 Restart the client.

How to Customize Your Mobile Client

This section describes the optional parameters you can add to the *Mobile IPX* section of a mobile client's NET.CFG file. The parameters enable you to customize your NetWare Mobile IPX configuration.

An example of how these parameters are used is provided in “[Example NET.CFG File](#)” on page 105.

Customizing Home Router Parameters

The following parameters enable you to customize the interaction between a mobile client and its Home Router:

- ◆ **Preferred HR=Home_Router_Name**

This command causes IPXODI to attempt to attach to the specified Home Router (HR). If the router does not exist or is not specified, the Home Router closest to the client is used.

This command enables some level of routing optimization to be achieved. Specifying a Home Router that is the user's preferred server, or specifying one in an inline routing path between the mobile client and most of its logged-in servers, causes packets sent back to the client to take a more direct path because all packets being sent to the client go through the Home Router first.

- ◆ **HR Time To Live=x** (where $x = 5$ to 10,080 minutes)

HR Time To Live specifies the time-to-live interval, in minutes, that the IPXODI module attempts to use with the Home Router. It defines how long the Home Router serves the mobile client before the router requires an update from the mobile client. If this is not defined, IPXODI uses a default value of 30 minutes.

NOTE: HR Time To Live can be overridden by the Home Router's Time To Live Override parameter.

If the client does not update the Home Router after the HR Time To Live value runs out, the Home Router stops serving the mobile client. Note that only NCP watchdog packets, *not* Sequenced Packet Exchange™ (SPX™) watchdog packets, are handled by the Home Router. If the mobile client roams out of range and comes back within the amount of time set by HR Time To Live, the timer is reset automatically; otherwise, if the client is not back within range when the timer runs out, the mobile client is dropped.

It is especially important to set HR Time To Live to a large value if the mobile client is out of network range for a long time. While the time-to-live value is still active in the Home Router, the router responds to server NCP watchdogs on behalf of the client so that client sessions do not time out while network connectivity is lost. If this value is too small, the Home Router stops serving the mobile client before it returns within network range, and all server connections are lost.

The only reasons the Home Router might not see a NetWare Mobile IPX watchdog packet from a mobile client are that the client is off, in sleep mode, or out of range.

◆ **Allow HR Change=[On|Off]**

This command determines what IPXODI does when the current Home Router is no longer reachable.

If you set Allow HR Change to On, IPXODI tries to sign on with the first available Home Router, even if it is not the same as the current one.

If you set Allow HR Change to Off, IPXODI continues trying to reestablish a connection to the Home Router to which the client is currently attached.

If Allow HR Change is not defined, IPXODI assumes Allow HR Change is set to Off; if it is set to On, and the current Home Router stops operating, IPXODI obtains a different virtual address while signing on with a different router. Most applications available today cannot operate

gracefully through the address change; as a result, connections might be terminated.

Specifying an Alternate Board

The NET.CFG parameters described in this section enable you to specify an alternate board to be used in the portable computer. Alternate board parameters enable IPXODI to use a second board for mobile communications if the primary board loses connectivity. The driver for the alternate board must be specified in the Mobile IPX section of NET.CFG. Note that the equal sign (=) is optional; however, it can be used to quickly find the parameter values.

IMPORTANT: The first three parameters (Alt Name , Alt Board Number , and Alt Frame) must *all* be specified; if one is missing, no alternate board setting is used. Additionally, the frame type specified by the Alt Frame parameter must be set under the Link Driver Heading for the desired alternate board.

- ◆ **Alt Name**=*Alternate_Driver_Name*

Alt Name specifies the name of the driver supporting the alternate board.

- ◆ **Alt Board Number**=*Alternate_Board_Number*

Alt Board Number specifies the Link Support Layer™ (LSL™) board number of the alternate board displayed when the driver loads.

The board number of a driver changes if there is a change in the order of MAC driver load commands. Therefore, it is important to always load the primary driver first, followed by the alternate driver, to ensure that the Alt Board Number parameter always refers to the alternate board.

After the MAC drivers have been loaded, you can get board numbers and other information by entering the following command:

```
MAC_Driver_Name /s
```

This command displays information about all Open Data-Link Interface™ (ODI™) drivers currently loaded. The information you see is similar to the following example:

The following LAN drivers are loaded in memory:

MAC driver name and version information

IRQ 5, Port 300, Mem D0000, Node Address 4096003F53 L

Max Frame 1514 bytes, Line Speed 2 Mbps

Board 1, Frame ETHERNET_II, LSB Mode

◆ **Alt Frame**=*Alternate_Frame_Type*

Alt Frame specifies the frame type for the alternate board (for example, ETHERNET_802.2).

The driver loaded first in STARTNET.BAT becomes the primary driver. If IPXODI cannot locate the alternate driver during initialization, an error is issued and only the primary board is used. If there is no primary board either, IPXODI issues an error and fails to load.

The drivers for both the primary and alternate boards must be loaded before IPXODI. The driver load order is not important; however, it is harder to determine the LSL board numbers to enter for the IPX BIND statement or alternate board configurations if the primary driver is not loaded first. If PCMCIA adapters are used and the card vendor's drivers are written to support card in/card out events, or driver initialization without the PCMCIA card inserted, the PCMCIA cards need not be inserted into the system until network connectivity is needed.

Here is an example STARTNET.BAT file that shows the load order of the primary driver, alternate driver, and IPXODI:

```
@ECHO OFF
SET NWLANGUAGE=ENGLISH
CD NWCLIENT
LSL
NESL
primary driver
alternate driver
IPXODI /M
VLM /ps=server_name
CD \
```

Specifying Watchdog Protocol Operation

The NET.CFG file allows you to specify whether the SPX Watchdog protocol will be run to validate SPX connections periodically. The one-line entry to specify the SPX Watchdog protocol behavior follows the Heading *PROTOCOL IPX* and has the following format:

SPX WATCHDOGS = ON|OFF

where *ON* specifies that the Watchdog protocol will be run to validate SPX connections periodically and *OFF* specifies that it will not. For more

information on the use of the Watchdog protocol, refer to “[Watchdog Packet Spoofing](#)” on page 30.

Example NET.CFG File

This section provides an example NET.CFG file that shows the format of mobile client customization parameters.

```
LINK DRIVER Wireless
    FRAME = ETHERNET_802.2
LINK DRIVER Wireless2
    FRAME = ETHERNET_802.2
NETWARE DOS REQUESTER
    NETWARE PROTOCOL = NDS BIND
    FIRST NETWORK DRIVE = F
    SHOW DOTS = ON
    USE DEFAULTS = ON
    VLM = AUTO.VLM
MOBILE IPX
    PREFERRED HR = Home_Router_Name
    ALT NAME = Wireless2
    ALT BOARD NUMBER = 2
    ALT FRAME = ETHERNET_802.2
PROTOCOL IPX
    SPX WATCHDOGS = OFF
```

Configuring the MacIPX Gateway

The Novell Internet Access Server 4.1 routing software includes MACIPXGW.LAN, a LAN driver that enables your router to operate as a gateway between Internetwork Packet Exchange (IPX) networks and Macintosh* clients running MacIPX[®] applications on AppleTalk networks. Macintosh clients use the *MacIPX gateway* to exchange data with NetWare clients and to use the resources available on IPX networks.

NOTE: MacIPX provides support for the IPX protocol on Macintosh computers. It does not enable Macintosh users connected to the IPX network to log in to a

NetWare server or print documents on NetWare printers. Users and developers must rely on NetWare for Macintosh software for NetWare file and print services.

This topic contains the following sections:

- ◆ “Configuring and Binding the Gateway Driver” on page 107
- ◆ “Restricting Gateway Service to Selected Networks” on page 109

You can use the MacIPX gateway if your networks have the following characteristics:

- ◆ You have IPX and AppleTalk networks that you want to connect and these networks are part of a LAN running NetWare 3.11 or later, or the NetWare MultiProtocol Router™ 2.0 software or later. The MacIPX gateway must run on one of these networks.
- ◆ Your AppleTalk networks support MacIPX clients.
- ◆ One or more of the networks use only the AppleTalk protocol family to connect Macintosh clients to the network.

From a user perspective, the MacIPX gateway is required only for Macintosh users who select the AppleTalk icon in the MacIPX Control Panel. If all Macintosh users select either the Ethernet or Token Ring icon, and if IPX traffic is permitted on those networks, then you do not need a MacIPX gateway.

MacIPX applications automatically select an IPX gateway *only* when the gateway is in the zone that contains the Macintosh client running MacIPX. If this is not the case, use the MacIPX Control Panel to configure MacIPX to look for IPX gateways in specific zones.

You should locate a MacIPX gateway so that the amount of configuration required by MacIPX is minimized. For example, if you have an AppleTalk network for dial-in users that provides service for AppleTalk Remote Access (ARA), you should ensure that a MacIPX gateway serves the AppleTalk zone that includes the ARA network so that Macintosh clients using ARA do not require MacIPX configuration.

Configuring and Binding the Gateway Driver

Configuring the MacIPX gateway is similar to configuring a typical LAN board and binding a network protocol to the board.

Before you begin, you must complete the following tasks:

- ♦ Ensure that your router has at least 65 KB of RAM available.
- ♦ Ensure that APPLETLK.NLM is loaded and configured.
- ♦ Ensure that IPX packet forwarding is turned on.
- ♦ Know the network number of the IPX network to which the gateway interface is attached.
- ♦ Know the number of MacIPX clients that will be served by the gateway.

To configure the MacIPX gateway, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Boards > Press Ins > MACIPXGW

2 Configure the MACIPXGW driver.

2a Select Board Name and assign a name to the gateway driver.

2b Select Gateway Name and assign a name to the MacIPX gateway.

This name is used to advertise the MacIPX gateway on the AppleTalk network. Because the name appears in the MacIPX Control Panel, it should be one that users recognize easily.

If you do not provide a name, the MacIPX gateway uses the name of the router on which the MacIPX gateway is installed.

2c Select Unicast Threshold and enter a value between 1 (the default value) and 4294967295.

This parameter controls how the MacIPX gateway propagates IPX broadcast packets to AppleTalk networks.

If you want to send IPX broadcast packets to all AppleTalk networks with MacIPX clients, enter a number less than the number of MacIPX clients served by the gateway.

Macintosh systems not running MacIPX applications do not understand IPX broadcast packets and discard them. When this

option is used, unnecessary packets are distributed to non-MacIPX clients on the network.

If you want to send IPX broadcast packets to each MacIPX client, enter a number equal to or higher than the number of MacIPX clients served by the gateway.

If the number of clients exceeds this threshold, the MacIPX gateway starts sending broadcast packets. Using this option can increase network traffic because a single IPX broadcast packet could become many AppleTalk unicast packets, depending on the number of MacIPX clients.

2d If you want to enter a note or comment about the gateway, select **Comment** and enter the information.

2e Press **Esc** to return to the Internetworking Configuration menu; save your changes when prompted.

3 Bind IPX to the gateway.

Binding IPX to the gateway causes the MacIPX Gateway icon to appear in the MacIPX Control Panel.

3a Select the following parameter path:

Select **Bindings** > Press **Ins** > **IPX** > the MacIPX interface

3b Select **IPX Network Number** and enter the network number of the IPX network to which the interface is attached.

4 Press **Esc** to return to the Internetworking Configuration menu; save your changes when prompted.

5 If you want these changes to take effect immediately, select **Reinitialize System** .

If you want to configure other parameters, do so now, then reinitialize the system when you are finished.

Restricting Gateway Service to Selected Networks

The MacIPX gateway, by default, serves all AppleTalk networks that make up the AppleTalk cloud. If you want the gateway to serve only selected AppleTalk networks, you must use a configuration file called MACIPXGW.DAT, which resides in SYS:SYSTEM.

You indicate the AppleTalk networks you want the gateway to serve—or not serve—by placing commands in MACIPXGW.DAT with the following syntax:

```
[exclude | include]  
  <net_number >[-net_number ]  
  . . . ]
```

The first line in the preceding example is a keyword that specifies the following modes of inclusion:

- ♦ *exclude* —Directs the MacIPX gateway to serve all AppleTalk networks *except* those whose numbers are listed on the following lines.
- ♦ *include* —Directs the MacIPX gateway to serve only networks whose numbers are listed on the following lines; this is the default mode if no keyword is specified.

The AppleTalk networks can appear as a number or range. You specify additional network numbers or ranges by placing each network on its own line. For example, a MACIPXGW.DAT file can contain the following command:

```
exclude  
  10-20
```

This directs the gateway to serve all AppleTalk networks *except* 10-20 and 100. Alternatively, the MACIPXGW.DAT file can contain the following command:

```
include  
  10-20
```

This directs the gateway to serve *only* AppleTalk networks 10-20 and 100, excluding all others.

NOTE: The network numbers in these examples are AppleTalk network numbers, not IPX network numbers.

If no MACIPXGW.DAT file is found in SYS:SYSTEM, the MacIPX gateway serves all AppleTalk networks.

To restrict gateway service to selected AppleTalk networks, complete the following steps:

1 Use a DOS ACSII text editor to create a file called MACIPXGW.DAT in the router's SYS:SYSTEM directory.

2 Place commands in the file using the following syntax:

```
[exclude | include]
  <net_number >[-net_number ]
  . . . ]
```

3 Save and close the file.

4 To put your changes into effect, enter

```
UNLOAD MACIPXGW
REINITIALIZE SYSTEM
```

4 Managing

This chapter describes how to monitor IPX™ LAN and WAN connections using the available router management consoles and utilities.

Using the IPXCON Utility

IPXCON is an NLM utility that provides access to statistics and information about the status of various components of the IPX protocol. IPXCON uses SNMP over IPX or UDP/IP to monitor remote servers, routers or network segments.

To launch IPXCON, enter **LOAD IPXCON** at the system console prompt or load NIASCFG and follow this path:

Select View Status for NIAS > Protocols and Routing > IPX Protocol Stack

You can use IPXCON to perform the following tasks:

- ♦ Monitor and troubleshoot IPX routes and network segments throughout your IPX internetwork
- ♦ Display the status of any IPX router or network segment on your internetwork
- ♦ Display all paths through which IPX packets can flow
- ♦ Locate all active IPX routers on your internetwork
- ♦ Display operational circuits for IPX
- ♦ Monitor remote routers running Novell® Internet Access Server router software

Using the IPXPING Utility on the Server

The IPXPING.NLM program enables you to check connectivity to an IPX server on your internetwork.

The IPXPING utility sends a packet request to the target node, an IPX server or workstation. After the target node receives the packet, it sends an IPXPING reply packet to the system that sent the request packet.

To use IPXPING, type a command similar to the following at the server prompt:

```
load ipxping
```

The system displays the New Target window. The fields of the New Target window are described in [Table 5](#).

Table 5 IPXPING New Target Window

Field	Description
Network	Select a target IPX server by entering its internal IPX address.
Node	Select a target IPX server by entering the target node number. You must enter both the internal IPX address and node number to select the server.
Seconds to pause between pings	Specify the number of seconds between each packet transmission.

After entering the network address, node address, and number of seconds to pause between pings, press `Esc` to start sending request packets. The sending node continues to send request packets and collect response time statistics until you press `Esc` again and exit IPXPING.

Request and reply packets use the same format; each packet contains the standard IPX header.

To select additional IPX servers, press `Insert`. Enter the IPX address of the server in the Address field. Press `Esc` to start sending packets.

Using the IPXPING Utility on the Workstation

The program is a DOS utility that determines the time to transport IPX packets to a specific server and back. IPXODI (or an equivalent IPX interface) must be loaded before IPXPING can be used. The time is displayed in milliseconds (although resolution is 1/18 second). The user can specify the number of pings to send and the interval between sends. A summary of the high, low, and average times is displayed when the program terminates. The program terminates when the acknowledgment to the last ping is processed or when the user presses Ctrl +C or Break .

This topic contains the following sections:

- ◆ [Syntax](#)
- ◆ [Parameters](#)
- ◆ [Example](#)

Syntax

```
IPXPING <netaddr> [/s=<size>] [/r=<repeat>] [/d=<delay>]  
          [/c=<char>]
```

All parameters except netaddr are optional. If no parameters or an invalid parameter is entered, help is displayed. The parameters are case-insensitive. A minus (-) sign can be used instead of a slash (/).

Parameters

The netaddr parameter is the internal network address of the server to ping. It is an eight-digit hexadecimal value. The server must support the IPX ping protocol.

The size parameter is the number of data bytes (excluding the IPX PING header) in the ping data packet. The default size is 100 bytes. The size is limited by the maximum packet size supported by the driver. If an invalid size is specified, a packet overflow error occurs.

The repeat parameter is the number of times to send the ping. The default is 1.

The delay parameter is the number of seconds to delay between successive pings. The default is 1; you can set the value to 0.

The char parameter is the character to fill the ping buffer with. The default character is P.

Example

To use IPXPING.EXE, type the following command at the DOS prompt:

```
ipxping 2e64afe3
```

The following information is displayed:

```
IPXPING: estimated time to send a 0.5K packet to 2E64AFE3 is  
166 ms.
```

```
IPXPING: sent ping packet number 1.
```

```
IPXPING: packet 1 response received in 166 ms.
```

```
IPXPING: 1 send, 1 received, low 166 ms high 166 ms, average  
166 ms.
```

NOTE: If you don't receive a response to the final ping, press Ctrl+C to terminate the program. Otherwise, IPXPING will wait until it receives a response.

Using the SPFCON Utility

The SPFCON utility enables you to monitor Sequenced Packet Exchange™ (SPX™) spoofing statistics. *Spoofing* is the process of preserving the transport end point connection by imitating keep alive packets and responding to watchdog request packets without passing this traffic across on-demand WAN links. Using SPX spoofing can help you maintain lower costs over on-demand WAN links.

NOTE: SPX spoofing is implemented only on PPP interfaces at this time.

To launch SPFCON, enter the command **load spfcon** at the NetWare console server.

Four windows make up the SPFCON user interface. The four windows are

- ◆ **Main Window**
- ◆ **Interfaces Window**
- ◆ **Connections Window**
- ◆ **Spoofing Statistics Window**

Main Window

The Main Window offers two options:

- ◆ Interfaces enable users to view the spoofing statistics on a per interface basis
- ◆ SPX spoofing enables users to view the spoofing statistics of all SPX connections

Interfaces Window

The Interfaces Window displays a list of WAN interfaces. At this time, only PPP interfaces will be included on this list. You can select either one or all listed interfaces. When you select an interface, the Connections Window is displayed.

Connections Window

The Connections Window displays all connections for the selected interface, either SPX or NCP. The Connections Window shows each connection's interface name, source node, destination node, source ID, and destination ID.

Spoofing Statistics Window

The Spoofing Statistics Window displays the SPX and NCP spoofing statistics for a selected connection. [Table 6](#) describes the fields on this window that require explanation.

Table 6 Spoofing Statistics Fields

Field	Meaning
Spoofing State	Initial indicates that spoofing has not yet started; active indicates that spoofing has started and is active; inactive indicates that spoofing has ended because Remaining Spoofing Time expired.
Number of Spoofing Started	Indicates the number of times spoofing has been started for this SPX connection.
Remaining Spoofing Time	Indicates the time remaining in this spoofing session; when this time expires, the spoofing state will change to inactive.
Sequence Numbers	Shows the sent and received sequence number.

Field	Meaning
Acknowledge Numbers	Shows the sent and received acknowledged number.
Allocation Numbers	Shows the sent and received allocation number.
Spoofing packets	Shows the number of packets spoofed.
Keep Alive Packets Dropped	Shows the number of keep-alive packets dropped during this active spoof session.

NOTE: In the spoofing statistics fields, Sequence Numbers, Acknowledge Numbers, Allocation Numbers, and Keep Alive Packets Dropped don't apply for NCP™.

Viewing NetWare IPX Configuration Information

To see how Internetwork Packet Exchange™ (IPX™) is configured, load IPXCON and select the following options:

- ◆ IPX Router Information—To view whether the RIP, SAP, and NetWare Link Services Protocol™ (NLSP™) protocols are configured for this router.
- ◆ NLSP Information—To view NLSP configuration for the system and the NLSP network.
- ◆ Circuits—To view the IPX circuits configured for the router.

Determining Whether a Remote IPX Router Is Reachable

To determine whether a remote router is reachable, you can run an IPX Echo test. To run an Echo test, load IPXPING and perform the following steps:

- 1** Specify the target router address in the Network field.
- 2** Specify the target router node number in the Node field.
- 3** Specify the number of seconds between each transmission in the Seconds to pause between pings field.
- 4** Press Esc to begin transmitting.

IMPORTANT: To run the IPX Echo test, both the machine originating the echo packets and the machine responding to the echo packets must support IPXPING. To support IPXPING, both machines must have IPXRTR.NLM, which is included in NetWare 4.1 software and Novell Internet Access Server 4.1 software.

Determining Which IPX Services Are Reachable

If you want to know whether a specific IPX service is available, you need to find out which IPX services are reachable. To determine the services available to a router, load IPXCON and follow this path:

Select Services > service you want to reach to get information about

If a service is visible, it is reachable. If you want to see information about a service, select the service and press Enter. The Service Information window displays Name, Type, Network Number, Node, Socket, and additional path information for NLSP systems under the selection Destination Information.

Checking an IPX Network for Inactive Routers

You can use IPXCON to identify routers on your network that are inactive—that is, not routing—for some reason. This information can often help you locate a defective network interface. To check your network for inactive routers, load IPXCON and follow this path:

Select NLSP Information > Routers

IPXCON lists the NLSP routers known to the system you are monitoring.

Any router labeled Unreachable might either be down or have a defective network interface. This labeling might also indicate that some other router in the path has one of these problems.

Any router labeled Overloaded has run out of memory and can no longer process NLSP routing information.

Routers not labeled Unreachable or Overloaded are operating properly.

Checking the IPX Routing Table

To check the IPX routing table and information associated with each route, load IPXCON and follow this path:

Select Forwarding > Display entire forwarding table

The Forwarding Table window shows you all known IPX destination networks. The list shows the following information about each item:

- ◆ Network number of the destination
- ◆ Routing protocol through which the destination was learned

- ◆ First hop circuit to the destination
- ◆ Name of the destination

The Destination Information window expands on this by showing information about the Next Hop (Name, Circuit, NIC Address), and by giving access to the services on the destination router that are available to the current router.

If the destination is an NLSP destination, you can determine potential paths to the destination by selecting Potential Paths.

Checking an IPX Network for Duplicate Network Numbers

Each external network number that identifies a LAN on your IPX network must be unique. Incorrect configurations and other problems can cause two LANs to have the same external network number.

One way in which this can occur is when a bridge connecting two LAN segments fails—a condition known as a *split LAN*. When the bridge fails, each segment becomes a separate LAN but retains the same network number. As a result, routers forwarding packets to that network number see two destinations and simply choose the nearest one.

To check your network for duplicate network numbers, load IPXCON and follow this path:

Select NLSP Information > LANs

The LANs are listed in numeric order by network number. To look for duplicate network numbers, scroll through the list of LANs.

Checking an IPX Network for Duplicate System IDs

To check your network for duplicate system IDs, load IPXCON, select a system, and select the following path:

Select NLSP Information > System Information > field associated with Detailed NLSP System Information

Check the numbers associated with Sequence Number Skips; if the number is increasing, two or more NLSP routers on your network have the same system ID.

Determining Where NLSP Is Running in Your Network

If you have migrated your network to NLSP, you can identify which LANs on your IPX network are using NLSP or RIP—or both—as the routing protocol.

If you partitioned your IPX network into routing areas, note the following points:

- ♦ The procedure in this section shows only the LANs and NLSP routers that are within the same routing area as the system you are currently monitoring.
- ♦ RIP should be running only on the routers at area boundaries.

To find LANs on which NLSP or RIP is running, load IPXCON, select a system, and follow this path:

Select NLSP Information > LANs

The Known LANs window displays the following information about each NetWare LAN of which the local system is aware:

- ♦ Network Number— External network number of the LAN
- ♦ Throughput— Number of Mbps reported by the LAN board
- ♦ Delay— Time, in microseconds, required for packets to reach the LAN

If NLSP is importing RIP routes to a LAN, the entry is labeled RIP Active. If the entry is labeled Unreachable, the LAN is no longer accessible from the local system. If there is no label, then NLSP is the only routing protocol running on the LAN and the LAN is reachable.

You can select one of the LANs to see which NLSP routers are on it.

Finding NLSP Routers with Insufficient Memory

To check your network for routers that have run out of memory, load IPXCON and follow this path:

Select NLSP Information > Routers

If a router does not have enough memory to process routing information, the entry is labeled Overloaded.

You can also select System Information and check the following information:

- ♦ Level 1 Overloaded—It should read No. If it reads Yes, the router does not have enough memory to process routing information.
- ♦ Detailed NLSP System Information—The Level 1 Database Overloads field indicates how many times the router has run out of memory.

Finding the Designated Router on a LAN

The Designated Router is an NLSP router elected by its peers to represent and keep track of the connectivity of its LAN. The Designated Router handles exchanges of link state information on behalf of all other NLSP routers on the LAN. Only broadcast (LAN) circuits have Designated Routers.

To find the NLSP Level 1 Designated Router on a LAN, load IPXCON, select a system, and follow this path:

Select Circuits > a broadcast circuit

The Circuit Information window appears and displays the name, type, and state of each circuit.

The Circuit Information window displays, among other information, the name of the Designated Router and the external network number of the LAN it represents.

To see more information about the Designated Router, complete the following steps:

- 1** Record the name of the Designated Router.
- 2** Press Esc until you return to the Available Options window.
- 3** Select SNMP Access Configuration.
- 4** Select IPX as the Transport protocol.
- 5** Type the name of the Designated Router in the Host Address field, then press Enter .
- 6** Press Esc to return to the Available Options window.

In a few moments, IPXCON begins displaying statistics for the Designated Router.

Monitoring Error Counters

Error counters are monitored to make sure they are not increasing rapidly, because a rapid increase indicates a problem. For information about troubleshooting these problems, refer to [Chapter 5, “Troubleshooting,” on page 125](#). You can monitor error counters for IPX interfaces in the following ways:

- ◆ By using MONITOR to view counters such as Checksum Errors, Send and Receive Packet Errors, and interface-specific errors. To view these counters, load MONITOR and follow this path:

Select LAN/WAN Information > interface you want to view

- ◆ By using PPPCON for WAN connections to view the following PPP counters:
 - ◆ Bad Address Fields
 - ◆ Bad Control Fields
 - ◆ Bad FCS Values
 - ◆ Packets Too Long

To view these counters, load PPPCON and follow this path:

Select PPP Interfaces > interface you want to view > PPP Error Statistics

- ◆ By using IPXCON to view the following IPX counters:
 - ◆ Too Many Hops
 - ◆ Header Errors
 - ◆ Unknown Sockets
 - ◆ Decompression Errors
 - ◆ Malformed Requests
 - ◆ Compression Errors
 - ◆ Open Socket Failures
 - ◆ Maximum Sockets

To view these counters, load IPXCON and follow this path:

Select IPX Information > Detailed IPX Information

Viewing the MacIPX Gateway Configuration

To view information about the configuration and operation of a MacIPX[®] gateway, enter the following command at the server prompt:

```
LOAD MACIPXGW SHOW=YES
```

This command does not reload the MacIPX gateway, but instead displays information about the MacIPX gateway and the AppleTalk networks that it serves, as in the following example:

```
MACIPXGW:
Unicast threshold set at 1.
AppleTalk nets this gateway is configured to serve:
10-20      111      2222-2223      3333-3335
AppleTalk nodes registered for IPX broadcasts:
IPX node: 0xffffffffffff
      Socket: 0x452
              10.238      1501.138      1502.168
      Socket: 0x453
```

The information in the preceding example includes the following items:

- ◆ The unicast threshold
- ◆ The network numbers of all AppleTalk networks served by this gateway
- ◆ All AppleTalk nodes currently registered with the MacIPX gateway for broadcasts and identified by the IPX socket

Viewing MacIPX Gateway Statistics

To view MacIPX gateway statistics, load MONITOR and select the following parameter path:

```
Select LAN/WAN Information > MACIPXGW
```

A screen displays the statistics explained in [Table 7 on page 123](#).

Table 7 MacIPX Gateway Custom Statistics

Statistic	Explanation
Received Tickle Packets	Number of tickle packets sent by MacIPX clients served by the gateway. MacIPX clients send tickle packets to the MacIPX gateway; the gateway sends IPX broadcast packets back to the clients.
IPX Broadcast Requests from IPX Stack	Number of IPX broadcast packets sent to the MacIPX gateway by the IPX stack in the NetWare server +79595or Novell router.
IPX Broadcast Requests from MacIPX Clients	Number of IPX broadcast packets sent to the MacIPX gateway by the MacIPX clients that the MacIPX gateway is servicing.
DDP Packets Broadcasted for IPX Broadcasts	Number of AppleTalk packets sent out as broadcast packets carrying IPX broadcast packets.
DDP Packets Unicastd for IPX Broadcasts	Number of AppleTalk packets sent out as unicast packets carrying IPX broadcast packets.
Received DDP Packets with Unknown Options	Number of AppleTalk packets received by the gateway that include unrecognized encapsulation demultiplexing options. This indicates corrupted packets or incompatible client software. Ensure that your network cabling is working correctly and that software on the Macintosh clients is compatible with this version of the MacIPX gateway.
Received DDP Packets with Wrong Type	Number of AppleTalk packets received by the gateway that include an incorrect AppleTalk packet type. This indicates the presence of corrupted packets or incompatible client software. Ensure that your network cabling is working correctly and that software on the Macintosh clients is compatible with this version of the MacIPX gateway.
Received Service Requests	Number of requests received by the gateway to provide service to MacIPX clients.
Transmitted Service Grants	Number of times the gateway granted service to MacIPX clients.
Transmitted Service Refusals	Number of times the gateway refused service to MacIPX clients.
Memory Allocation Failure	Number of times the gateway could not allocate memory. You might need to add memory to the NetWare server or Novell router to fix this problem.

5

Troubleshooting

This section discusses IPX troubleshooting information that is divided into four categories:

- ◆ Troubleshooting tools
- ◆ Configuration tips
- ◆ Troubleshooting checkpoints
- ◆ Common problems

If a problem that is general in nature occurs, the procedure described in [“Troubleshooting Checkpoints” on page 127](#) will help you isolate and resolve the problem. If a problem with a specific symptom occurs, refer to [“Common Problems” on page 131](#).

Troubleshooting Tools

The IPX-specific troubleshooting tools are explained in the following sections:

- ◆ [“IPXCON” on page 125](#)
- ◆ [“System Console Commands” on page 126](#)

IPXCON

IPXCON is a NetWare[®] Loadable Module™ (NLM™) utility that provides access to statistics and information about the status of various components of the IPX™ protocol. It uses SNMP to access this information from any local or remote system on the network. IPXCON operates over IPX and TCP/IP networks, and uses the User Datagram Protocol (UDP) to run over the

networks. For more information on IPXCON, refer to [Chapter 4, “Managing,”](#) on page 111.

Enter **LOAD IPXCON** at the system console prompt. You can use IPXCON to perform the following tasks:

- ◆ Monitor and troubleshoot IPX routes and network segments throughout your IPX internetwork
- ◆ Display the status of any IPX router or network segment on your internetwork
- ◆ Display all paths through which IPX packets can flow
- ◆ Locate all active IPX routers on your internetwork
- ◆ Display operational circuits for IPX
- ◆ Monitor remote routers running the Novell Internet Access Server 4.1 routing software

System Console Commands

Specific commands are available from the system console prompt that prove useful in troubleshooting IPX connection problems. The following commands are examples:

- ◆ **DISPLAY SERVERS**

This command lists all known NetWare server names and the number of hops (IPX routers that must be crossed) to reach each server. This information is similar to the information shown in the IPXCON forwarding table, but it is less comprehensive than IPXCON.

- ◆ **DISPLAY NETWORKS**

This command shows the IPX network number, the number of hops needed to reach the network, and the estimated time, in ticks (1/18 of a second), for a packet to reach a network. The number of known networks is shown at the end of the list. For NetWare servers, both the internal IPX network numbers and the cabling network numbers are displayed. This information is similar to the information shown in the IPXCON services table, but it is less comprehensive than IPXCON.

- ◆ **RESET ROUTERS**

This command resets the IPX routing table in the file server if the table has become inaccurate or corrupted.

- ◆ TRACK ON

This command displays three types of RIP and SAP information: *Server*, *Network*, and *Connection Requests*. This information is formatted according to whether the router/server is receiving the information (IN), broadcasting the information (OUT), or receiving a connection request. Refer to Utilities Reference for more details.

Configuration Tips

We recommend the following guidelines for configuring IPX:

- ◆ Each server and router in the entire internetwork must have a unique internal internetwork number.
- ◆ Each LAN in the entire internetwork must have a unique IPX network number, even if you have configured a PPP unnumbered WAN link that connects two LANs.
- ◆ Each WAN in the entire internetwork must have a unique internal internetwork number.

Troubleshooting Checkpoints

Observe the procedures described in the following sections when you are configuring IPX or NLSP for the Novell Internet Access Server 4.1 routing software:

- ◆ “IPX Checkpoints” on page 127
- ◆ “NLSP Checkpoints” on page 131

IPX Checkpoints

To isolate and resolve problems with IPX, complete the following steps:

- 1** Verify that workstations can connect to all desired servers.

If a problem with LAN connectivity occurs, refer to “[IPX Connectivity Problems \(Duplicate ID or Network Number\)](#)” on page 129. If you are using the NetWare Mobile IPX™ software, refer to “[NetWare Mobile IPX Client Loses Connectivity to the Server](#)” on page 151.

- 2** Verify that the IPX network number is different for each LAN across a WAN link.

IMPORTANT: Each LAN segment must have a unique IPX network number. It is a common error to incorrectly use the same IPX network number on each side of an unnumbered PPP or WAN link.

- 3** Verify that all servers and routers in the entire internetwork have unique internal network numbers.

In addition, each network segment has a unique network number, and all servers and routers on the same segment must have their interfaces configured with the same IPX network number.

- 3a** Look at the routing table in IPXCON on any NetWare Link Services Protocol™ (NLSP™) system in the suspected area (parameter path: Select NLSP Information > Routers).

Determine whether there are routers that appear and disappear from the table (these routers might also become unreachable for brief periods of time). Then establish a remote connection with RCONSOLE and check for any error messages indicating duplicate internal network numbers.

Typically, if you can log in to a server, but cannot establish a connection to the same server using RCONSOLE, then the server is configured with a duplicate IPX address.

- 3b** Enable SAP on one of the interfaces in a router in the NLSP area.

Between SAP periodic updates, you should see two routers (or more, if more than one router has the same internal network number) being listed as unreachable and then reachable. This should occur every 5 to 10 seconds.

- 4** After you have identified the NLSP systems with the problem, load IPXCON to determine to which networks the servers or routers are connected (parameter path: Select NLSP Information > Routers).
- 5** Select the NLSP router that is the source of the problem. You might need to select the router several times because its connectivity is intermittent.

If the router is a Novell router, then the internal network number is probably a duplicate.
- 6** Change one of the router's internal network numbers and restart the system.
- 7** For WAN links, verify that third-party routers use IPXWAN™ software (RFC 1362, 1551, or 1634).

To establish an IPX connection to third-party routers over a WAN, the third-party routers must support IPXWAN; otherwise, problems with initiating, maintaining, or terminating the IPX connection occur.

- 8** Verify that the IPX network number is different for each WAN link unless an unnumbered RIP is used (in which case, the IPX network number is zero).

IPX Connectivity Problems (Duplicate ID or Network Number)

To isolate and resolve IPX connectivity problems, complete the following steps:

- 1** Find the server or router that is connected to the same segment as the workstation.

If more than one server or router connects the workstation to the network, look at each system to determine if it has proper connectivity. If the system with which you are communicating is a server, then use that system.

- 2** Check the forwarding table on both systems.

2a If you are on the server or router representing network A, then find network B in the forwarding table. Load IPXCON (parameter path: Select Forwarding > Display entire forwarding table).

2b If network B is not displayed in the table, then exit and enter the forwarding table until it appears.

2c If network B still does not appear in the table, probably a router in the path either is malfunctioning or has a duplicate system ID.

2d If the route shows up intermittently, then probably a router in the path has a duplicate system ID.

2e If the route shows up consistently, then look at the route and select Potential Paths by selecting the network on the Forwarding Table menu.

- 3** Make sure that the potential path leads to the correct network by looking at the intermediate routers.

If a duplicate network number exists, you can determine the location of the duplicate number from this window.

3a If all the routers in the path seem to be correctly configured, then write down the addresses of all the LANs and routers in the path.

- 3b** Use IPXPING to check all routers or servers in the end-to-end path (do this from one side only). Also, check the end points.

If connectivity is occurring from a workstation, make sure that the workstation can log in to the first server or router in the path or that it has access through the router in question.

If the connectivity loss is only temporary (for example, you occasionally get abort retries on the workstation), then let IPXPING run for several minutes. Check for packet loss during this time, then examine the router at which packet loss occurred.

It is also possible that a router is malfunctioning in the end-to-end path. Usually, IPXPING can help you determine where the fault is occurring.

- 3c** Once you have found the router that has the problem, check its potential paths.

All downstream routes from the first router to the router that has the problem should also be potential paths on this router. If this is not the case and the router does not quickly acquire the downstream routes, then the system probably has a software error in it. Contact the router manufacturer's technical support for further assistance. To help minimize problems like this, you should purchase only NLSP-certified routers.

- 4** If connectivity loss occurs outside an NLSP area, check each router in the end-to-end path for an external RIP route.

RIP can be used between NLSP areas. Therefore, it is necessary to check the end-to-end path in a more tedious way, as follows:

- 4a** Find the next-hop router from each of the servers.

- 4b** Look at that system's forwarding database.

- 4c** Find the next-hop router from that system, and so on, until you have found where the route leads. Do the same from the other side of the path as well.

This process is difficult with the current implementation of RIP and SNMP for Novell, because RIP shows only the next-hop LAN and Network Interface Card (NIC) address (over LANs) instead of the internal network number of the system. SNMP cannot receive packets that are addressed to a NIC; the packets must be addressed to the internal network number. You must work backward, first finding all routers attached to the LAN, and then finding the receiving LAN

card on each router. One of the routers you should start with is the next-hop router. Repeat these steps until you find the destination network number. If you do not find a duplicate network number in either direction, check each link in the path for errors.

NLSP Checkpoints

To isolate and resolve problems with NLSP, complete the following steps:

- 1** Determine connectivity.
 - ◆ Verify that all neighbors are displayed under the NLSP Neighbors option in IPXCON. This will determine whether there is local connectivity.
 - ◆ Verify that there are sufficient potential paths within each area.
 - ◆ Verify that all LANs are listed in the NLSP LANs table in IPXCON.
 - ◆ Verify that all NLSP routers are listed in the NLSP Routers table in IPXCON.
- 2** Determine whether RIP is active.
 - ◆ Verify that the NLSP LANs window indicates that RIP packets are being absorbed.
 - ◆ Verify that the Circuits table indicates the state of any system.

Common Problems

This topic discusses the following common problems and their potential solutions:

- ◆ [Login Times Out \(page 132\)](#)
- ◆ [Load Balancing over IPX Is Not Working \(page 133\)](#)
- ◆ [Only One IPX Packet Is Sent and Received \(page 133\)](#)
- ◆ [IPXCON Counters Are Increasing \(Duplicate ID or Network Number\) \(page 133\)](#)
- ◆ [Error Messages Are Displayed \(Duplicate ID or Network Number\) \(page 134\)](#)
- ◆ [NLSP Decision Process Is Running Frequently \(Duplicate System ID\) \(page 135\)](#)

- ◆ Other Router Names Are Not Displayed (page 135)
- ◆ System Frequently Appears and Disappears on the LAN (page 135)
- ◆ Multiple Systems on a LAN Become Unreachable Intermittently (page 138)
- ◆ Connectivity Across a Point-to-Point Link Has Been Lost (page 139)
- ◆ An NLSP Server on a LAN Cannot Be Accessed (page 140)
- ◆ LAN Is Partitioned (page 141)
- ◆ No Communication Occurs between Two Networks (page 142)
- ◆ Services Are Inaccessible in the Area (page 143)
- ◆ Number of Routes and Services on a System Shows Local Connectivity Only (page 144)
- ◆ Services or Routes are Fluctuating Excessively (page 145)
- ◆ Heavy Network-Layer Traffic Occurs on a Point-to-Point Link (page 146)
- ◆ Applications Perform Poorly (page 147)
- ◆ CALLMGR Shows an IPX Circuit but IPXCON Does Not (page 149)
- ◆ Single System Is Entering an Overloaded State (page 149)
- ◆ Many Systems Are Entering an Overloaded State (page 149)
- ◆ Connectivity Is Lost on Only One LAN (page 150)
- ◆ NetWare Mobile IPX Client Loses Connectivity to the Server (page 151)
- ◆ Re-establishing the Connection (page 151)

Login Times Out

The remote IPX LAN number might have a static route on the dial-in side. Load STATICON and select Dynamically Configure Static Routing Tables to dynamically configure local and remote routing tables. To initiate dynamic configuration with a remote router, the call to that router must be Connected. Press **Ins** to use the Make Call option to attempt to make a call to a currently Not Connected WAN destination.

Load Balancing over IPX Is Not Working

Add the following line to the workstation's AUTOEXEC.BAT file:

```
Set loadbalance local lan = on
```

Verify that the router is configured to use NLSP compatible with RIP/SAP. Set maximum path splits to 8.

Only One IPX Packet Is Sent and Received

If only one IPX packet is sent and received each time the routing software attempts to establish a connection, decrease the user data size value so that it is equal to or less than the size used by the frame relay switch. Also, ensure that the user data size is equal to or less than the physical receive packet size.

IPXCON Counters Are Increasing (Duplicate ID or Network Number)

- ◆ The Link State Packets (LSPs) Received counter in IPXCON is increasing (parameter path: Select NLSP Information > System Information > Detailed NLSP System Information).

Two NetWare servers or routers have conflicting internal network numbers and both systems are in the same NLSP area, resulting in a duplicate NLSP system ID. A number of activities occur when this situation exists. If router A and router B have the same system ID, both routers attempt to assert that they own the system ID. First, router A issues LSPs that supersede router B's LSPs and purges any LSPs of router B that it does not have. Then, router B does the same to router A. It is possible that this increases the amount of LSP traffic in the network considerably, particularly if either router A or router B has many LSPs (for example, if either router is importing many routes and services).

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the node that is causing the problem, refer to [“IPX Connectivity Problems \(Duplicate ID or Network Number\)” on page 129](#).

- ◆ The Sequence Number Skips counter in IPXCON is increasing on both routers (parameter path: Select NLSP Information > System Information > Detailed NLSP System Information).

It is normal for a system to have some sequence number skips, but the Sequence Number Skips value should not increase after the first five

minutes of a router's operation, unless there is a duplicate NLSP system ID.

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the node that is causing the problem, refer to [“IPX Connectivity Problems \(Duplicate ID or Network Number\)” on page 129](#).

- ◆ The Own LSP Purges counter in IPXCON (parameter path: Select NLSP Information > System Information > Detailed NLSP System Information) is increasing on both routers.

There is a duplicate NLSP system ID, and many systems have fluctuating counts of routes and services because the services available through one or the other router become unreachable. If one of the systems in question can route, then all systems in the network are running the NLSP decision process frequently.

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the malfunctioning node, refer to [“IPX Connectivity Problems \(Duplicate ID or Network Number\)” on page 129](#).

Error Messages Are Displayed (Duplicate ID or Network Number)

- ◆ The console displays the following message:

```
Router name has the same internal network number of number
but a system ID of system_ID .
```

A duplicate internal network number has resulted in a duplicate system ID.

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the malfunctioning node, refer to [“IPX Connectivity Problems \(Duplicate ID or Network Number\)” on page 129](#).

- ◆ The console displays the following message:

```
LSP graph inconsistency detected in stored LSP from system
name length number . There has been a memory corruption
or software error.
```

Cause 1 —Errant application is corrupting the NLSP graph or LSP database.

Contact technical support.

Cause 2 —NLSP has a software error that is either corrupting the graph or causing the graph to be represented incorrectly.

Contact technical support.

NLSP Decision Process Is Running Frequently (Duplicate System ID)

The NLSP decision process is running frequently.

To observe this symptom, obtain access to a NetWare server or router running NLSP in the network and enter **SET ISUL DEBUG=256** at the system console prompt. Every time the decision process runs, an entry is displayed at the system console. If the decision process runs at least every 30 seconds, there might be a duplicate system ID.

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the malfunctioning node, refer to [“IPX Connectivity Problems \(Duplicate ID or Network Number\)”](#) on page 129.

Other Router Names Are Not Displayed

If other router names do not display when the DISPLAY SERVERS command is used, the NLSP Local Area Addresses might be different. Load NIASCFG and select Configure NIAS > Protocols and Routing > Network Interfaces > IPX > IPX Expert Configuration option to set the IPX network number and area mask to zeros.

System Frequently Appears and Disappears on the LAN

A system frequently appears and disappears on the LAN.

Cause 1 .—System is not transmitting its packets to the Designated Router properly.

- ♦ At the Designated Router, check the Neighbor State Changes option in IPXCON (parameter path: Select Circuits > a specific circuit > Detailed Circuit Information) and monitor it for the circuit (LAN).
- ♦ If the number is increasing but there are no new systems on the network and systems are not being bound and unbound or restarted, then a local connectivity problem probably exists.

- ◆ After you have determined that there is a local connectivity problem, check whether there are any systems in the Initializing state in IPXCON (parameter path: Select NLSP Information > Neighbors). You might need to exit and enter the Neighbors window several times over several minutes. Also check for any systems that do not have names associated with them. Any system without a name has not transmitted its LSP to the Designated Router.
- ◆ If you find any system that enters the Initializing state, then you have identified a connectivity problem between that system and the Designated Router. If you have not identified any such system, then it is still possible for a router to lose connectivity occasionally. Select the entry in the NLSP Neighbor table of the router that does not have connectivity (parameter path: Select NLSP Information > Neighbors). There is an initial holding time for the system. By default, every system sends a Hello packet on a LAN every 15 to 20 seconds. You can see whether the Designated Router is receiving all the Hello packets from the system by comparing the packets sent to the packets received.
- ◆ If you still cannot determine the problem with the system, and if you have the routing software located on the LAN, then load IPXPING at the console prompt. Set the PING send rate to zero and check to determine whether packets are being dropped. A packet dropped every once in a while should not cause concern; however, if more than 1 percent of the packets are dropped, there is a problem with the router or server. The problem could be caused by the software or hardware. To determine whether there is a problem with the software, restart the PC. If the problem continues, install a new interface board.
- ◆ Check connectivity between the Designated Router and another system. The Designated Router might be dropping packets or be the source of the problem.

Cause 2 —Problem with the underlying media.

Look at LAN/WAN information in MONITOR and check for errors. Errors are specific to the media; therefore, press F1 (for online help) to see what different errors mean. Most errors indicate that there is a problem with the server's or router's network interface board. These errors could be caused by the software or hardware. To determine whether there is a problem with the software, restart the PC. If the problem continues, install a new interface board.

Cause 3—System misconfiguration.

By default, NLSP timers are set so that a system becomes unreachable when three packets are dropped. Look at the system's configuration to ensure that this setting has been used.

Cause 4—One or both of the systems are dropping packets.

Check the interface boards to determine whether packets are being dropped because of insufficient Event Control Blocks (ECBs). Increasing the maximum number of physical receive packets might help stop the system from dropping packets. To reach this option load install, select NCF file options > edit startup.ncf. Increase the maximum number of physical receive packets to at least 1524. However, the system might be incapable of handling the system load. In this case, increase the processor power of the system that is dropping packets, or reduce the load on the server by either removing NLM files or decreasing the number of users on the system. You can determine whether the system is using too much CPU processing power by using MONITOR and viewing utilization.

Cause 5—NLM on the server is not relinquishing control of the CPU frequently enough.

This is a rare occurrence. To determine whether this is occurring, select Performance in MONITOR. Look for processes that exceed many millions of cycles per iteration. You can also determine that an NLM is malfunctioning by removing the NLM from the server and observing whether the problem is resolved.

Cause 6—Internal error in NLSP.

If you have exhausted all other possibilities, you should document your system configuration, number of users, and error frequency and send a copy of the system configuration (SYS: SYSTEM\CONFIG.TXT), including the NLSP configuration file (usually located in \ETC\NLSP.CFG), to technical support.

Multiple Systems on a LAN Become Unreachable Intermittently

If multiple systems on a LAN become unreachable intermittently, the Designated Router might be the source of the problem.

Cause 1—Designated Router does not have enough system memory to represent the LAN.

To determine whether this is the case, you can check whether the Designated Router is overloaded when the problem occurs or whether it has been overloaded in the past. Check the Level 1 Data Base Overloads statistic in IPXCON (parameter path: Select NLSP Information > System Information > Detailed NLSP System Information). If the Designated Router does not have sufficient memory to represent the LAN, then the LAN loses connectivity as new systems are added. You must then select another system on the network to become the Designated Router by increasing its priority in NIASCFG and issuing the REINITIALIZE SYSTEM command. In addition, you must add memory to the system that was overloaded. You might also want to check whether other systems in the network are overloaded. Refer to [“Many Systems Are Entering an Overloaded State.”](#)

Cause 2—Designated Router or some other system on the LAN is causing the network outage. Another system can cause this problem by electing itself as the Designated Router on the LAN.

A single malfunctioning system on a LAN might be causing all systems on the LAN to become unreachable intermittently. Using IPXCON, check the Designated Router Changes counter for all systems on the LAN (parameter path: Select Circuits > a specific circuit > Detailed Circuit Information). If a single system has a large value displayed in the Designated Router Changes counter, the system probably has connectivity problems with other systems on the LAN. Check whether the counter increases over time. If the Designated Router is not being restarted or unbound from a LAN, the counter should not increase. If the Designated Router Changes counters of all systems on the LAN are increasing, the system that should be the Designated Router probably has a connectivity problem. To determine whether the problem is particular to the system or to the network itself, remove the system from the LAN or decrease its Designated Router priority.

Cause 3—Two systems are contending to be the Designated Router for the LAN.

In this case, the Own LSP Purges counter increases (parameter path: Select NLSP Information > System Information > Detailed NLSP System

Information). However, unless you are using ARCnet* or some other media that does not have IEEE addresses, only one system has the highest priority on the LAN (the MAC address is used as a tie breaker and IEEE addresses are unique). If necessary, change the priority on one of the contending routers.

Connectivity Across a Point-to-Point Link Has Been Lost

You cannot bring up an IPX point-to-point link, but IP is working.

Cause 1 —System on the other end of the link does not support IPXWAN, or IPXWAN is not supported over the media that you use.

Contact the router manufacturer to verify that its product supports IPXWAN.

Cause 2 —Link has excessive errors.

Cause 3 —One of the IPXWAN implementations has an error.

Load MONITOR and view LAN/WAN Information under a specific NIC or LAN adapter. Determine whether the link has excessive errors by viewing discrepancies in packet error counts.

Issue the SET ISLL DEBUG=ON command and capture the IPXWAN exchanges. Contact the manufacturer of the router that appears to be in violation of the IPXWAN specification.

Cause 4 —Link has excessive errors.

Load MONITOR and view LAN/WAN Information under a specific NIC or LAN adapter. Determine whether the link has excessive errors by viewing discrepancies in packet error counts.

Cause 5 —Timers are misconfigured, causing the link to drop packets. (Typically, the defaults are used for PPP.)

Check in MONITOR under Driver Statistics to determine whether this is the cause. Set the timers so that the values match those set on the remote node.

Cause 6 —System is limited by the amount of memory or by the capacity of the CPU or bus.

Load MONITOR and view memory utilization to determine if the capacity of the CPU or memory is limiting the function.

Cause 7 —Link itself is corrupting data.

If the link is corrupting data, the corrupt LSPs statistic increases. To check this, load IPXCON (parameter path: Select NLSP Information > System Information > Detailed NLSP System Information). Even a single corrupt LSP indicates a serious problem because LSPs are transmitted infrequently. Note that this counter is a global counter and it is possible that some other media is corrupting the data link.

An NLSP Server on a LAN Cannot Be Accessed

- ◆ You cannot access a server on the LAN.

Cause 1 —Area address is set to the wrong value.

In this case, the number of destinations (known networks and services) implies local connectivity only. Also, the Initialization Failures statistic in IPXCON increases (parameter path: Select Circuits > Detailed Circuit Information). If you have configured area addresses, make sure that all systems that should be in communication have the same area addresses. It is acceptable for systems to have different addresses, if that is the desired configuration.

Cause 2 —RIP is not enabled.

If you are running multiple areas on a LAN and are using RIP to interconnect the systems, verify that RIP is enabled on those servers that are interconnecting areas. Load NIASCFG and select Configure NIAS > Protocol and Routing > Bindings > IPX Binding > Expert Bind Options > RIP Bind Options > RIP State. If the RIP State is set to Auto in NIASCFG, there is a small chance that RIP will fail. To avoid having RIP fail, set RIP State to On.

Cause 3 —Hub or bridge has failed.

Use IPXPING to check whether there is data-link connectivity.

Cause 4 —LAN board driver does not support multicast, even though the driver's documentation claims that it does.

Use a driver that supports multicast or set the MAC Channel option to Broadcast (parameter path: Select Bindings > a specific binding > Expert Bind Options > NLSP Bind Options).

Cause 5 —LAN board has failed. This system does not see other systems on a LAN, or it does not have any adjacencies in the Up state. However, it declares itself as being attached to the LAN.

In IPXCON, look for a system that is in the Initializing state. If a system on the LAN appears in the Initializing state on all other systems but has no neighbors itself, then the system can send but not receive. Check that system, particularly if it is the Designated Router. Use IPXPING to help determine the actual source of the problem. An interface board with a conflicting interrupt is a common source of this problem.

Cause 6—System is declaring itself the owner of the LAN, even though it is not the owner.

Reinitialize the system.

Cause 7—LAN has become partitioned temporarily during normal NLSP operation. This should occur only during an NLSP system's startup, and the error should be corrected within a few minutes.

Verify that the condition does not persist. If it does, check for a hardware problem or an NLSP software incompatibility with other systems.

- ◆ Clients running UnixWare software, OS/2 Named Pipes, and NetWare/IP™ software have problems with connectivity.

The systems are not properly configured. Refer to [Chapter 1, “Understanding,” on page 11](#) and [Chapter 3, “Setting Up,” on page 51](#) for information about the solution.

LAN Is Partitioned

A LAN is partitioned when there are duplicate LANs or one of the LANs is declared unreachable in IPXCON (parameter path: Select NLSP Information > LANs).

Cause 1—A hub or bridge has failed.

Use IPXPING to check whether there is data-link connectivity.

Cause 2—LAN board has failed. This system does not see other systems on a LAN, or it does not have any adjacencies in the Up state. However, it declares itself as being attached to the LAN.

In IPXCON, look for a system that is in the Initializing state (parameter path: Select NLSP Information > Neighbors). If a system on the LAN appears in the Initializing state on all other systems but has no neighbors itself, then the system can send but not receive. Check that system, particularly if it is the Designated Router. Use IPXPING to help determine the actual source of the problem.

Cause 3—System is declaring itself the owner of the LAN, even though it is not the owner.

Reinitialize the system.

Cause 4—LAN has become partitioned temporarily during normal NLSP operation. This should occur only during an NLSP system's startup, and the error should be corrected within a few minutes.

Verify that the condition does not persist. If it does, check for a hardware problem or for NLSP software incompatibility with other systems.

No Communication Occurs between Two Networks

- ◆ Connectivity is lost because a router is missing required routes. This router's table is not consistent with the routing tables of other routers.

The router has not converged because it is configured for multicast and the driver does not support multicast, even though the driver's documentation claims that it does. Set the MAC Channel option to Broadcast (parameter path: Select Bindings > a specific binding > Expert Bind Options > NLSP Bind Options).

- ◆ A workstation cannot communicate with a server on a different connecting LAN, but other systems on the LAN can communicate.

Cause 1—Physical problem on the workstation (for example, a broken LAN card).

Replace the malfunctioning hardware.

Cause 2—Packet filter that is discarding the system's packets has been implemented somewhere in the network.

Check each intervening router and correct the filter configurations.

- ◆ A server cannot communicate with another server on a different connecting LAN, but other systems on the LAN can communicate.

Cause 1—Server is misconfigured.

Correct the server's configuration and make sure that it has connectivity by verifying that it has the routes and services typical for your network.

Cause 2—Connectivity exists, but it is so poor that the transports above IPX cannot maintain connectivity.

Set up an IPXPING test between the two systems. If the rate of dropped packets is high, the connectivity problem is probably caused by a

malfunctioning link between the two networks. To determine which link has the problem, refer to “[Applications Perform Poorly](#)” on page 147.

Cause 3—There is a duplicate network number. This can cause a duplicate system ID, provided that both systems are in the same area and the duplicate network numbers are two internal network numbers on two NetWare implementations of NLSP. The console probably displays the following message:

```
System server_name with internal network number number
has my system ID in it.
```

Change the internal network number of one of the conflicting systems, or remove one of the systems from the network immediately. For information about how to find the malfunctioning node, refer to “[IPX Connectivity Problems \(Duplicate ID or Network Number\)](#)” on page 129.

Cause 4—Packet filtering has been implemented on a router. This can cause symptoms similar to those caused by duplicate network numbers (for example, network A might be visible from network B, but network B is not visible from network A).

If communication does not occur within an NLSP area, it is usually easy to determine whether the network fault is caused by packet filtering or a duplicate network number. Use IPXCON to view duplicate LAN network numbers (parameter path: Select NLSP Information > LANs). If your system does not have a matching network number, use FILTCFG to remove the packet filtering.

Services Are Inaccessible in the Area

Services are inaccessible in the area.

Cause 1—Services are being blocked by filters.

Examine the IPXCON Services option of each router in the path to isolate the router that is filtering the services.

Cause 2—Network connectivity problems.

Check that the network to which the service is attached exists. If the network does not exist, look for link connectivity in the path between the area in the network that is missing the network number and the area that is generating the service.

Cause 3—Service name conflict. This occurs when you have the same service name and the same type (for example, file server). If the service is a file service, then the user logging in might not have appropriate rights and, consequently, the login is rejected.

- ◆ Use the Services option in IPXCON to verify that the service has the appropriate network number.
- ◆ If the service and network numbers are both visible, refer to “[No Communication Occurs between Two Networks](#)” on page 142.

Cause 4—Under rare circumstances, the server from which you are logging in has insufficient space to store the service in the bindery.

Increase the disk space on the file server.

Cause 5—If there are many services, a third-party router might be unable to transmit the entire SAP table before the next periodic update. A third-party router can start transmitting the services again from the beginning of the table, instead of completing the current update.

Contact technical support.

Number of Routes and Services on a System Shows Local Connectivity Only

The number of routes and services on a system shows that there is local connectivity only.

Cause 1—If the network to which the system is attached is an NLSP- only network, and if the system is not configured for RIP mode only, then the system might not receive RIP updates. If the following message is displayed at the system console, there is a RIP mode misconfiguration and the two routers cannot communicate:

```
Router server_name claims network number is really number
```

- ◆ Turn on the tracking screens and check to determine whether there are any RIP and SAP updates being sent to the server. If there are, it is possible that you need RIP on the network but that you have set the RIP State option to Off. You can check this by loading IPXCON (parameter path: Select Circuits > select a circuit name). Check the RIP information to see whether you have accidentally set RIP State to Off.
- ◆ If you should be receiving RIP on the network but are not, it is possible that the other routers on the network have been configured with RIP State

set to Off. If this is the case, and if you want to run RIP between the servers, then set RIP State to On.

Cause 2 —Two NLSP systems are configured with different area addresses. In this case, the Initialization Failures counter in IPXCON should be increasing (parameter path: Select Circuits > Detailed Circuit Information).

If you are on an area boundary and you are using RIP as the interarea protocol, configure RIP on both systems on the interface through which they are communicating.

Cause 3 —One of the routers has the RIP State option set to Auto. If you intended to use RIP for interarea routing on the network, this condition is potentially serious. If the two NLSP routers are in communication with each other, they continue to run RIP. However, if they are connected together with a bridge or hub and that hardware fails and is brought up again sometime later, NLSP does not detect the condition and RIP does not turn on again.

If you are running RIP and SAP on the network and the routes and services fluctuate, refer to “**Services or Routes are Fluctuating Excessively.**”

Services or Routes are Fluctuating Excessively

The number of routes and services are fluctuating excessively. Some fluctuation is normal in a large network. Change is occurring constantly as systems are brought down for maintenance and other reasons. However, hundreds of routes and services appearing and disappearing indicates a network error. All the following problems are solved with the same procedure:

Cause 1 —Misconfiguration with RIP and SAP.

Cause 2 —Problem with a link.

Cause 3 —Error in NLSP.

Cause 4 —Two systems are competing for a system ID.

Cause 5 —LAN is generating errors.

- ◆ First try to determine which systems are affected. It might be helpful to determine which services are appearing and disappearing and to trace the paths backward to their location. If many services are appearing and disappearing, it might work better to find the set of affected systems. In either case, you should be able to determine the boundary routers.

In NLSP, this determination is easy to make. Simply look at an affected system's neighbors to see whether they have the problem, too. Examine each system as you move outward. If just a single NLSP system is experiencing the problem, then the cause is probably a local connectivity problem. Check the system for neighbor state changes and disappearing services.

- ◆ Try to determine whether the problem has a boundary. If you see a system in which one set of routers or services is disappearing and see an adjacent system in which a different set of routers or services is disappearing, you probably have found the boundary. Check to determine whether any of these systems has problems with its neighbors, has fluctuating links, or has a duplicate system ID.
- ◆ If you find that the problem is isolated to a LAN, then follow the procedures described under [“Multiple Systems on a LAN Become Unreachable Intermittently” on page 138](#). You might find that one system that is supplying routing information to the network is losing its neighbors, particularly the Designated Router. This system is probably the source of the problem.
- ◆ If you find that the problem is caused by some interaction with RIP and SAP, check all systems for consistency on the LAN regarding the following configurable RIP and SAP parameters in NIASCFG: Packet Size Override, Periodic Update Interval, and Aging Interval Multiplier (parameter path: Select Bindings > a specific interface > Expert Bind Options > RIP (or SAP) Bind Options).

You should never let a periodic multiplier be less than 4. Because of timer skew, this means that after three packets containing the same route are dropped, the system loses this route. If any of the previously listed RIP and SAP parameters are different, you must reconfigure the values. Better still, use the default values, especially on LANs. Using different values for the timers is too risky to justify the savings on a LAN.

Heavy Network-Layer Traffic Occurs on a Point-to-Point Link

- ◆ There is heavy Network-layer traffic on a point-to-point link and you are using RIP and SAP.
Migrate to NLSP.
- ◆ There is heavy Network-layer traffic on a point-to-point link and you are using NLSP.

Cause 1 —Network-layer packets are being retransmitted because there is a software error or because two important timers are misconfigured in NIASCFG. The important timers are the Minimum Non-Broadcast LSP Transmission Interval timer (which indicates the amount of time before an LSP is retransmitted when there is no acknowledging Partial Sequence Number Packet [PSNP]) and the Partial SNP Interval timer (parameter path: Select Configure NIAS > Protocols and Routing > Protocols > Expert Configuration Options > NLSP Convergence Rate Configuration).

The latter timer should be set to a value much smaller than the value set for the former timer because it acknowledges LSPs; if it is set too high, the LSP transmitter responds as if the LSP is lost and retransmits the LSP. If the problem persists after you reconfigure the timers, call technical support.

Cause 2 —Many changes are occurring in your network because there is too much RIP activity in your network.

Migrate more of your network to NLSP.

Cause 3 —Some systems are sending too many updates.

- ◆ If possible, migrate more of your network to NLSP.
- ◆ If you cannot migrate more of your network to NLSP, find the boundary routers (those that are importing RIP) by looking at the NLSP LANs window. Increase the Maximum LSP Generation Interval option on these systems.

Applications Perform Poorly

You are experiencing poor application performance on systems in your network.

Cause 1 —Suboptimal path has been selected by NLSP.

- ◆ Check the end-to-end path of the connection and make sure that the links that you thought would be chosen are being used to forward data. In an NLSP area, look at the Potential Paths window by loading IPXCON (parameter path: Select Forwarding > a specific destination). Outside an NLSP area, perform the procedure described in [“IPX Connectivity Problems \(Duplicate ID or Network Number\)” on page 129](#).
- ◆ If you find an incorrect path caused by RIP, then you can increase the RIP cost by manually changing the cost of the RIP link. Refer to "Configuring RIP and SAP" for information about how to do this.

Cause 2—Application relies on ticks to retransmit its packets. This should not happen with the routing software, but it is possible that some other manufacturer's router does not comply with the ticks value.

If this is the case, increase the cost of the routing software to match the value of the router in question. This procedure should not affect other paths much, but it should help to stop the application from retransmitting packets.

Cause 3—Router is in an NLSP area and you have routers with load sharing enabled. This causes the application to retransmit packets needlessly.

If this is the case, turn off load sharing to see whether the situation improves.

Cause 4—Link speed is too slow. You might be choosing the optimal path, but throughput is still not adequate.

- ♦ If you have a Novell router and the protocol is windowed, you might want to enable the IPX Header Compression option or experiment with PPP data compression, if it is being used. However, the application might require more bandwidth than you have available.
- ♦ Make sure that the problem is not caused by latency. Compression adds latency, which can slow down protocols that do not have windowing, such as the Sequenced Packet Exchange™ (SPX™) protocol. Also, older versions of the NetWare shells do not have windowing. If you experiment with the Packet Burst™ protocol, you might be able to reduce latency and increase throughput.
- ♦ If this is an X.25 problem, you might be able to remove some of the latency by increasing the X.25 window size, the physical frame size, or both.
- ♦ If the previous suggestions do not work, change the type of line that you have. Some kinds of frame relay lines have relatively low latency, as do leased lines. X.25 and other WAN technologies sometimes have high latency. If an X.25 problem exists, you can remove some of the latency by increasing the X.25 window size, the physical frame size, or both. Over slow PPP lines, increasing the frame size continually can hurt performance. Because this can cause packets to be retransmitted, it can take a long time to transmit a single frame across the link.

Cause 5—Malfunctioning routers in the end-to-end path or a link that is causing problems in the end-to-end path.

Determine the routers in the end-to-end path and check each router and link for abnormal behavior.

Cause 6 —Load sharing is enabled between dissimilar paths.

Verify that the two paths have comparable media and data rates.

CALLMGR Shows an IPX Circuit but IPXCON Does Not

CALLMGR shows an IPX circuit but IPXCON does not.

IPXCON does not show a circuit until after IPXWAN has completed negotiation.

Check the link for errors, and make sure that both sides of the IPX link are implementing IPXWAN properly.

Single System Is Entering an Overloaded State

A single system is in an overloaded state.

Cause 1 —System is running out of memory.

Cause 2 —Another system is experiencing database overload. It is possible that another system in an overloaded state is causing your system to go into an overloaded state.

On each suspect system, use MONITOR to check whether the Alloc Memory Pool is set too low. If the value is too low, increase the value set for the Alloc Memory Pool or remove some applications.

Cause 3 —Transient condition on that router.

If the problem persists, the system is running out of memory. If the value is too low, increase the value set for the Alloc Memory Pool or remove some applications.

Many Systems Are Entering an Overloaded State

Many systems are entering an overloaded state.

Cause 1 —Systems are being overrun with routing information. Possibly the number of systems in the NLSP area has exceeded the number you originally intended.

Using IPXCON, look at the number of routers in the area (parameter path: Select NLSP Information > Routers). Add this number to the number of LANs in the area. This sum is a good indicator of the amount of memory that is

required by an NLSP area. We recommend that you do not exceed 400 LANs and routers (total) in any single NLSP area. It is also possible that two areas have merged when they should not have. Determine whether routers are in the area that should not be there. Prevent the areas from merging or use area addresses.

Cause 2 —Backbone has been imported multiple times into the NLSP area. NLSP is careful about the way that it imports external routes and services into the NLSP network. For example, only the Designated Router on a LAN imports information. Usually, if two NLSP systems are connected to the same RIP backbone but they are on different LANs, a conflict does not occur. If RIP reports two different routes to the same location, only the RIP route with the shortest hop count is imported into the NLSP network. However, it might be that most of the backbone is imported more than once. This can occur if there is more than one equal cost path from the NLSP network to the RIP network. To determine whether most of the backbone has been imported more than once, look carefully at all routers that are importing RIP.

Using IPXCON, find the systems that are importing RIP into the NLSP area, then determine whether RIP Active is displayed for the LAN (parameter path: Select NLSP Information > LANs). Find the Designated Router in each LAN. Look at the Forwarding table of each Designated Router. If more than one Designated Router is on the LAN, this is probably because you have turned off NLSP on the routers. In this case, you must run NLSP between the routers to reduce the amount of imported information on each LAN.

Connectivity Is Lost on Only One LAN

Connectivity is not possible on a single LAN, but it is possible on other LANs.

This almost certainly indicates a duplicate network number. Refer to [“No Communication Occurs between Two Networks” on page 142](#) for the solution. Usually, the network number that is the duplicate does not have the connectivity.

NetWare Mobile IPX Client Loses Connectivity to the Server

The NetWare Mobile IPX client loses connectivity to the server.

If you lose connectivity before you start an operation, you will see messages such as `Access Denied`, or it might look as if access is not available on your network drives. If you lose connectivity after you start an operation, you will receive a DOS critical error message that asks you whether you want the

operation to abort, retry, or fail. The method of reestablishing the connection depends on whether you lose connectivity before or after you start an operation. This method is the same for each of the following causes.

Cause 1 —The NetWare Mobile IPX client was out of range of wireless coverage for too long.

Return the client to the range of wireless coverage and reestablish the connection as explained in [“Re-establishing the Connection” on page 151](#).

Cause 2 —The driver used with NetWare Mobile IPX is not Network Event Service Layer (NESL) aware, so IPX is not sent the receive notification of data-link events.

Use only NESL-aware drivers. Reestablish the connection as explained in [“Re-establishing the Connection” on page 151](#).

Cause 3 —The wireless board was removed when a process was running in the background or the board was not correctly plugged in. This can happen when you swap boards.

Simply reinsert the wireless or Personal Computer Memory Card International Association (PCMCIA) board. Reestablish the connection as explained in [“Re-establishing the Connection” on page 151](#).

Cause 4 —The PCMCIA board was swapped when the portable was in a low-power state. This tends to confuse the card and socket services, and events are not sent to the drivers. This invariably causes computer lockup.

Do not swap the PCMCIA board when the portable is in a low-power state. Reestablish the connection as explained in [“Re-establishing the Connection” on page 151](#).

Re-establishing the Connection

If you lose connectivity before you start an operation, you can usually reestablish a connection by selecting Open/Save. If the HR Time To Live timer has expired, selecting Open/Save will not reestablish the connection and you must log in again.

If you lose connectivity after you start an operation, reestablish your connection as follows:

- ◆ In the Windows environment, you are asked to either retry or cancel the current operation. Selecting Cancel terminates the connection to the server you were accessing, and you must log in again to reestablish a

connection. If you want the operation to be completed, you must select the Retry option when you return to the network. In the meantime, the system is unusable. If the HR Time To Live timer has expired, selecting Retry will not reestablish the connection and you must log in again.

- ◆ In the DOS environment, you have three options: Abort, Retry, or Fail. Abort and Fail both terminate the connection to the server you were accessing, and you must log in again to reestablish a connection. If you want the operation to be completed, you must select the Retry option when you return to the network. In the meantime, the system is unusable. If the HR Time To Live timer has expired, selecting Retry will not reestablish the connection, and you must log in again.