

IPX Router Specification

September 2, 1992

Revision A

Part Number: 107-000029-001

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of this manual, 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.

© Copyright 1992 by Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express prior written consent of the publisher.

Novell, Incorporated
122 East 1700 South
Provo, Utah 84606

Trademarks

Novell has made every effort to supply trademark information about company names, products, and services mentioned in this document. Trademarks indicated below were derived from various sources.

ARCNET is a trademark of Datapoint Corporation.

Ethernet is a trademark of Digital Equipment Corporation, Intel Corporation, and Xerox Corporation.

IBM, NetBIOS, and Token-Ring are trademarks of International Business Machines Corporation.

Intel is a trademark of Intel Corporation.

Novell, NetWare, and Internetwork Packet Exchange (IPX), are registered trademarks of Novell, Inc.

Omninet is a trademark of Corvus Systems, Inc.

Xerox is a registered trademark of Xerox Corporation.

Table of Contents

Chapter 1: Introduction

Purpose	1-2
NetWare Communication Protocols	1-2
Medium Access Protocols	1-3
Internetwork Packet Exchange (IPX)	1-4
Routing Information Protocol (RIP)	1-4
Service Advertising Protocol (SAP)	1-5
NetBIOS	1-6
Summary	1-6

Chapter 2: Internetwork Packet Exchange (IPX)

Introduction	2-2
Packet Format	2-2
Router Implementation Guidelines	2-6
Packet Delivery	2-7
Sending Node's Responsibility	2-7
Router's Responsibility	2-9
Mixed Topology Routing	2-10
Burst Mode Protocol	2-10
Routing Information Gathering	2-10

Chapter 3: Routing Information Protocol (RIP)

Introduction	3-2
Packet Format	3-3
RIP Operations	3-5
Router Implementation Guidelines	3-6
Selecting the Best Route	3-7
Routing Information Broadcasts	3-8
Split-Horizon Algorithm	3-8
Router Initialization	3-9
Router Shutdown	3-12
Router Information Maintenance	3-13
RIP Aging	3-14
Lost Route Algorithm	3-15
RIP Request Handling	3-15
RIP Interpacket Gap	3-16
RIP and SAP Interoperability	3-16

Chapter 4: Service Advertising Protocol (SAP)

Introduction	4-2
Packet Format	4-3
SAP Operations	4-6
Router Implementation Guidelines	4-7
Server Information Table	4-7
Router Initialization	4-8
Router Shutdown	4-11
Server Information Maintenance	4-11
SAP Aging	4-12
SAP Request Handling	4-13
SAP Interpacket Gap	4-14
RIP and SAP Interoperability	4-14

Chapter 5: IPX Type 20 Propagation Packet

Introduction	5-2
Packet Format	5-2
Router Implementation Guidelines	5-3

Appendix A: RIP and SAP Bandwidth Requirements A-1

Index	Index-1
-------------	---------

Preface

This document describes the network protocols, processes and algorithms that govern routing in the native NetWare environment.

The document is divided into the following chapters:

Chapter 1 provides an introduction to the NetWare environment and the protocols involved in NetWare routing.

Chapter 2 describes the Internetwork Packet Exchange (IPX) protocol and provides guidelines for implementing IPX in a router.

Chapter 3 describes the Routing Information Protocol (RIP).

Chapter 4 describes the Service Advertising Protocol (SAP).

Chapter 5 describes the special IPX Type 20 packet and provides guidelines for the handling these packets require by an IPX router.

Appendix A describes RIP and SAP bandwidth requirements.

Manual Conventions

All numbers in this document are decimal unless otherwise specified. Hexadecimal numbers are identified by a trailing 'h' (i.e. FFh). Where bit fields within a byte are specified, bit 0 is assumed to be the low-order bit.

Note that numeric fields composed of more than 1 byte can be in one of two formats: high-low or low-high. High-low numbers contain the most significant byte in the first byte of the field, the next most significant byte in the second byte, and so on, with the least significant byte appearing last. Low-high numbers are stored in exactly the opposite order. The Intel microprocessors store numbers in low-high order. Unless otherwise noted, all numeric fields described in this document are in high-low format.

Introduction

Purpose	1-2
NetWare Communication Protocols	1-2
Medium Access Protocols	1-3
Internetwork Packet Exchange (IPX)	1-4
Routing Information Protocol (RIP)	1-4
Service Advertising Protocol (SAP)	1-5
NetBIOS	1-6
Summary	1-6

Purpose

When networks were simple one-trunk systems with a single file server, the communication processes that facilitated interaction between workstations and file servers were not a serious issue in the network's design. More time was spent trying to get single-user software packages to run in a multiuser environment than worrying about internetworking issues.

Since then, networks have become larger and more complex. Products such as multiprotocol routers and MAC-layer bridges and protocol implementations such as Token Ring source routing and TCP/IP are becoming necessary elements in the design of large internetworks. To successfully integrate these and other elements with NetWare and IPX/SPX, a solid understanding of the mechanics of NetWare's communication environment is essential.

This document examines closely the processes, protocols, and algorithms that govern NetWare routing. It gives a detailed description of the various protocols involved and examines the actual mechanics of packet routing and the administration of routing and server information in the native NetWare environment.

Netware Communication Protocols

Any in-depth discussion of the mechanics of NetWare communications and routing requires a good understanding of the various protocols NetWare uses. These protocols define the format of the packets that are used to exchange information.

Basically, seven different protocols are used for communications with the native NetWare environment. These include:

- Medium Access Protocols
- Internetwork Packet Exchange (IPX)
- Routing Information Protocol (RIP)
- Service Advertising Protocol (SAP)
- NetWare Core Protocol (NCP)
- Sequenced Packet Exchange (SPX)
- NetBIOS

Additional protocols also exist that may be used in the NetWare environment, however only those protocols applicable to routing are described in this document.

Figure 1.1 displays the general architecture model for the NetWare protocol stack. As in the case of the Open Systems Interconnection (OSI) model, the upper NetWare protocols (RIP, SAP, NCP and SPX)

rely on the lower protocols (IPX and the medium access protocols) to take care of the lower-level communication issues. This modularized approach is the key to NetWare's compatibility with numerous network interface cards.

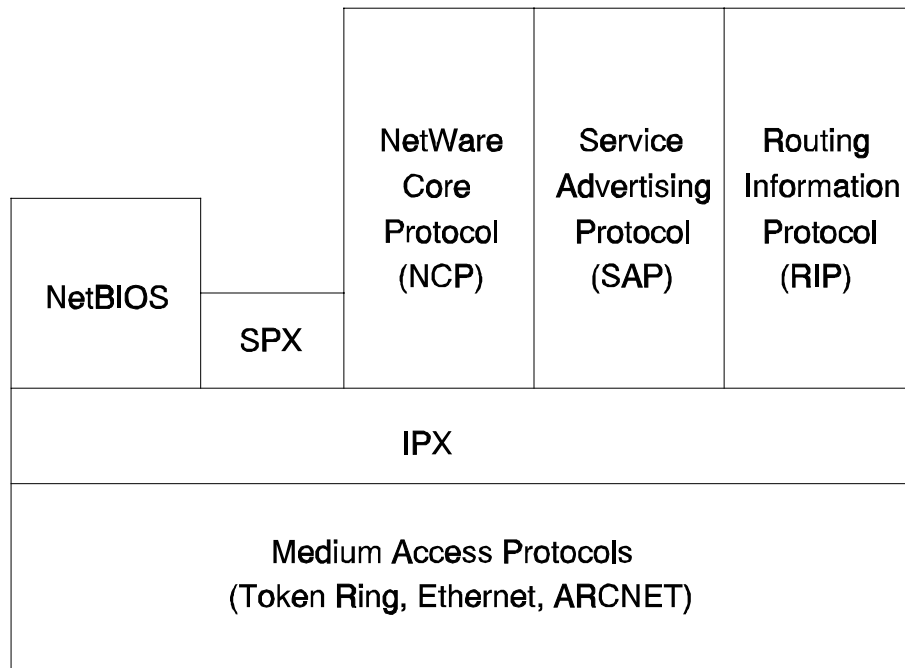


Figure 1.1 NetWare Protocol Model

Each of the protocols displayed in Figure 1.1, with the exception of NCP and SPX, play some part in the design of NetWare routers. These protocols are described briefly here; detailed descriptions can be found in the chapters that follow.

Medium Access Protocols

A number of medium access protocols have been defined, many of which can be used with NetWare. Among the most common medium access protocol implementations in existence are Token Ring 802.5, Ethernet 802.3, and ARCNET. The 802.x protocols have been defined by the Institute of Electrical and Electronic Engineers (IEEE). ARCNET was developed by Datapoint, Inc. As detailed specification documents currently exist which deal with medium access protocol implementations, this document provides only introductory information.

Medium access protocol implementations are primarily concerned with transporting packets from one node to another on a single network segment. They provide bit-level error checking in the form of a frame check sequence (FCS) which is appended to every packet transmitted, thus assuring that the large majority of the packets successfully

received will be free of corruption.

These implementations also define the addressing that distinguishes each node on a network. Node addressing is implemented within the hardware of each Network Interface Card (NIC). To get a packet to the proper node on a network, a Medium Access Control (MAC) header is placed at the front of every packet. This header contains source and destination node address fields to indicate where the packet originated and where it is going. Each NIC checks the destination address in the MAC header of every packet sent on the network. If the destination address matches the NIC's own node address, or if the packet is identified as a broadcast packet intended for all nodes, the NIC copies the packet. The packet is usually then passed up to a higher layer process (such as an IPX process) for further examination.

Internetwork Packet Exchange (IPX)

Novell adapted IPX from Xerox Network System's (XNS) Internet Datagram Protocol (IDP). IPX is a datagram, connectionless protocol that does not require an acknowledgment for each packet sent. Any packet acknowledgment, or connection control, must be provided by protocols above IPX.

Whereas the medium access protocols define only node addressing, IPX defines internetwork and intranode addressing schemes. IPX internetwork addressing is based on network numbers which are assigned to each network segment on a NetWare internetwork. The IPX intranode address comes in the form of socket numbers. Since several processes are normally operating within a node, socket numbers provide a sort of mail slot so that each process can distinguish itself to IPX.

Routers interconnect different network segments and by definition are network layer devices. In other words, routers receive their instructions for forwarding a packet from one segment to another from a network layer protocol. IPX, with the help of RIP and SAP which are introduced next, performs these Network layer tasks. These tasks include addressing, routing, and switching information packets to move single packets from one location to another on an internetwork.

A detailed discussion of IPX as well as guidelines for implementing IPX in a router can be found in chapter 2.

Routing Information Protocol (RIP)

Like IPX, RIP was derived from XNS. However, a change in the packet structure prohibits the straight integration of NetWare's RIP with other undeviating XNS implementations.

The Routing Information Protocol facilitates the exchange of routing information on a NetWare internetwork. IPX routers use RIP to create and dynamically maintain a database of internetwork routing information.

RIP allows a router to exchange routing information with a neighboring router. As a router becomes aware of any change in the internetwork layout, this information is immediately broadcast to any neighboring routers. Routers also send periodic RIP broadcast packets containing all routing information known to the router. These broadcasts keep all routers on the internetwork synchronized and provide a means of "aging" those networks which might become inaccessible due to a router going down abnormally.

Besides providing information to routers, RIP allows a workstation to locate the fastest route to a distant network. A full description of RIP is given in chapter 3.

Service Advertising Protocol (SAP)

The Service Advertising Protocol is similar in concept to RIP. While RIP allows routers to exchange internetwork routing information, SAP provides routers and servers (which contain SAP agents) with a means of exchanging internetwork service information.

Through SAP, servers advertise their services and addresses. Routers gather this information and share it with other routers. This allows routers to create and dynamically maintain a database of internetwork service information. Clients on the network can determine what services are available on the network and obtain the internetwork address of the nodes (servers) where they can access those services. Clients require this information to initiate a session with a file server.

SAP allows a router SAP agent to exchange information with a neighboring SAP agent. As a router SAP agent becomes aware of any change in the internetwork server layout, this information is immediately broadcast to any neighboring SAP agents. SAP broadcast packets containing all server information known to the SAP agent are also sent periodically. These broadcasts keep all routers on the internetwork synchronized and provide a means of "aging" those servers which might become inaccessible due to a router or server going down abnormally.

Like RIP, SAP uses IPX and the medium-access protocols for its transport. A full description of SAP is given in chapter 4.

NetBIOS

NetBIOS was written for IBM by Sytek in 1984 and has become a de facto industry standard. NetBIOS performs tasks applicable to the Transport and Session layers of the OSI model.

An in-depth discussion of NetBIOS is not given in this document, as a detailed understanding of NetBIOS is not critical in understanding NetWare routing. However, in order for NetBIOS (and perhaps some other protocol implementations as well) to function in the NetWare environment, routers must allow a packet to be propagated on an internetwork. NetWare accomplishes this by defining a special IPX Packet Type of 20 (14h). These packets require unique handling on the part of a router. Guidelines for doing this are given in chapter 5.

Summary

NetWare routers shoulder other responsibilities besides just passing packets from one network segment to another. They must act as "good citizens", sharing the information that they have with the other routers on the network, as well as being centers for internetwork information on their respective segments.

Each of the protocols introduced in this document play a different role in allowing the router to perform its various functions. Medium-access protocols and IPX provide the addressing mechanism that allow packets to be delivered to a desired destination. RIP and SAP provide the means whereby routers gather internetwork information and share that information with other routers. The major components of an IPX router are summarized in Figure 1.2.

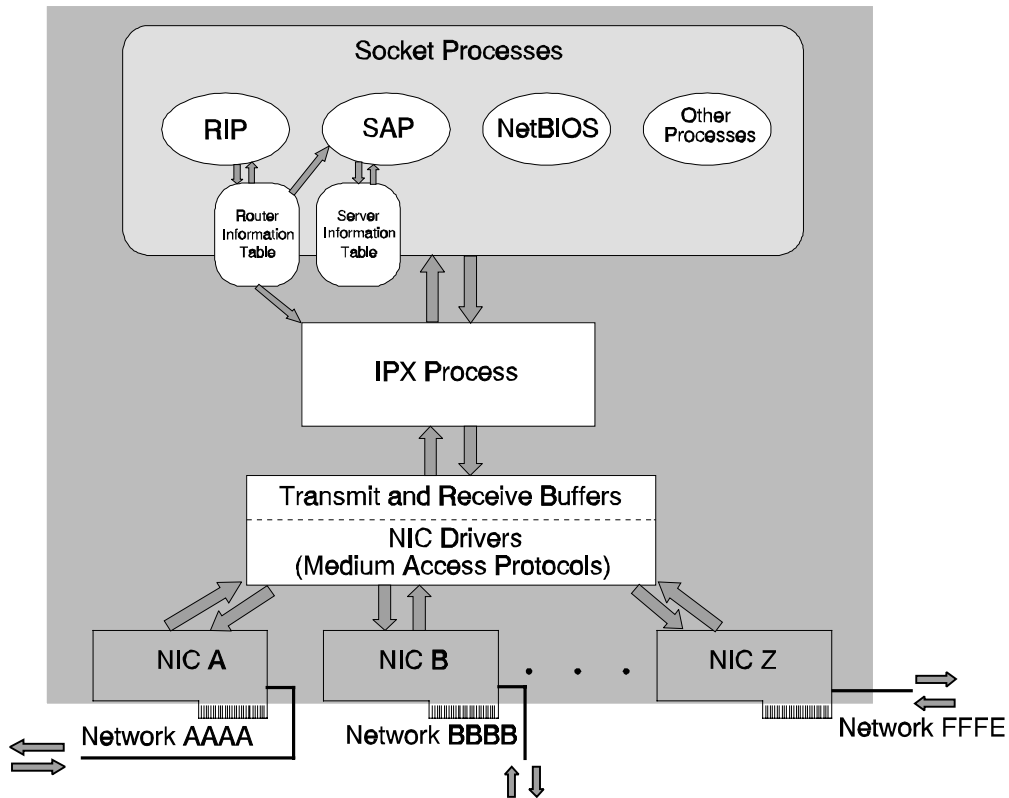


Figure 1.2 Internal Components of an IPX Router

Internetwork Packet Exchange (IPX)

Introduction	2-2
Packet Format	2-2
Router Implementation Guidelines	2-6
Packet Delivery	2-7
Sending Node's Responsibility	2-7
Router's Responsibility	2-9
Mixed Topology Routing	2-10
Burst Mode Protocol	2-10
Routing Information Gathering	2-10

Introduction

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

The tasks performed by IPX include addressing, routing, and switching information packets to move single packets from one location to another on the network. IPX defines internetwork and intranode addressing schemes, while relying on the network hardware for the definition of node addressing.

Network numbers are the basis of IPX's internetwork addressing. Each network segment on an internetwork must be assigned a unique network number. This network number is used by routers to forward packets to their final destination segment.

The IPX intranode address comes in the form of socket numbers. Since several processes are normally operating within a node, socket numbers provide a sort of mail slot so that each process can distinguish itself to IPX. As a process needs to communicate on the network, it requests that a socket number be assigned to it. Any packets that IPX receives that are addressed to that socket are passed on to the process. Hence, socket numbers provide a quick method of routing packets within a node.

Packet Format

The IPX packet is similar to a Xerox IDP packet. It consists of two parts: a 30-byte header and a data portion. The network, node, and socket addresses for both the destination and the source are held within the packet's IPX header.

The minimum IPX packet length, not including the length of the MAC header is 30 bytes (IPX header only). Historically, the maximum IPX packet size has been specified to be 576 bytes (IPX header plus data). While IPX implementations must be able to handle at least a 576 byte packet, enhancements to IPX now allow support for packet sizes up to 65,535 bytes. Media constraints will typically limit the actual maximum packet size allowed to something less than 65,535. Ethernet 802.3 packets, for example, are limited in size to 1500 bytes (not including

802.3 MAC header).

The IPX header is placed after the Media Access Control (MAC) header and before the packet data. (Packet data often includes the header of a higher-level protocol.) Figure 2.1 illustrates the structure of an IPX packet.

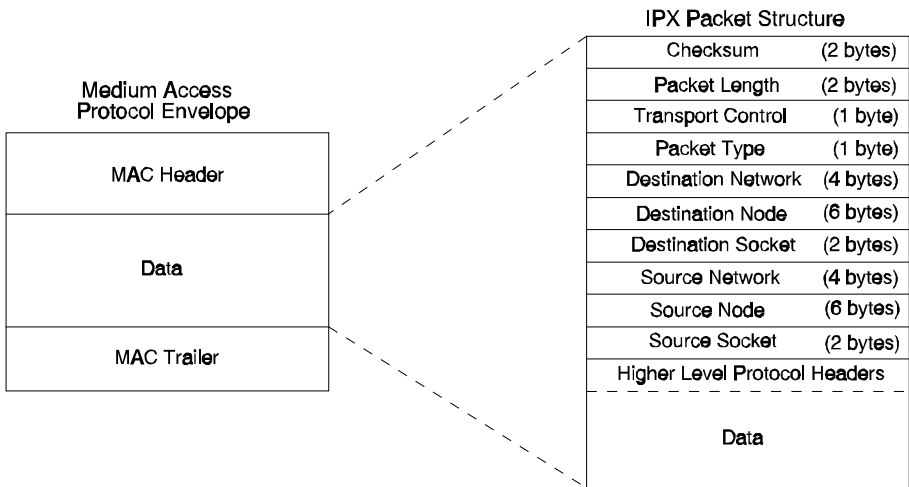


Figure 2.1 IPX Packet Structure

Checksum

Historically, Novell has not used checksums in the IPX header and has required that this field be set to FFFFh. However, there is no guarantee that checksums will not be used in future versions of NetWare. At present, routers should set this field to FFFFh when generating an IPX packet but should make no assumptions regarding what value this field will have for received packets.

Packet Length

This field contains the length of the complete internetwork packet, which is the length of the IPX header plus the length of the data.

Transport Control

This field indicates how many routers a packet has passed through on its way to its destination. Packets are discarded when this value reaches 16. Sending nodes always set this field to zero when building

an IPX packet. When a router receives a packet that requires further routing, this field is incremented by one.

Packet Type

This field indicates the type of service offered or required by the packet. Novell currently uses the following packet types:

0	Unknown Packet Type	(Used for all packets not classified by any other type)
1	Routing Information Packet	(see ch. 3)
4	Service Advertising Packet	(see ch. 4)
5	Sequenced Packet Protocol	(Used for SPX packets)
17	NetWare Core Protocol Packet	(Used for NCP packets)
20	Propagated Packet	(See ch. 5)

Note: For some versions of NetWare, the Packet Type field is not a reliable indicator of the type of packet encapsulated by the IPX header. The Source and Destination Socket fields should be used to determine the packet type, should this determination be required. (Propagated packets are an exception to this rule. Packet Type should be checked to determine if a packet is a propagated packet. See chapter 5.)

Destination Network

This field contains the network number of the network to which the destination node belongs. When a sending node sets this field to 0, the destination node is assumed to be on the same local network segment as the sending (or source) node. Internetwork routers should never set this field to 0 nor should they propagate packets which have this field set to 0.

A special case exists when a workstation sends SAP Get Nearest Server and RIP Get Local Target (or Route Request) broadcast requests at boot-up. Since the workstation does not yet know which network it belongs to it will set both the Source and Destination Network fields to 0 for these requests. When a router receives one of these requests it will send a reply directly to the sending workstation, filling in the Source and Destination Network fields with the appropriate network numbers. RIP and SAP requests/replies are described in detail in chapters 3 and 4.

There is no broadcast Network number (such as FFFFFFFFh) allowed in IPX.

Destination Node

This field contains the 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 specify its address, while a node on an Omninet network requires only 1 byte. If a physical network needs less than 6 bytes to specify a node address, the address should occupy the last of the 6 bytes of this field and the first bytes should be set to zero.

A node address of FF FF FF FF FF FF (6 bytes of value FFh) broadcasts the packet to all nodes on the destination network.

Destination Socket

This field contains the socket address of the packet's destination process. Sockets route packets to different processes within a single node. Novell has reserved several sockets for use in the NetWare environment:

File Servers		
451h		NetWare Core Protocol (NCP) Packet
Routers		
452h		Service Advertising Protocol (SAP) Packet
453h		Routing Information Protocol (RIP) Packet
Workstations		
4000h-7FFFh		Dynamic sockets; used for interaction with file servers and other network communications
8000h-FFFFh		Well-known sockets; assigned by Novell
455h		NetBIOS packets
456h		Diagnostics packets

Novell administers a list of sockets that are well-known in all NetWare environments. Software developers writing software for NetWare can obtain socket assignments by contacting Novell at 800-729-4357 or (512) 794-1795. Well-known socket numbers assigned by Novell begin at 8000h. Socket numbers above 8000h should not be used by an application unless they have been registered for that application with Novell.

There is no broadcast Socket number (such as FFFFh) allowed in IPX.

Source Network

This field contains the number of the network to which the source node belongs. If a sending node sets this field to zero, it means the local network to which the source is connected is unknown. In the case of

routers, the rules which apply to the Destination Network field apply to this field as well except that routers may propagate packets with this field set to zero after filling the field with the appropriate source address.

Source Node

This field is set to the physical address of the source node. Broadcast addresses are not allowed.

Source Socket

This field is set to the socket address of the process that transmits the packet. Processes communicating in a peer-to-peer fashion need not send and receive on the same socket number.

In a client/server situation, the server node 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 may originate from any socket number.

As in the case of destination sockets, these numbers can be static or dynamic. Source socket numbers follow the same conventions as those for destination sockets.

Router Implementation Guidelines

Routers interconnect different network segments and by definition are network layer devices. In other words, routers receive their instructions for forwarding a packet from one segment to another from a network layer protocol. IPX, with the help of RIP and SAP, performs these network layer protocol tasks for routers in the NetWare environment (IPX performs tasks outside the network layer as well). Figure 2.2 graphically illustrates the interaction of IPX with the other components in an IPX router.

NetWare-compatible routers may be configured to interconnect two or more segments. Each of these segments, however, must be assigned a unique IPX network number to distinguish it from other segments on the internetwork. A segment's network number must be configured into each of the routers connected to that segment. The network number serves as a common address for each node connected to a segment.

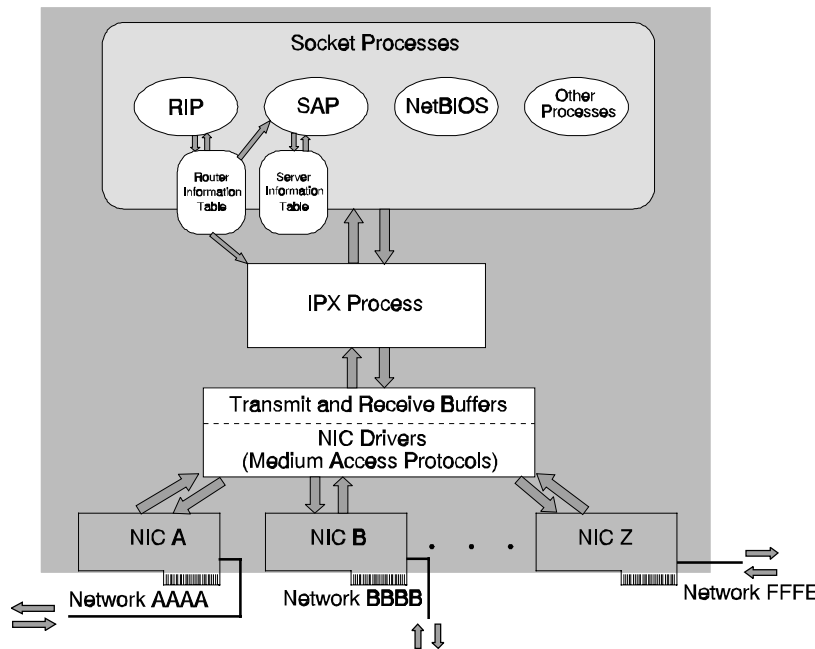


Figure 2.2 Internal Components on an IPX Router

Packet Delivery

On a NetWare network, the successful delivery of a packet depends on the proper addressing of the packet and the internetwork configuration (whether it is a single segment network or series of segments interconnected by repeaters, bridges and/or routers). The addressing of the packet is handled in its media-access protocol header and IPX header address fields. To send a packet to another node, the sending node must know the full internetwork address (network, node, and socket) of the node it desires to send to (the destination node). Once the sending node has the destination node's address it can proceed with addressing the packet. The way the MAC header of that packet is addressed, however, depends on whether the sending and destination nodes are separated by a router.


Sending Node's Responsibility

When a node wants to send information to another node with the same network number (both nodes are on the same segment), the sending node can simply send packets directly to the destination node. However, if the two nodes have different network numbers (reside on different network segments), the sending node must find a router on its own segment that can forward packets to the destination node's network segment.

To find this router, a workstation broadcasts a RIP packet requesting the fastest route to the destination node's network number (RIP requests/responses are discussed in detail in chapter 3). This RIP request is responded to by the router residing on the sending nodes segment with the shortest path to the desired segment; in the response, the router includes its network and node address in the IPX header. In the case when the sending node is a router rather than a workstation, the router will be able to get this information from its internal routing tables and need not send the RIP request.

Once the sending node knows the router's node address, it is prepared to send packets to the destination node. The sending node addresses these packets in the following way:

- It places the destination node's internetwork address (network, node and socket number) in the destination address fields of the IPX header.
- The sending node places its own internetwork address in the source address fields of the IPX header. (All other fields in the IPX header must be filled out as well.)
- The sending node places the node address of the router (the one that responded to the RIP request if the sending node is a workstation) in the destination address field of the MAC header.
- The sending node places its own node address in the source address field of the MAC header and sends the packet. (See Figure 2.3.)



Note: Only those fields in the IPX header which deal with packet addressing and routing are shown. However, all other fields must be set appropriately as described earlier in the chapter.

Figure 2.3 Packet Addressing Through a Router

Router's Responsibility

When a router receives a packet, its IPX process does the following:

- The IPX header Destination Address fields (net, node, socket) are checked to see if the packet needs to be handled by an internal socket process. Packets requiring such handling include any packets addressed to the router directly or broadcasts to any one of the router's directly connected network segments. Type 20 packets are also a special case which require unique handling on the part of a router. Chapters 3-5 provide guidelines for handling RIP, SAP and Type 20 packets.
- The IPX header is then examined to see if the Transport Control field is 16 or greater. If so the packet is discarded.

Note: If either the Source or Destination Network field of the IPX header is set to zero, the IPX process should fill the field with the appropriate network number before passing the packet to an internal socket process and/or routing the packet onward.

If the packet does not need to be handed internally or if the packet is a broadcast to one the router's directly connected segments other than the segment from which the packet was received it requires further routing. The router can now take one of two possible actions. **If the packet is destined for a network number that the router is directly connected to**, the router performs the following steps:

- It places the destination node address from the IPX header in the destination address field of the MAC header.
- It places its own node address in the source address field of the MAC header.
- It increments the Transport Control field of the IPX header, and transmits the packet on the destination node's segment (See Figure 2.4.)

If, however, the router is not directly connected to the segment that the final destination node resides on, it will send the packet to the next router in the path to the destination node. To forward the packet to another router:

- 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 a Router Information Table, see ch. 3).
- It places its own node address in the source address field of the MAC header.

- It increments the Transport Control field in the IPX header and sends the packet to the next router.

Note that in the two cases just described, **under normal conditions the only modification that the router makes to the IPX header when routing a packet is incrementing the Transport Control field.** All other fields are left as initially set by the sending node. The only exception to this is if the Source Address field in the IPX header has been set to zero by the sending node. In this case the router must fill in the appropriate source address before propagating the packet. (If the Destination Address field is set to zero the packet will not be routed.) Of course if the router is generating one of its own packets it will need to fill in the entire IPX header and perform the functions of a sending node.

Mixed Topology Routing

As mentioned earlier in the chapter, IPX routers must be able to handle packet sizes up to 576 bytes. Novell also highly recommends that routers take advantage of IPX enhancements which allow packets sizes larger than 576 by attempting to route packets of any size. A problem arises when a large packet from one network topology requires routing to another topology which cannot handle the large packet (for example, a 2 kbyte packet on a Token Ring network is destined to an Ethernet network that can only handle 1500 byte packets). In this scenario the packet should simply be dropped.

Burst Mode Protocol

The Novell Burst Mode Protocol is a relatively new protocol that employs enhancements to the NetWare Core Protocol. Burst Mode, which is built on top of IPX, was developed explicitly to improve NetWare's performance when transferring large files over wide area networks. Although no special handling on the part of a router is required for Burst Mode to function properly, the performance of Burst Mode – as well as any other packet sequence handling protocol for that matter – will be enhanced if routers always transfer IPX packets in the order they are received.

Routing Information Gathering

For a router to perform the routing tasks mentioned above it must maintain a database of information about the layout of the internetwork. All of this information is created and maintained using the RIP and SAP protocols which are described in detail in chapters 3 and 4.

Routing Information Protocol (RIP)

Introduction	3-2
Packet Format	3-3
RIP Operations	3-5
Router Implementation Guidelines	3-6
Selecting the Best Route	3-7
Routing Information Broadcasts	3-8
Split-Horizon Algorithm	3-8
Router Initialization	3-9
Router Shutdown	3-12
Router Information Maintenance	3-13
RIP Aging	3-14
Lost Route Algorithm	3-15
RIP Request Handling	3-15
RIP Interpacket Gap	3-16
RIP and SAP Interoperability	3-16

Introduction

The Routing Information Protocol (RIP) facilitates the exchange of routing information on a NetWare internetwork. NetWare routers use RIP to create and maintain a database of internetwork routing information (commonly called a Router Table). Like IPX, RIP was derived from XNS. However, an extra field was added to the packet structure to improve the decision criteria for selecting the fastest route to a destination. This change prohibits the straight integration of NetWare's RIP with other undeviating XNS implementations.

The single packet structure defined by RIP allows the following exchanges of information:

- Workstations locate the fastest route to a network number by broadcasting a route request.
- Routers request routing information from other routers to update their own internal tables by broadcasting a route request.
- Routers respond to route requests from workstations and other routers.
- Routers perform periodic broadcasts to make sure that all other routers are aware of the internetwork configuration.
- Routers perform broadcasts whenever they detect a change in the internetwork configuration.

The major components of a RIP implementation are graphically described in Figure 3.1.

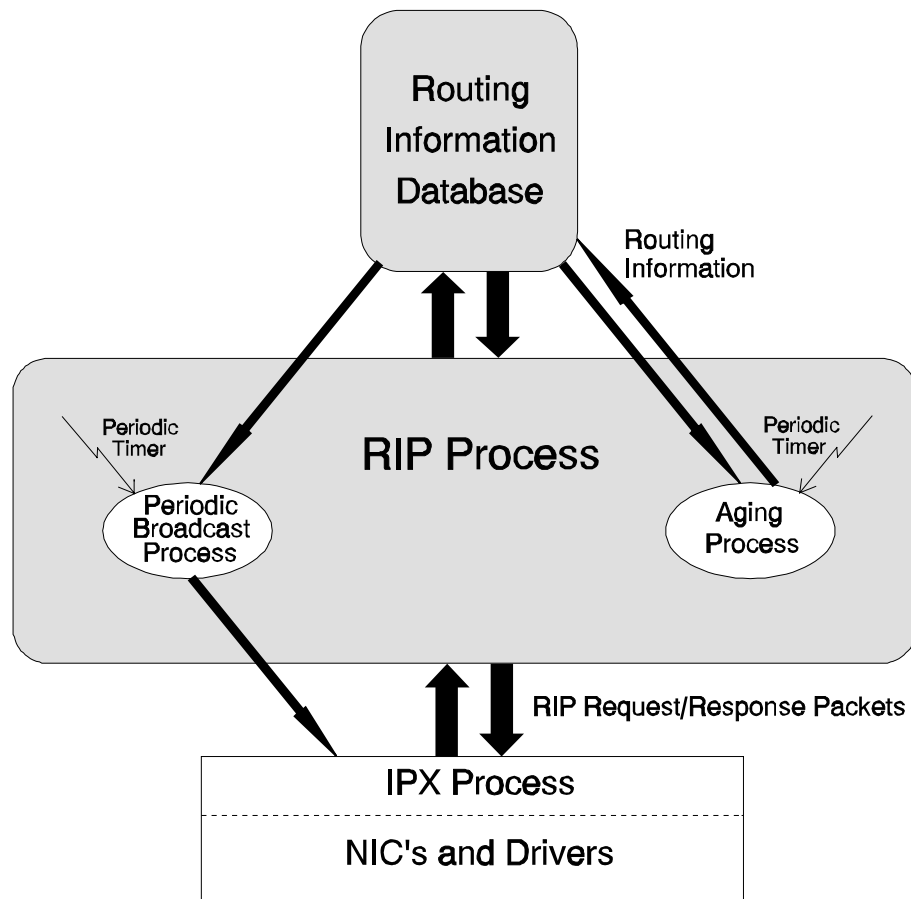


Figure 3.1 RIP Implementation Components

Packet Format

The RIP packet structure is shown in Figure 3.2. As with most of the higher-level protocols discussed in this document the packet structure is encapsulated within the data area of IPX. RIP packets are defined in the IPX header as Packet Type 1 and Socket Number 453h.

Note: Earlier version of NetWare did not always set the Packet Type to 1 in the IPX header for RIP packets. Also RIP requests may only have the Destination Socket set to 453h and not the Source Socket. Routers should check the Destination Socket field to determine if a packet is a RIP packet rather than depend on Packet Type. Similarly, in responding to a RIP request whose Source Socket is not 453h, the router should set the Destination Socket to whatever the sending node's Source Socket is and set the Source Socket to 453h.

RIP packet fields are described below.

Operation

This field indicates whether the packet is a request or a response: 1 = Request, 2 = Response. RIP requests and responses are discussed later in this chapter.

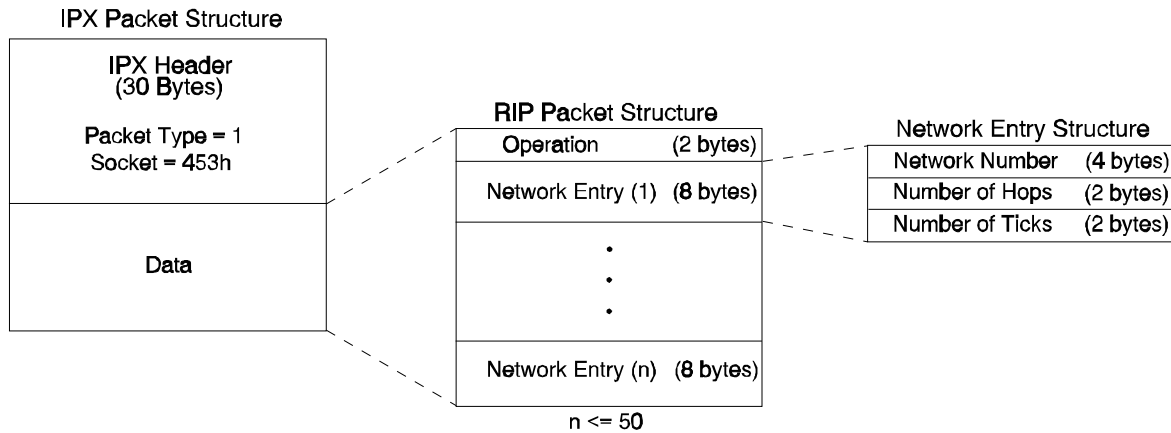


Figure 3.2 RIP Packet Structure

The Operation field can be followed by one or more sets of information. Each set of information is 8 bytes in length. A RIP packet can contain a maximum of 50 sets of network number information. Thus depending on the amount of network information contained in the packet the size of a RIP packet will vary anywhere from a minimum size of 40 bytes (IPX header plus one RIP network entry) to a maximum size of 432 bytes (IPX header plus 50 entries). Each RIP network entry consists of the following fields:

Network Number

This field indicates the Network Number that is the subject of the RIP request or response. On requests this field can be set to FFFFFFFFh to indicate a general request.

Number of Hops

This field indicates the number of routers that must be passed through to reach the network number associated with this field entry. For response packets this value must be at least 1.

Number of Ticks

This field indicates how much time, in ticks, that the packet takes to reach the network number associated with this field entry. A Tick is roughly 1/18 of a second (there are 18.21 ticks in a second, to be precise). The number in this field is always at least one for response packets.

The original XNS definition of the RIP did not include the Number of Ticks field. This field was added by the developers of NetWare so that the NetWare shell could estimate how long it should wait for a response from a file server. Also, if multiple routes exist to a network number, a router uses the route with the shortest number of Ticks when forwarding packets to that network number. If two routes exist with equal Tick values, the route with the least hops should be used.

For Novell routers, the initial time estimate is the responsibility of the router's LAN drivers. Each driver reports this estimate to its router. This time estimate is used by the router in its periodic broadcasts to indicate the time necessary to deliver a 576 byte packet to a node on that segment.

The method that Novell LAN drivers use for estimating the time delay on a segment depends on segment type. For local segments with more than 1 Mb/sec of bandwidth (Token-Ring, Ethernet, ARCNET, and so on), this delivery time is usually one tick. For remote segments (T1, 64 kbps, X.25, and asynchronous), the driver will periodically poll to determine the current time delay. For instance, the delay for a T1 link normally ranges from six to seven Ticks. If this delay changes, the driver will inform its router. As information about the segment is passed along throughout the network (by way of periodic broadcasts), routers will add any additional delay that they impose to the initial time estimate for the segment. (Note that this time estimate does not include any dynamic queuing delay.)

RIP Operations

RIP packets can be either requests or responses as indicated by the Operation field. If a RIP packet is a request for information (Operation = 1), only the Network Number field has meaning. However the Number of Hops and Number of Ticks fields must still be included for each Network Number entry in the packet. The values assigned to the Hops and Ticks fields for requests are irrelevant.

RIP requests can be categorized as either *general* or *specific*. A request is a general request if one the Network Number entries has a value of FFFFFFFFh (normally there is only 1 entry for a general request). General requests are broadcast by routers to obtain information about all networks that exist on an internetwork. Specific requests are sent

out to obtain information about specific networks. In the case of a specific request, one or more (to a maximum of 50) entries are filled in with the unique network numbers of those networks being requested.

RIP response packets (Operation = 2) come in one of two forms: 1) A reply to a general or specific request from a router or workstation, or 2) An informational broadcast by a router. These information broadcasts occur upon router start-up, shutdown, and when a router becomes aware of a routing change in the internetwork.

Routers also broadcast informational RIP response packets periodically which contain all routing information known to the router. This allows all routers on the internetwork to remain synchronized and also provides routers with a means of aging networks which might become inaccessible suddenly due to a router going down abnormally. Routers will become aware very quickly of a change of this nature since each periodic broadcast received will no longer contain information for lost routes. These routes are then aged and after a certain period of time deleted from the router's information database.

Note: If a RIP response entails providing information for more than 50 networks, multiple RIP response packets will be required.

Router Implementation Guidelines

To forward packets by the best possible route, NetWare routers maintain a Routing Information Table that holds information about all the network segments on the internetwork. Figure 3.3 shows a sample Routing Information Table (only the fields pertinent to this discussion have been included). Each entry in the Routing Information Table gives the router forwarding information for a network segment. This example table is simply intended to show the general form such a table might take and does not serve as a strict specification.

The first field contains the network numbers for segments that the router is currently aware exist. The router simply matches the destination network number in the packet's IPX header with an entry in this field to get its forwarding instructions for the packet. The second field indicates the number of routers that must be traversed to reach the network segment.

An estimate of the time necessary to reach the destination segment is recorded in the third field. The NIC field of the Routing Information Table records which NIC in the router the network segment can be reached through. The Immediate Address field contains the node address of the router that can forward packets to each segment. If the segment is directly connected to the router, this field will remain empty. The final field is used to make sure that information about the segment is current.

Network Number	Hops to Network	Ticks to Network	NIC	Immediate Address of Forwarding Router	Aging Timer
00000001	1	2	A		0
00000002	1	2	B		0
FEED0038	1	20	C		0
FEED0035	2	3	B	00001B029927	1
000000FF	2	3	A	00001B0349B2	2
FEED0036	3	4	A	00001B0349B2	2

Figure 3.3 Sample Portion of Routing Information Table

Selecting the Best Route

On large internetworks, there may be multiple routes to a single network. The criteria used by routers in selecting the "best" route to a network when choosing between alternate routes is shown in Figure 3.4.

The Routing Information Table might contain a list of all alternate routes for each network number in case the best route to a network number goes down. In other words, if the router can reach the segment through more than one of its NICs, it will make a record of both routes. Another option would be to maintain alternate routes only if these alternate routes require the same amount of Ticks to reach the segment as the primary route. This reduces the size of the Routing Information Table. Currently, developers of IPX routers are free to choose if the router will store only the best route, all known routes, or a subset of the latter.

Criteria routers use for selecting the Best Route:

- 1) Select the route that requires the lowest number of Ticks.
- 2) If multiple routes exist with the number of Ticks equal, select the route with the lowest number of Hops.
- 3) If routes exist with both Ticks and Hops equal the router is free to choose any of the routes as the Best route.

Figure 3.4 Best Route Criteria

Routing Information Broadcasts

On an internetwork, routers are constantly exchanging information with each other to make sure that their Routing Information Tables reflect up-to-the-minute changes in the layout of the internetwork. To accomplish this, routers transmit a series of broadcasts from the time they come up until they are brought down. These broadcasts can be separated by the time at which they occur:

- Initial broadcast of directly connected network segments
- Initial request to receive routing information from other routers
- Periodic broadcasts (every 60 seconds) of current list of active network numbers
- Broadcast of change in internetwork configuration
- Final broadcast when brought down

Although the broadcasts occur at different times and, for the most part, contain different information, they must follow two important rules. First, each broadcast must be a local broadcast, addressed so that it will not be immediately passed on, intact, by the routers that receive it. This reduces the network traffic created by these information exchanges. Second, routers must follow the "split-horizon" algorithm when providing information to other routers through a broadcast (since the second broadcast listed above is a request for information, this rule does not apply to it).

Split-Horizon Algorithm

The purpose of routing information broadcasts is twofold: to allow a router to share its current impression of the layout of the internetwork with other routers, and to inform the routers of an internetwork change so the routers can update their tables. A router sends routing information broadcasts to every network segment that it is directly connected to. The first rule of the split-horizon algorithm dictates that

a router about to broadcast to a particular network segment should not include any information about other segments that it has received from the segment to which the information is being sent.

For example, if Router 1 in Figure 3.5 is going to broadcast a routing information broadcast to network segment BB, it will not include information that it received from FS1 about network segment AA. If it did, someone on segment BB might erroneously conclude that there are two paths to segment AA – one through FS1 and another through Router 1.

The split-horizon algorithm also states that routers should not include information about the network segment that they are sending routing information broadcasts to. For example, Router 1 would not include information about BB in its broadcast onto BB.

Taking these rules into account, the information that Router 1 would broadcast onto segment BB would be information about segments CC and DD.

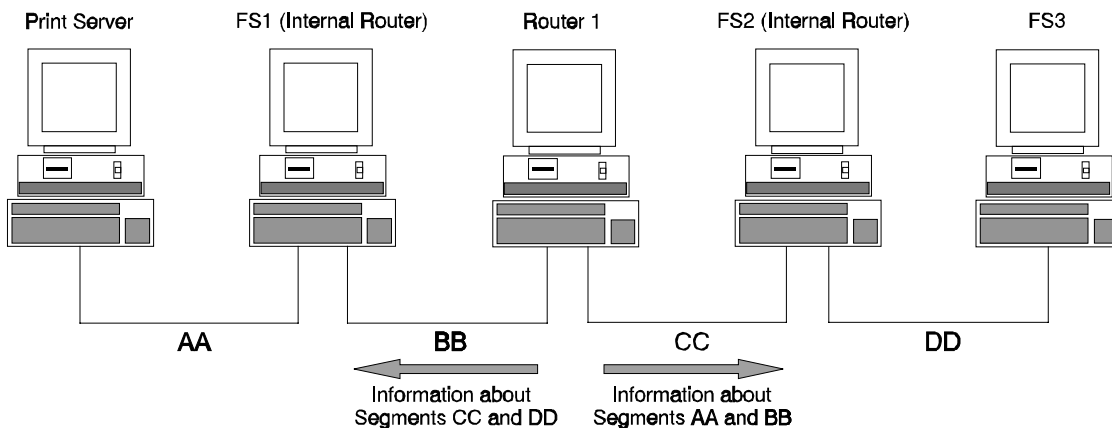


Figure 3.5 The Split-Horizon Algorithm

Note: The Split-Horizon Algorithm is also known as the Best Information Algorithm.

Router Initialization

When a router is first brought up, the following sequence of events takes place:

- The router places the network numbers of its directly connected segments into its Routing Information Table.
- Following the split-horizon algorithm, the router sends a RIP

broadcast response packet to inform the routers on its directly connected segments of the network segments that the router will be making available. (See Figure 3.6.)

- The router next broadcasts a general request to each of its directly connected segments for information about all other network segments that exist on the internetwork. (See Figure 3.7.)
- The general request is responded to by all the routers (each using split-horizon) on these directly connected segments. (See Figure 3.8.)
- The router places the information gleaned from these responses in its Routing Information Table.
- The router now begins to send out RIP broadcast response packets containing all the information in its Routing Information Table (except that excluded by the split-horizon algorithm) to each of its connected network segments every 60 seconds. (See Figure 3.9.)




Figure 3.6 Initial RIP Broadcast




Figure 3.7 Initial RIP Request




Figure 3.8 RIP Response from Local Routers




Figure 3.9 RIP Sixty Second Broadcasts

Once the router has performed these tasks it is ready to accept routing requests and route packets. The periodic broadcast packets are sent to ensure that all routers on the internetwork remain synchronized.

Note: Because of the lower bandwidth of X.25 and asynchronous links, Novell routers historically have not performed 60 second broadcasts on these links — only initial broadcasts, changed information broadcasts and final broadcasts are sent over these links.

Router Shutdown

When a router is brought down "cleanly" (for a NetWare router, "cleanly" would mean the DOWN command was issued at the console as opposed to someone just shutting the power off or some other power failure), the router will inform its directly connected segments of the fact before discontinuing service. To do this the router issues broadcasts (as always, using the split-horizon algorithm) that indicate that the network segments which the router had made available will no longer be accessible through this router. The format of these broadcast packets is identical to that of the sixty second broadcasts except that the Hops value for each network entry in the RIP packet is set to 16 (see Figure 3.10).

Of course a router will be unable to broadcast itself as DOWN should a hardware failure, power glitch, or power outage occur. In order for other routers to become aware of an internetwork change of this type an "aging" mechanism has been built into NetWare routers. Aging is described later in this chapter.




Figure 3.10 RIP DOWN Broadcasts

Router Information Maintenance

When a router receives a RIP broadcast response indicating a change in the internetwork configuration (in other words a network has gone down, been brought up, or is accessible through a different route), the router must update its Router Information Table and inform other routers of these changes.

To relay this information to the rest of the internetwork the router sends a broadcast (using split-horizon) immediately to all of its directly connected segments except the segment that the router received the information from. This broadcast packet will contain information regarding the change that occurred in the network; this information will also be reflected in all future periodic broadcasts. (A "downed" network requires special consideration, see section on Lost Route Algorithm.)

Note: If a router becomes aware of a network route with Number of Hops = 15, the router should store this route in its Router Information Table (unless a better route already exists to this network, in which case the router may not be interested in storing the route). However, the router should not propagate further any information about the route since in doing so the Hops would now be reported as 16; all routers receiving this information would interpret this to mean the route had gone down.

A similar procedure should be followed if the router becomes aware of a network change by some means other than by receiving a RIP broadcast. An example would be any time a directly connected network segment becomes operational or goes down. In the case when a directly connected segment goes down a broadcast packet is immediately sent to all other directly connected segments informing them that all

networks accessible through the segment that went down are now DOWN (Hops = 16).

In the case when a directly connected segment has been added, in addition to sending a RIP broadcast to all other directly connected segments informing them of the newly added network segment the router must do the following:

- A RIP broadcast is sent to the newly added segment (using split-horizon) containing all network information known to the router.
- A RIP general request is broadcast to the newly added segment to find out about all distant networks that can be reached through this new route.
- If any response is received in reply to the general request this information is immediately added to the Router Information Table and then broadcast to other directly connected segments using the split-horizon algorithm.

These procedures will insure that information about any change in the internetwork will propagate throughout the entire internetwork in a relatively short amount of time.

RIP Aging

Routers must implement an aging mechanism to handle those conditions (hardware failure, power glitch, power outage, etc) which would cause a router to go down suddenly without sending a DOWN broadcast. Routers might maintain a timer for each entry in their Routing Information Table which keeps track of how much time has elapsed since information was received concerning a particular table entry. Every time information is received for a table entry the timer is reset to zero. If the timer reaches 3 minutes, the router assumes the route for that network entry is no longer valid and deletes it from the table. The router will then send an immediate broadcast (split-horizon style) to all directly connected segments informing them that the network is DOWN (Hops = 16). Since this information is new or changed, the routers that receive this information will pass it on immediately and the change will quickly permeate the internetwork. The router should also follow the Lost Route Algorithm, described next.

Note: In the future, NetWare routers will allow the RIP aging and RIP periodic broadcast intervals to be configurable. IPX Router developers may want to include similar functionality in their RIP implementations.

Lost Route Algorithm

When a router becomes aware that a network route has been eliminated (from a DOWN broadcast or through the aging process), besides sending a broadcast packet informing other routers of this change — as described earlier — the router should follow a procedure known as the Lost Route algorithm. This algorithm states that after becoming aware that a route to a network has been eliminated, a router should wait 10 ticks (or close to it; this is just over a half a second) and check its Routing Information Table for an alternate route. Should an alternate route be found the router then sends a RIP broadcast (using split-horizon) informing other routers of the alternate route.

RIP Request Handling

When a router receives a general request (Network Number entry = FFFFFFFFh), a RIP response packet containing information about all networks known to the router is sent to the sending source using the split-horizon algorithm. (Basically this would include the same information as is sent out in each periodic broadcast.) For large internetworks, multiple RIP response packets may be required.

If the request however is a specific request for 1 or more networks, the router will send a RIP response containing the specific network information asked for by the requesting source (as far as the router knows this information) using split-horizon. (In other words, the router would respond with any information normally sent in a periodic broadcast packet that would be applicable to the request.)

One exception to using the split-horizon algorithm occurs if a specific RIP request is received from a node on a directly connected network segment. Normally this type request can be handled like any other specific request, however if the request is for information regarding the same directly connected network that the request was received from (for example, a request is received from a workstation on directly connected network segment AAAA with Network Number entry = AAAA in the RIP portion of the packet), the router should respond with the requested information. Although the RIP response information is useless as far as routing packets is concerned, the address information from the IPX header of each response allows a node to learn the node addresses of all routers on its local network segment.

For any of the responses just mentioned, routers should not include any information for networks for which the Number of Hops is 16.

RIP Interpacket Gap

NetWare servers and routers may have a problem properly processing a sequence of RIP packets if the interpacket gap time between those packets is not sufficient. Currently, the minimum interpacket gap time allowed for RIP packets is 55 milliseconds.

RIP and SAP Interoperability

Routers can improve their decision making and information broadcasting ability to some degree by allowing RIP and SAP processes to be aware of each other. SAP, as well as RIP/SAP interoperability issues are discussed in the next chapter.

Service Advertising Protocol (SAP)

Introduction	4-2
Packet Format	4-3
SAP Operations	4-6
Router Implementation Guidelines	4-7
Server Information Table	4-7
Router Initialization	4-8
Router Shutdown	4-11
Server Information Maintenance	4-11
SAP Aging	4-12
SAP Request Handling	4-13
SAP Interpacket Gap	4-14
RIP and SAP Interoperability	4-14

Introduction

The Service Advertising Protocol allows service-providing nodes – such as file servers, print servers, and gateway servers – to advertise their services and addresses. SAP makes the process of adding and removing services on an internetwork dynamic. As servers are booted up, they advertise their services using SAP; when they are brought down, they use SAP to indicate that their services will no longer be available.

Through SAP, routers create and maintain a database of internetwork service information. This allows clients on the network to determine what services are available on the network and obtain the internetwork address of the nodes (servers) where they can access those services. This is an important function, since a workstation cannot initiate a session with a file server without first having that server's address.

A gateway server, for instance, will periodically broadcast a SAP packet onto the network segment it is connected to. The SAP agent in each router on that segment copies the information contained in the SAP packet into an internal table (commonly called a Server Information Table). Because the SAP agent in each router keeps up-to-date information on available servers, a client wanting to locate the gateway server can access a nearby router for the correct internetwork address.

Like the RIP, SAP uses IPX and the medium-access protocols for its transport. The packet structure provided by SAP allows for five different functions:

- A workstation request for the name and address of the nearest server of a certain type.
- A general request, by a router, for the names and addresses of either all the servers or all the servers of a certain type on the internetwork .
- A response to either a nearest server request or a general request.
- Periodic broadcasts by servers and routers.
- Changed server information broadcasts.

Note that SAP agents in non-routing servers are only required to perform periodic broadcasts and respond to SAP requests. All other SAP functionality described hereafter applies to router SAP agents only. Figure 4.1 graphically describes the major components of a SAP implementation.

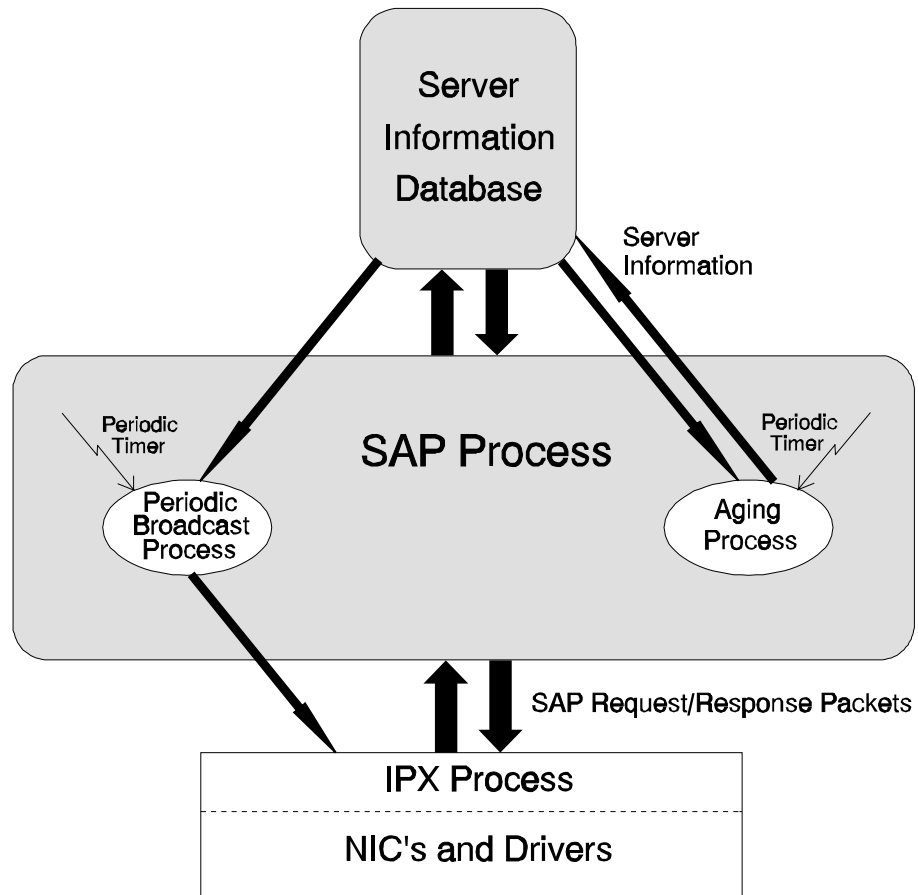


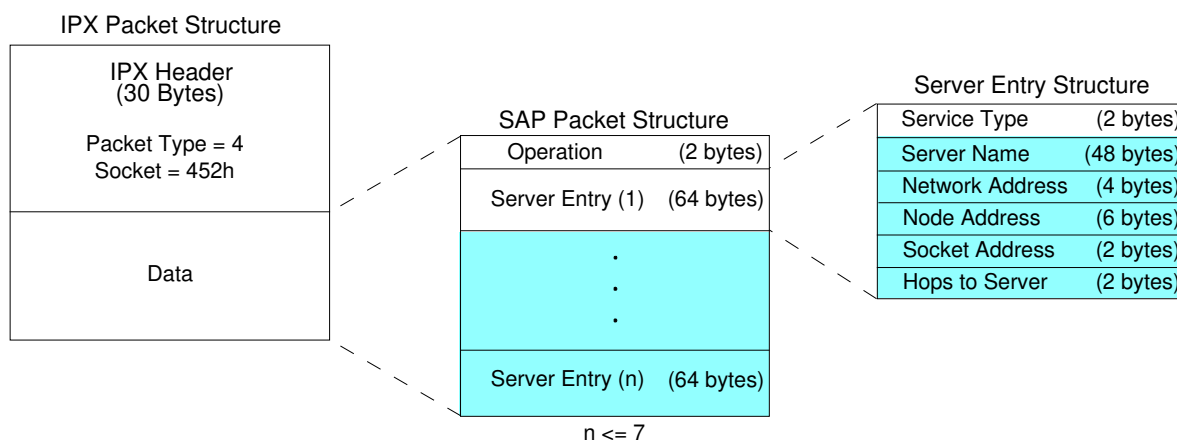
Figure 4.1 SAP Implementation Components

Packet Format

The SAP packet structure is shown in Figure 4.2. As with most of the higher-level protocols discussed in this document the packet structure is encapsulated within the data area of IPX. SAP packets are defined in the IPX header as Packet Type 4 and Socket Number 452h.

Note: Earlier version of NetWare did not always set the Packet Type to 4 in the IPX header for SAP packets. Also SAP requests may only have the Destination Socket set to 452h and not the Source Socket. Routers should check the Destination Socket field to determine if a packet is a SAP packet rather than depend on Packet Type. Similarly, in responding to SAP request whose Source Socket is not 452h, the router should set the Destination Socket to whatever the sending node's Source Socket is and set the Source Socket to 452h.

SAP packet fields are described below.



Note: Shaded fields apply only to SAP response packets.

Figure 4.2 SAP Packet Structure

Operation

This field indicates the type of operation the SAP packet performs and can be set to one of the following values:

- 1 = Request
- 2 = Response
- 3 = Get Nearest Server Request
- 4 = Get Nearest Server Response

SAP requests and responses are discussed later in this chapter.

Depending on the type of operation, the Operation field will be followed by either a single field or 1 or more sets of fields. For all SAP requests (operations 1 and 3), only the first Service Type field is included in the packet. This means that ALL SAP requests will be of length 34, not including media header (IPX header plus SAP Operation and Service Type fields = 30 + 2 + 2 = 34 bytes). All other fields apply ONLY to SAP responses (operations 2 and 4). A SAP response may include from one to seven sets of fields. This means that SAP response packets may vary in size from 96 bytes (IPX header plus one server entry) to 480 bytes (IPX header plus 7 entries). Each server entry includes information regarding a particular server and consists of the following fields:

Service Type

This field identifies the type of service the server provides. Novell assigns each type of server a unique Server Type. Software developers

must contact Novell to obtain Server Type assignments for any value-added server product they create. Server Types can be obtained from Novell by calling 800-729-4357 or (512) 794-1795. For example, a Novell file server advertises itself as type 4. This value becomes the object type for this server as it is found in the NetWare bindery.

Note: Although IPX routers use SAP, routers typically do not act as servers and require no Server Type assignment.

Following is a list of some well-known server (object) types:

Wild	FFFFh (-1)
Unknown	0000h
Print Queue	0003h
File Server	0004h
Job Server	0005h
Print Server	0007h
Archive Server	0009h
Remote Bridge Server	0024h
Advertising Print Server	0047h
Reserved Up To	8000h

Server Name

This field contains the 48 byte object name that is assigned to a server. The Server Name, in combination with the Service Type, uniquely identifies a server on an internetwork. Routers need not be too concerned about the naming conventions for servers, since all server name information of interest to the router will be contained in SAP response packets received from other router and server SAP agents. Routers must make sure this information is stored in the exact format (all 48 bytes, including any NULL padding) as it is received so that any subsequent SAP response packets sent by the router will maintain the correct format for any Server Name entries.

Network Address

This field contains the address of the network on which the server resides.

Node Address

This field contains the address of the node on which the server resides.

Socket Address

This field contains the socket number on which the server will receive service requests.

Hops to Server

This field indicates the number of routers that must be passed through to reach the server associated with this field entry. Each time the packet passes through an intermediate network, the field is incremented by one.

SAP Operations

If the Operation field is 1, the SAP packet is either a *specific* or *general* request. For either of these requests only the Operation and Service Type fields are used and packets are limited in size to 34 bytes. Specific requests are used to find out about all servers of a specific type. In this case the Service Type field is set to a specific value (for example Service Type = 4 for file servers). General requests are used to find out about all servers of any type. This type request will occur if the Service Type field is set to FFFFh.

If the Operation is 2, the packet is a SAP response which might come in one of two forms: 1) A reply to a specific or general request from a server or router or 2) An informational broadcast. These informational broadcasts occur upon server start-up and shutdown, and when a change occurs in server information on the internetwork. Router SAP agents also broadcast informational SAP response packets periodically which contain all server information known to the SAP agent. This allows all SAP agents on the internetwork to remain synchronized and also provides router SAP agents with a means of aging servers which might become inaccessible suddenly due to a router or server going down abnormally. SAP agents will become aware very quickly of a change of this nature since each periodic broadcast received will no longer contain information for lost servers. These servers are then aged and after a certain period of time deleted from the SAP agents server information table.

Note: Should a SAP response require information for more than 7 servers, multiple SAP response packets must be sent.

If the Operation is 3, the packet is a SAP Get Nearest Server request. This request is used by a workstation to find the nearest server of a particular type. In this context, nearest refers to the server or router that is able to respond the quickest. This type of request packet, like all other SAP request packets is limited in size to 34 bytes. Also the Service Type field must be set to a specific Type, FFFFh is not allowed.

Operation 4 signifies a SAP Get Nearest Server response (sometimes called a Give Nearest Server packet). When a server or router receives a Get Nearest Server request it responds with this type of packet. The packet will be 96 bytes in length (IPX header plus one server entry).

More information on the usage and handling of each of these types of SAP packets is given later in this chapter.

Router Implementation Guidelines

Using the SAP, servers on a NetWare network can advertise their services and addresses. The information that these servers broadcast is not directly used by clients but instead collected by a SAP agent within each NetWare router on the server's segment. The SAP agents store this information in a Server Information Table and, if they reside within a server, in their server's bindery. The clients can then contact the nearest router or file server SAP agent for server information.

The SAP broadcasts that servers and routers perform are local broadcasts and, therefore, only received by SAP agents on their connected segments. Consequently, SAP agents periodically broadcast their server information so that all SAP agents on the internetwork have information about all servers that are active on the internetwork. Note that this is the same broadcast method used by routers to distribute and exchange network number (RIP) information, and, as in the case of RIP broadcasts the split-horizon algorithm applies to all SAP broadcasts. The split-horizon algorithm is described in chapter 3.

Server Information Table

The table that SAP agents use to store information received in SAP broadcasts is called the Server Information Table. If all SAP agents on the internetwork are exchanging SAP information properly, each agent's Server Information Table should have information about all the servers on the internetwork, thus providing clients with nearby access to the addresses of all the servers on the internetwork. Figure 4.3 shows a sample portion of a Server Information Table. This example table is simply intended to show the general form such a table might take and does not serve as a strict specification.

The first 4 fields reflect the information obtained from SAP response packets (with the appropriate adjustments made to the Hops to Server field). The network card that the information about the server was received on is specified in the NIC field. The Aging Timer field is used for aging servers that have unexpectedly gone down.

Server Name	Server Address (Net:Node:Socket)	Server Type	Hops to Server	NIC	Aging Timer
FSERVER1	00000123:00001B029888:0451	4	1	A	0
GSERVER3	00000002:00001B0349B2:0451	6	2	B	1
PSEVER2	FEED0038:00001B023456:0456	2	3	C	2

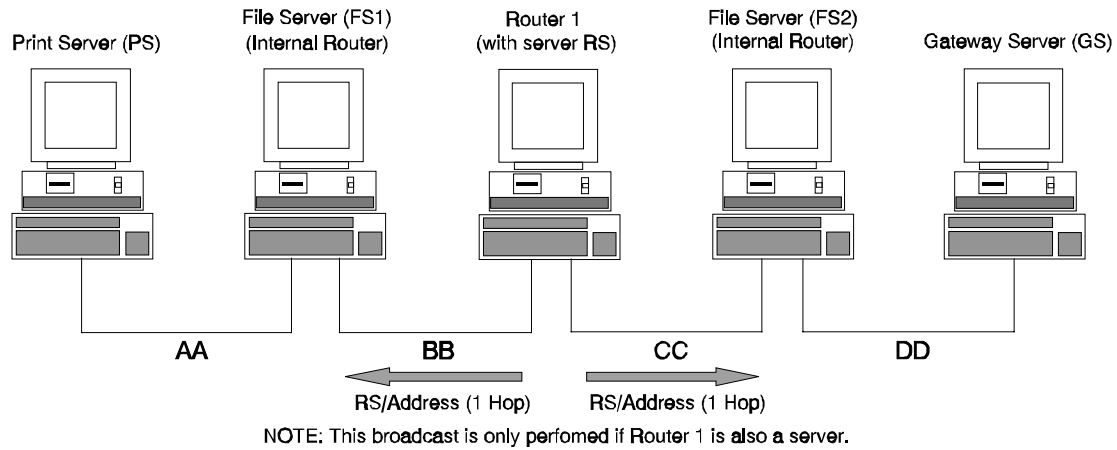

Figure 4.3 Sample Portion of Server Information Table

As in the case of the Router Information Table, IPX router developers are free to choose whether to keep track of only the best route to a server, all known routes, or a subset of the latter.

Router Initialization

When a router is first brought up, the following sequence of events takes place:

- **If the router is also acting as a server** (this is not required by a router and will probably not be the case and routers may skip this step), its internal SAP agent places this server information in the agent's Server Information Table and then sends a SAP broadcast, following the split-horizon algorithm, onto each of its directly connected segments to inform the SAP agents on those segments that a new server has become available. (See Figure 4.4.)
- The SAP agent then broadcasts a general request onto each of its directly connected segments for information about all other servers that exist on the internetwork. (See Figure 4.5.)
- The general request is responded to by all the SAP agents on these directly connected segments. (See Figure 4.6.)
- The SAP agent places the information received in these responses in its Server Information Table.
- The SAP agent now begins to send out SAP broadcast response packets containing all the information in its Service Information Table (except that excluded by the split-horizon algorithm) to each of its connected network segments every 60 seconds. (See Figure 4.7.)

*Figure 4.4 Initial SAP Broadcast**Figure 4.5 Initial SAP Request*

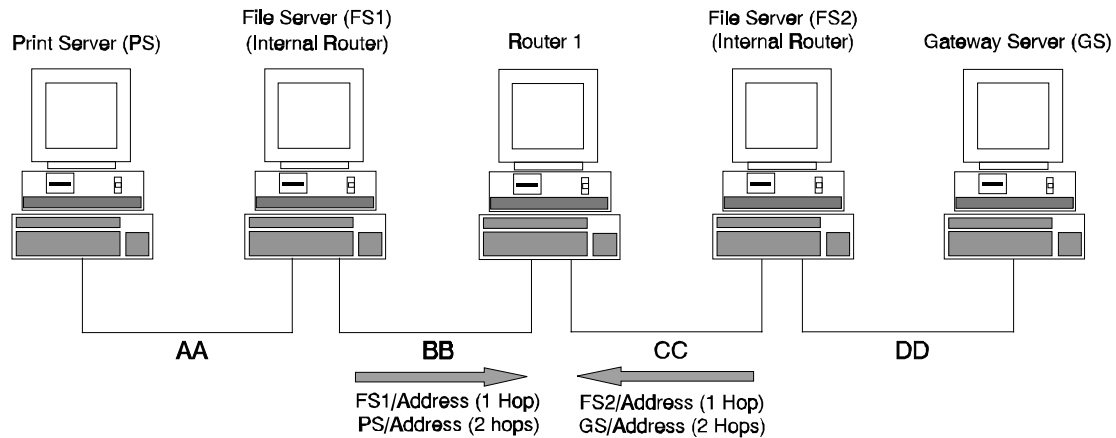
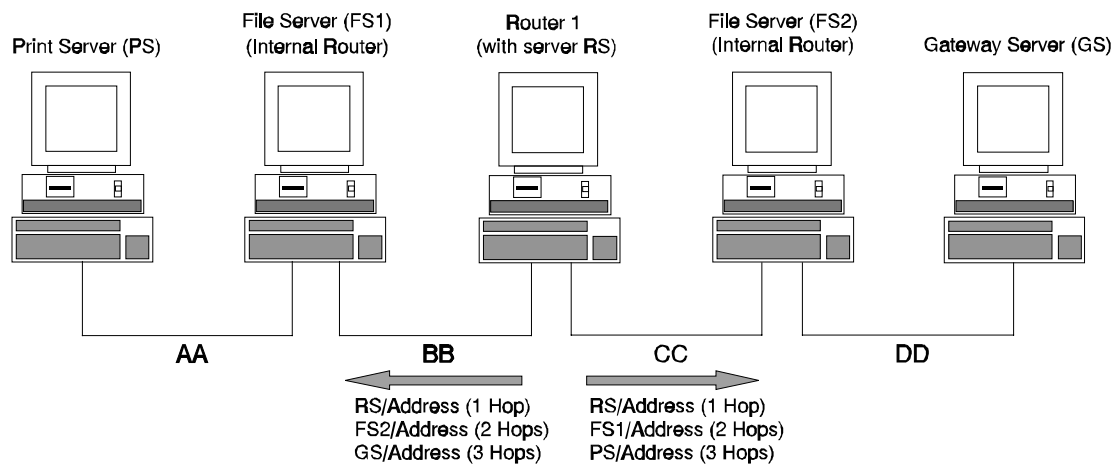


Figure 4.6 SAP Response from SAP Agents



NOTE: The broadcast of RS is only performed if Router 1 is also a server.

Figure 4.7 SAP Sixty Second Broadcasts

Once the router's SAP agent has performed these tasks it is ready to accept SAP requests and broadcasts. The periodic broadcasts ensure that all SAP agents on the internetwork remain synchronized.

Note: Because of the lower bandwidth of X.25 and asynchronous links, Novell routers historically have not performed 60 second broadcasts on these links – only initial broadcasts, changed information broadcasts and final broadcasts are sent over these links.

Note: Routers may want to stagger RIP and SAP 60 second broadcasts every 30 seconds so that the RIP and SAP traffic load is more evenly distributed.

Router Shutdown

When a router is brought down "cleanly" (for a NetWare router, "cleanly" would mean the DOWN command was issued at the console as opposed to someone just shutting the power off or some other power failure), the router will inform its directly connected segments of the fact before discontinuing operation. To do this the router issues broadcasts (as always, using the split-horizon algorithm) that indicate that the servers which the router had made available will no longer be accessible through this router. The format of these broadcast packets is identical to that of the sixty second broadcasts except that the Hops value for each server entry in the SAP packet is set to 16 (see Figure 4.8).

Of course a router will be unable to broadcast a DOWN packet should a hardware failure, power glitch, or power outage occur. In order for other routers to become aware of an internetwork change of this type an "aging" mechanism has been built into NetWare SAP agents. Aging is described later in this chapter.

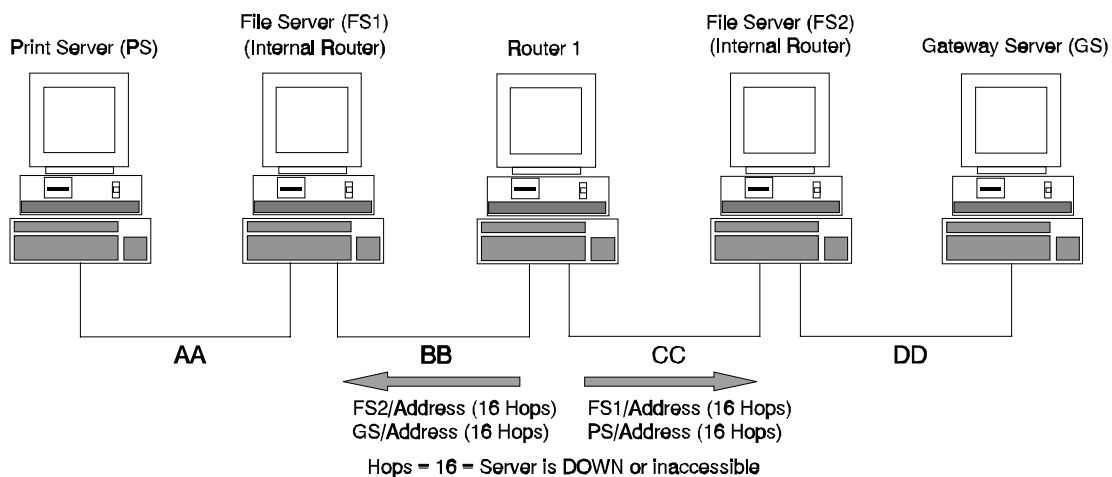


Figure 4.8 SAP DOWN Broadcasts

Server Information Maintenance

When a router's SAP agent receives a SAP broadcast response indicating a change in the internetwork server configuration (in other words a server has gone down, been brought up, or is accessible through a different route), the agent must update its Server Information Table and inform other SAP agents of these changes.

To relay this information to the rest of the internetwork the SAP agent sends a broadcast (using split-horizon) immediately to all of its directly connected segments except the segment that the router received the

information from. This broadcast packet will contain information regarding the server change that occurred; this information will also be reflected in all future periodic broadcasts.

A similar procedure should be followed if the router SAP agent becomes aware of a server change by some means other than by receiving a SAP response. An example would be any time a directly connected network segment becomes operational or goes down. In the case when a directly connected segment goes down a broadcast packet is immediately sent to all other directly connected segments informing them that all services accessible through the segment that went down are now DOWN (Hops = 16).

In the case when a directly connected segment has been added, the router's SAP agent must do the following:

- A SAP broadcast is sent to the newly added segment (using split-horizon) containing all server information known to the router.
- A SAP general request is broadcast to the newly added segment to find out about all servers that can be reached through this new route.
- If any response is received in reply to the general request this information is immediately added to the Server Information Table and then broadcast to other directly connected segments using the split-horizon algorithm.

These procedures will insure that information about any change in internetwork servers will propagate throughout the entire internetwork in a relatively short amount of time.

SAP Aging

Router SAP agents must implement an aging mechanism to handle those conditions (hardware failure, power glitch, power outage, etc) which would cause a SAP agent to go down suddenly without sending a DOWN broadcast. SAP agents maintain a timer for each entry in their Server Information Table which keeps track of how much time has elapsed since information was received concerning a particular table entry. Every time information is received concerning a table entry the timer is reset to zero. If the timer reaches 3 minutes, the SAP agent assumes that server entry is no longer valid and deletes it from the table. The SAP agent will then send an immediate broadcast (split-horizon style) to all directly connected segments informing them that the server is DOWN (Hops = 16). Since this information is new or changed, the SAP agents that receive this information will pass it on immediately and the change will quickly permeate the internetwork.

Note: In the future, NetWare routers will allow the SAP aging and SAP periodic broadcast intervals to be configurable. IPX Router developers may want to include similar functionality in their SAP implementations.

SAP Request Handling

When a SAP agent receives a general request a SAP response packet containing information about all servers of any type known to the SAP agent is sent to the sending source. (Basically, this would include the same information sent out in a periodic broadcast.)

If the request is a specific request, the SAP agent will send a SAP response directly to the requesting node containing information regarding all servers of the type asked for by the requesting source (as far as the router knows this information).

Note: For large internetworks with many servers, multiple SAP response packets may need to be sent to satisfy the general/specific request.

If a SAP Get Nearest Server request is received, the SAP agent responds with the nearest server of the type requested. The criteria used to select the nearest server is shown in Figure 4.9.

Criteria SAP agents use for selecting the Nearest Server:

- 1) Server is owned by the SAP agent. In the case of a router this will only be true if the router is also a server.
- 2) Server with the best route. This is determined from the Router Information Table. (See RIP/SAP Interoperability section later in chapter; also see chapter 3 section on Selecting the Best Route.)
- 3) Server with the least number of Hops. This is determined from the Server Information Table.

SAP agents first try to find a server which meets criteria 1. If no server can be found to match this criteria, the SAP agent moves to the 2nd criteria, then the 3rd. If no server can be found which meets any of these criteria the SAP agent is free to choose any server as the Nearest Server.

Figure 4.9 Nearest Server Criteria

Note: In responding to all requests: if the requests are broadcast requests (in other words, the IPX header Destination Node address is FFFFFFFFh), responses follow the split-horizon algorithm. However, in responding to any direct requests (IPX header Destination Node is not broadcast or multicast), the split-horizon algorithm is not

followed.

SAP Interpacket Gap

NetWare servers and routers may have a problem properly processing a sequence of SAP packets if the interpacket gap time between those packets is not sufficient. Currently, the minimum interpacket gap time allowed for SAP packets is 55 milliseconds.

RIP and SAP Interoperability

Routers can improve their decision making and information broadcasting ability by allowing RIP and SAP processes to be aware of each other. For example, in the process of determining the nearest server when responding to a Get Nearest Server request, a SAP agent must be capable of consulting the Router Information Table to determine which route to a server is the fastest (or has the lowest Tick count). Also before accepting a server a SAP agent should make sure a route exists to the server. Similarly if a network goes down, all SAP services associated with that network should immediately be marked as DOWN by the SAP agent. Of course if this is not done, the SAP services will be set DOWN eventually anyway by the aging process. However there would be about to a 4 minute delay before this would happen during which time the RIP and SAP information being propagated on the internetwork would be out-of-sync.

IPX Type 20 Propagation Packet

Introduction	5-2
Packet Format	5-2
Router Implementation Guidelines	5-3

Introduction

In order for certain protocol implementations, like NetBIOS, to function in the NetWare environment, routers must allow a broadcast packet to be propagated throughout an internetwork. The IPX Packet Type 20 is used specifically for this purpose. The special handling of this packet required by a router is described in this chapter.

Note: The handling of Type 20 packets described in this chapter applies only to broadcasts (Destination Node = FFFFFFFFh) and not to directed packets. Furthermore, Type 20 packets should not be propagated across slower links (like X.25 and asynchronous links) with bandwidths of less than 1 Mbps.

Packet Format

The IPX Type 20 packet structure as implemented by Novell is shown in Figure 5.1. Note that this packet is defined in the IPX header as Packet Type 20 (14h).

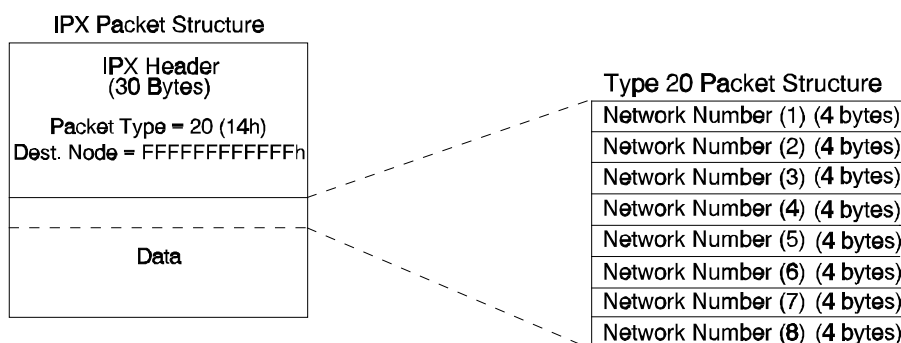


Figure 5.1 Type 20 Packet Structure

Network Number (1-8)

This 32-byte area is used to store up to 8 4-byte network addresses. As a Type 20 broadcast packet is propagated from the source network to another distant network across the internetwork the address of each network traversed is stored in this area. Note that the propagation of a Type 20 packet is limited to 8 networks.

Unused fields are filled with zeroes.

Router Implementation Guidelines

When a Type 20 broadcast packet is received by a router (this is indicated by a Packet Type of 14h, and Destination Node of FFFFFFFFFFh in the IPX header,) the following sequence of events should occur:

- The Transport Control field of the IPX header is examined; this value indicates the number of routers traversed by the packet so far. This field also indicates how many Network Number fields in the packet have been filled in. If this value is 8 or greater the packet is discarded, otherwise proceed to the next step.
- 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 a match is found the packet is discarded to prevent multiple traversals of the same network segment; if no match is found proceed to the next step.
- The router places the address of the network segment on which the packet arrived into the next available Network Number field. The offset of this field is easily calculated as $4 * n$ bytes past the end of the IPX header, where n is the value of the Transport Control field.
- The router increments the Transport Control field of the IPX header and broadcasts the packet **to all directly connected network segments that are NOT included in the Network Number fields.**

Note that for this type of packet the IPX header Destination Network field is meaningless. Notice also that if the last step mentioned above is obeyed by all routers, the second step becomes unnecessary. Routers should nonetheless perform this step to guard against improper handling of a Type 20 packet by a misbehaving router.

RIP and SAP Bandwidth Requirements

On large internetworks with several hundred servers, administrators become concerned with the load that RIP and SAP broadcasts will place on their network segments. These concerns can be appeased for asynchronous and X.25 links if only changed server and routing information is sent on these lines (no 60 second broadcasts). On other segment types the traffic caused by these broadcasts usually will not cause a significant load. The requirements and some examples are shown below. As can be seen, the SAP broadcasts for an internetwork containing 250 servers account for less than one percent of the total bandwidth (10Mb/s) of an Ethernet segment.

RIP 60 Second Broadcast Bandwidth Requirements

32-byte header
8 bytes per RIPv4 device
Maximum 50 entries/packet

3 Network Numbers	=	56 bytes/minute	
52 Network Numbers	=	480 bytes/minute	= 64 bits/second

SAP 60 Second Broadcast Bandwidth Requirements

32-byte header
64 bytes per SAPing device
Maximum 7 entries/packet

6 SAPing devices	=	416 bytes/minute	
250 SAPing devices	=	17,152 bytes/minute	= 2,286 bits/second

Total traffic load of routing and server information broadcasts on any given segment will be equal to broadcasting information about all the network segments and servers that exist on the internetwork. For example, on a T1 link between two NetWare routers, one router will broadcast information about all of the network segments and servers that it is making available to the other router (using the split-horizon algorithm). The other router will broadcast information about all the segments and servers that it is making available to the first router. The total of the two equals the total number of servers and segments that reside on the internetwork.

Index

A

Addressing 1-4, 1-6, 2-2, 2-7, 2-8
 Aging 1-5, 3-1, 3-6, 3-12, 3-14, 4-1, 4-6, 4-7,
 4-11, 4-12, 4-13,
 4-14
 Alternate routes 3-15

B

Best route 3-1, 3-7, 3-8, 4-8
 Burst Mode Protocol 2-1, 2-10

C

Changed information broadcasts 3-12, 4-10
 Connectionless 1-4, 2-2

D

Datagram 1-4, 2-2
 Destination Network 2-4, 2-5, 2-6, 2-9, 3-6,
 5-3
 Destination Node 1-4, 2-4, 2-5, 2-7, 2-8, 2-9,
 4-13, 5-2, 5-3
 Destination Socket 2-4, 2-5, 3-3, 4-3
 Directly connected network 2-9, 3-8, 3-13,
 3-15, 4-12, 5-3

F

Fastest route 1-5, 2-8, 3-2

G

General request 3-4, 3-5, 3-10, 3-14, 3-15,
 4-2, 4-6, 4-8, 4-12,
 4-13
 Get Nearest Server Request 4-4, 4-6, 4-7,
 4-13, 4-14
 Get Nearest Server Response 4-4, 4-7

H

Hardware failure 3-12, 3-14, 4-11, 4-12
 Hops 3-4, 3-5, 3-12, 3-13, 3-14, 3-15, 4-6,
 4-7, 4-11, 4-12

I

Internetwork Packet Exchange (IPX) 1-1,
 1-2, 1-4, 2-1
 IPX Checksum 2-3
 IPX Header 2-2, 2-3, 2-4, 2-7,
 2-8, 2-9, 2-10, 3-3,
 3-4, 3-6, 3-15, 4-3,
 4-4, 4-7, 4-13, 5-2,
 5-3
 IPX Router 1-6, 1-7, 2-6, 2-7,
 3-14, 4-8, 4-13
 Mixed Topology Routing 2-1,
 2-10
 Packet Delivery 2-1, 2-7
 Packet Format 2-1, 2-2
 Packet Length 2-2, 2-3
 Router's Responsibility 2-1, 2-9
 Routing Information Gathering
 2-1, 2-10
 Sending Node's Responsibility
 2-1, 2-7

L

Lost Route Algorithm 3-1, 3-13, 3-14, 3-15

M

MAC Header 1-4, 2-2, 2-3, 2-7, 2-8, 2-9
 Medium Access Protocols 1-1, 1-2, 1-3, 1-4
 Mixed Topology Routing 2-1, 2-10
 Multiple routes 3-5, 3-7

N

Nearest server 2-4, 4-2, 4-4, 4-6, 4-7, 4-13,
 4-14

NetBIOS 1-1, 1-2, 1-6, 2-5, 5-2
 NetWare Communication Protocols 1-1, 1-2,
 1-4, 1-6, 2-2, 2-7,
 2-8
 NetWare Core Protocol 1-2, 2-4, 2-5, 2-10
 NetWare protocol stack 1-2
 Network address 4-5
 Network number 2-2, 2-4, 2-6, 2-7, 2-8, 2-9,
 3-2, 3-4, 3-5, 3-6,
 3-7, 3-15, 4-7, 5-2,
 5-3
 Node address 1-4, 2-5, 2-8, 2-9, 3-6, 4-5,
 4-13
 Non-routing servers 4-2

P

Packet Delivery 2-1, 2-7
 Packet Type 1-6, 2-4, 3-3, 4-3, 5-2, 5-3
 Periodic broadcasts 3-2, 3-5, 3-8, 3-13, 4-2,
 4-10, 4-12

R

RIP and SAP Interoperability 3-1, 3-16, 4-1,
 4-14
 RIP Request 2-8, 3-1, 3-3, 3-4, 3-11, 3-15
 RIP Response 3-6, 3-11, 3-15
 Router Implementation Guidelines 2-1, 2-6,
 3-1, 3-6, 4-1, 4-7,
 5-1, 5-3
 Router Information Maintenance 3-1, 3-13
 Router Initialization 3-1, 3-9, 4-1, 4-8
 Router Shutdown 3-1, 3-12, 4-1, 4-11
 Routing 1-1, 1-2, 1-4, 1-5, 1-6, 2-1, 2-2, 2-4,
 2-5, 2-8, 2-9, 2-10,
 3-1, 3-2, 3-6, 3-7,
 3-6, 3-7, 3-8, 3-9,
 3-10, 3-12, 3-14,
 3-15, 4-2, A-1
 Routing Information Broadcasts 3-1, 3-8, 3-9
 Routing Information Gathering 2-1, 2-10
 Routing Information Protocol (RIP) 1-1, 1-2,
 1-4, 2-5, 3-1, 3-2
 Bandwidth Requirements A-1
 Packet Format 3-1, 3-3
 RIP Aging 3-1, 3-14
 RIP Interpacket Gap 3-1, 3-16
 RIP Operations 3-1, 3-5
 RIP Request Handling 3-1, 3-15
 Routing Information Broadcasts

 3-1, 3-8, 3-9
 Selecting the Best Route 3-1,
 3-7
 Routing Information Table 3-6, 3-7, 3-6, 3-7,
 3-9, 3-10, 3-14, 3-15

S

SAP Request 4-1, 4-3, 4-6, 4-9, 4-13
 SAP Response 4-4, 4-5, 4-6, 4-7, 4-10,
 4-12, 4-13
 Sequenced Packet Exchange (SPX) 1-2
 Server Information Maintenance 4-1, 4-11
 Server Information Table 4-1, 4-2, 4-6, 4-7,
 4-8, 4-11, 4-12
 Server name 4-5
 Service Advertising Protocol (SAP) 1-1, 1-2,
 1-5, 2-5, 4-1
 Bandwidth Requirements A-1
 Packet Format 4-1, 4-3
 SAP Agent 1-5, 4-2, 4-6, 4-7,
 4-8, 4-10, 4-11,
 4-12, 4-13, 4-14
 SAP Aging 4-1, 4-12, 4-13
 SAP Interpacket Gap 4-1, 4-14
 SAP Operations 4-1, 4-6
 SAP Request Handling 4-1,
 4-13
 Service type 4-4, 4-5, 4-6
 Socket address 2-5, 2-6, 4-6
 Socket number 2-2, 2-6, 2-8, 3-3, 4-3, 4-6
 Socket process 2-9
 Source Network 2-5, 5-2
 Source Node 2-5, 2-6
 Source Socket 2-6, 3-3, 4-3
 Specific request 3-6, 3-15, 4-13
 Split-Horizon Algorithm 3-1, 3-8, 3-9, 3-10,
 3-12, 3-14, 3-15,
 4-7, 4-8, 4-11, 4-12,
 4-13, A-1

T

Ticks 3-5, 3-7, 3-15
 Transport Control 2-3, 2-9, 2-10, 5-3
 Type 20 Propagation Packet 5-1
 Packet Format 5-1, 5-2