

Configuring Novell IPX

Novell Internet Packet Exchange (IPX) is derived from the Xerox Network Systems (XNS) Internet Datagram Protocol (IDP). IPX and XNS have the following differences:

- IPX and XNS do not always use the same Ethernet encapsulation format.
- IPX uses Novell's proprietary Service Advertisement Protocol (SAP) to advertise special network services. File servers and print servers are examples of services that are typically advertised.
- IPX uses ticks, while XNS uses hop count as the primary metric in determining the best path to a destination.

This chapter describes how to configure Novell IPX and provides configuration examples. For a complete description of the commands mentioned in this chapter, refer to the "Novell IPX Commands" chapter in the *Router Products Command Reference* publication. For historical background and a technical overview of Novell IPX, see the *Internetworking Technology Overview* publication.

Cisco's Implementation of Novell IPX

Cisco's implementation of Novell's IPX protocol has been certified as providing full IPX router functionality. A Cisco router connects Ethernet, Token Ring, and FDDI networks, either directly or through high-speed serial lines (56 kbps to T1 speeds), X.25, or Frame Relay. At this time, the Cisco X.25 and T1 support is not compatible with Novell. This means that our routers must be used on both ends of T1 and X.25 circuits.

Cisco supports the IPX MIB. The IPX Accounting group represents one of the local variables we support. This group provides access to the active database that is created and maintained if IPX accounting is enabled on a router.

Cisco routers also support IPX Enhanced IGRP, which provides the following features:

- Automatic redistribution. IPX RIP routes are automatically redistributed into Enhanced IGRP, and Enhanced IGRP routes are automatically redistributed into RIP. If desired, you can turn off redistribution. You also can completely turn off Enhanced IGRP and IPX RIP on the router or on individual interfaces.
- Increased network width. With IPX RIP, the largest possible width of your network is 15 hops. When Enhanced IGRP is enabled, the largest possible width is 224 hops. Because the Enhanced IGRP metric is large enough to support thousands of hops, the only barrier to expanding the network is the transport layer hop counter. Cisco works around this problem by incrementing the transport control field only when an IPX packet has traversed 15 routers and the next hop to the destination was learned via Enhanced IGRP. When a RIP route is being used as the next hop to the destination, the transport control field is incremented as usual.

- Incremental SAP updates. Complete SAP updates are sent periodically on each interface until an Enhanced IGRP neighbor is found and thereafter only when there are changes to the SAP table. This procedure works by taking advantage of Enhanced IGRP's reliable transport mechanism, which means that an Enhanced IGRP peer must be present for incremental SAPs to be sent. If no peer exists on a particular interface, periodic SAPs will be sent on that interface until a peer is found. This functionality is automatic on serial interfaces and can be configured on LAN media.

IPX Addresses

An IPX network address consists of a network number and a node number expressed in the format *network.node*.

The network number identifies a physical network. It is a four-byte (32-bit) quantity that must be unique throughout the entire IPX internetwork. The network number is expressed as eight hexadecimal digits. Our router software does not require that you enter all eight digits: you can omit leading zeros.

The node number identifies a node on the network. It is a 48-bit quantity, represented by dotted triplets of four-digit hexadecimal numbers.

The following is an example of an IPX network address:

```
4a.0000.0c00.23fe
```

In this example, the network number is 4a (more specifically, it is 0000004a), and the node number is 0000.0c00.23fe. All digits in the address are hexadecimal.

IPX Configuration Task List

To configure IPX routing, complete the tasks in the following sections. At a minimum, you must enable IPX routing. The remaining tasks are optional.

- Enable IPX Routing
- Control Access to IPX Networks
- Tune IPX Network Performance
- Configure IPX Accounting
- Configure IPX Enhanced IGRP
- Shut Down an IPX Network
- Configure IPX over WANs
- Monitor and Maintain the IPX Network

See the end of this chapter for configuration examples.

Enable IPX Routing

To enable IPX routing, you must perform the tasks described in the following sections:

- Enable IPX Routing on the Router
- Assign Network Numbers to Individual Interfaces

Enable IPX Routing on the Router

The first step in enabling IPX routing is to enable it on the router. If you do not specify the node number of the router, the router uses the hardware media access control (MAC) address currently assigned to it as its node address. This is the MAC address of the first Ethernet, Token Ring, or FDDI interface card.

To enable IPX routing on the router, perform the following global configuration task:

Task	Command
Enable IPX routing on the router.	ipx routing [<i>node</i>]

For an example of how to enable IPX routing, see the section Enabling IPX Routing Example.



Caution If you plan to use DECnet and IPX routing concurrently on the same interface, you should enable DECnet routing first, then enable IPX routing without specifying the optional Media Access Control (MAC) node number. If you enable IPX before enabling DECnet routing, routing for IPX will be disrupted because DECnet forces a change in the MAC-level node number.

Assign Network Numbers to Individual Interfaces

After you have enabled IPX routing on the router, you assign network numbers to individual interfaces. This has the effect of enabling IPX routing on those interfaces. When you enable IPX routing on an interface, you also can specify an encapsulation (frame type) to use for packets being transmitted on that network.

A single interface can support a single network or multiple logical networks. For a single network, you can configure any encapsulation type. Of course, it should match the encapsulation type of the servers and clients using that network number.

When assigning network numbers to an interface that supports multiple networks, you must specify a different encapsulation type for each network. Because multiple networks share the physical medium, this allows the router to determine which packets belong to which network. For example, you can configure up to four IPX networks on a single Ethernet cable, because four encapsulation types are supported for Ethernet. Again, the encapsulation type should match the servers and clients using the same network number.

The following sections describe how to enable IPX routing on interfaces that support a single network and those that support multiple networks.

Assign Network Numbers to Interfaces That Support a Single Network

To assign a network number to an interface that supports a single network, perform the following interface configuration task:

Task	Command
Enable IPX routing on an interface.	ipx network <i>network</i> [encapsulation <i>encapsulation-type</i>]

If you specify an encapsulation type, make sure you choose the one that matches that used by the servers and clients on that network.

For an example of how to enable IPX routing, see the section Enabling IPX Routing Example.

Assign Network Numbers to Interfaces that Support Multiple Networks

The first logical network you configure on an interface is considered the primary network. Any additional networks are considered secondary networks. Remember that each network on an interface must use a distinct encapsulation and that it should match the clients and servers using the same network number.

To configure multiple IPX networks on an interface, perform the following tasks in interface configuration mode:

Task	Command
Step 1 Enable IPX routing on the primary network.	ipx network <i>network</i> encapsulation <i>encapsulation-type</i>
Step 2 Enable IPX routing on a secondary network.	ipx network <i>network</i> encapsulation <i>encapsulation-type</i> secondary

To configure more than one secondary network, repeat step 2 as appropriate.

For an example of configuring multiple IPX networks on an interface, see the section Enabling and Disabling IPX Routing on Multiple Networks Example.

Table 19-1 lists the encapsulation types you can use on IEEE interfaces and shows the correspondence between the encapsulation type and the IPX frame type.

Table 19-1 Novell IPX Encapsulation Types on IEEE Interfaces

Interface Type	Encapsulation Type	IPX Frame Type
Ethernet	novell-ether (default)	Ethernet_802.3
	arpa	Ethernet_II
	sap	Ethernet_802.2
	snap	Ethernet_Snap
Token Ring	sap (default)	Token-Ring
	snap	Token-Ring_Snap
FDDI	snap (default)	Fddi_Snap
	sap	Fddi_802.2

Control Access to IPX Networks

To control access to IPX networks, you create access lists and then apply them with filters to individual interfaces.

There are four types of IPX access lists that you can use to filter various kinds of traffic:

- Standard access list—Restricts traffic based on the source network number. You can further restrict traffic by specifying a destination address and a source and destination address mask. Standard IPX access lists have numbers from 800 to 899.
- Extended access list—Restricts traffic based on the IPX protocol type. You can further restrict traffic by specifying source and destination addresses and address masks, and source and destination sockets. Extended IPX access lists have numbers from 900 to 999.
- SAP access list—Restricts traffic based on the IPX Service Advertisement Protocol (SAP) type. These lists are used for SAP filters and Get Nearest Server (GNS) response filters. Novell SAP access lists have numbers from 1000 to 1099.
- IPX NetBIOS access list—Restricts IPX NetBIOS traffic based on NetBIOS names, not numbers.

There are 13 different IPX filters that you can define for IPX interfaces. They fall into five groups:

- Generic output filters—Control which packets are routed out an interface based on the packet's source and destination addresses and IPX protocol type.
- Routing table filters—Control which Routing Information Protocol (RIP) updates are accepted and advertised by the router and which routers the local router accepts RIP updates from.
- SAP filters—Control which SAP services the router accepts and advertises and which Get Nearest Server (GNS) response messages it sends out.
- IPX NetBIOS filters—Control incoming and outgoing IPX NetBIOS packets.
- Broadcast filters—Control which broadcast packets are forwarded.

Table 19-2 summarizes the filters and the commands you use to define them. Use the **show ipx interfaces** command to display the filters defined on an interface.

Table 19-2 IPX Filters

Filter Type	Command Used to Define Filter
Generic filters	
Filter outbound packets based on protocol, address and address mask, and socket.	ipx access-group <i>access-list-number</i>
Routing table filters	
Control which networks are added to the routing table.	ipx input-network-filter <i>access-list-number</i>
Control which networks are advertised in routing updates.	ipx output-network-filter <i>access-list-number</i>
Control the routers from which updates are accepted.	ipx router-filter <i>access-list-number</i>
SAP filters	
Filter incoming service advertisements.	ipx input-sap-filter <i>access-list-number</i>
Filter outgoing service advertisements.	ipx output-sap-filter <i>access-list-number</i>
Control the routers from which SAP updates are accepted.	ipx router-sap-filter <i>access-list-number</i>
Filter list of servers in GNS response messages.	ipx output-gns-filter <i>access-list-number</i>
IPX NetBIOS filters	
Filter incoming packets by node name.	ipx netbios input-access-filter <i>host name</i>
Filter incoming packets by byte pattern.	ipx netbios input-access-filter <i>bytes name</i>
Filter outgoing packets by node name.	ipx netbios output-access-filter <i>host name</i>
Filter outgoing packets by byte pattern.	ipx netbios output-access-filter <i>bytes name</i>
Broadcast filters	
Control which broadcast packets are forwarded.	ipx helper-list <i>access-list-number</i>

Keep the following in mind when configuring IPX network access control:

- Access lists entries are scanned in the order you enter them. The first matching entry is used. To improve performance, it is recommended that you place the most commonly used entries near the beginning of the access list.

- An implicit *deny everything* entry is defined at the end of an access list unless you include an explicit *permit everything* entry at the end of the list.
- All new entries to an existing list are placed at the end of the list. You cannot add an entry to the middle of a list. This means that if you have previously included an explicit *permit everything* entry, new entries will never be scanned. The solution is to delete the access list and re-enter it with the new entries.
- Take care not to set up conditions that result in packets getting lost. One way this can happen is when a router or interface is configured to advertise services on a network that has access lists that deny these packets.

You perform the tasks in one or more of the following sections to control access to IPX networks:

- Create Access Lists
- Create Generic Filters
- Create Filters for Updating the Routing Table
- Create SAP Filters
- Create GNS Response Filters
- Create IPX NetBIOS Filters
- Create Broadcast Message Filters

Create Access Lists

To create access lists, you can perform one or more of the following tasks in global configuration mode:

Task	Command
Create a standard IPX access list (for generic, routing, and broadcast filters).	access-list <i>access-list-number</i> { deny permit } <i>source-network</i> [<i>source-node</i> [<i>source-node-mask</i>]] [<i>destination-network</i> [<i>destination-node</i> [<i>destination-node-mask</i>]]]
Create an extended IPX access list (for generic, routing, and broadcast filters).	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> [<i>source-network</i> [<i>source-node</i> [<i>source-network-mask</i> , <i>source-node-mask</i>]]] <i>source-socket</i> [<i>destination-network</i> [<i>destination-node</i> [<i>destination-network-mask</i> , <i>destination-node-mask</i>]]] <i>destination-socket</i>]
Create an IPX access list for SAP filters.	access-list <i>access-list-number</i> { deny permit } <i>network</i> [<i>node</i>] [<i>network-mask</i> <i>node-mask</i>] [<i>service-type</i> [<i>server-name</i>]]]
Create an access list for filtering IPX NetBIOS packets by node name.	netbios access-list <i>host name</i> { deny permit } <i>string</i>
Create an access list for filtering IPX NetBIOS packets by arbitrary byte pattern.	netbios access-list <i>bytes name</i> { deny permit } <i>offset</i> <i>byte-pattern</i>

Once you have created an access list, apply it to a filter on the appropriate interfaces as described in the sections that follow. This activates the access list.

Create Generic Filters

Generic filters determine which packets to send out an interface based on the packet's source and destination addresses, IPX protocol type, and source and destination socket numbers.

To create generic filters, perform the following tasks:

Step 1 Create a standard or an extended access list.

Step 2 Apply a filter to an interface.

To create an access list, perform one of the following tasks in global configuration mode:

Task	Command
Create a standard IPX access list.	access-list <i>access-list-number</i> { deny permit } <i>source-network</i> [. <i>source-node</i> [<i>source-node-mask</i>]] [<i>destination-network</i> [. <i>destination-node</i> [<i>destination-node-mask</i>]]]
Create an extended IPX access list.	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> [<i>source-network</i> [. <i>source-node</i> [<i>source-network-mask</i> . <i>source-node-mask</i>]]] <i>source-socket</i> [<i>destination-network</i> [. <i>destination-node</i> [<i>destination-network-mask</i> . <i>destination-node-mask</i>]] <i>destination-socket</i>]]]

To apply a generic filter to an interface, perform the following task in interface configuration mode:

Task	Command
Apply a generic filter to an interface.	ipx access-group <i>access-list-number</i>

For an example of creating a generic filter, see the section "IPX Network Access Example" later in this chapter.

Create Filters for Updating the Routing Table

Routing table update filters control the entries that the router accepts for its routing table and the networks that it advertises in its routing updates.

To create filters to control updating of the routing table, perform the following tasks:

Step 1 Create a standard or an extended access list.

Step 2 Apply one or more routing filters to an interface.

To create an access list, perform one of the following tasks in global configuration mode:

Task	Command
Create a standard IPX access list.	access-list <i>access-list-number</i> { deny permit } <i>source-network</i> [. <i>source-node</i> [<i>source-node-mask</i>]] [<i>destination-network</i> [. <i>destination-node</i> [<i>destination-node-mask</i>]]]
Create an extended IPX access list.	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> [<i>source-network</i> [. <i>source-node</i> [<i>source-network-mask</i> . <i>source-node-mask</i>]]] <i>source-socket</i> [<i>destination-network</i> [. <i>destination-node</i> [<i>destination-network-mask</i> . <i>destination-node-mask</i>]] <i>destination-socket</i>]]]

To apply routing table update filters to an interface, perform one or more of the following tasks in interface configuration mode:

Task	Command
Control which networks are added to the routing table when IPX routing updates are received.	ipx input-network-filter <i>access-list-number</i>
Control which networks are advertised in routing updates sent out by the router.	ipx output-network-filter <i>access-list-number</i>
Control the routers from which routing updates are accepted.	ipx router-filter <i>access-list-number</i>

You can apply one of each of these filters to each interface.

Create SAP Filters

A common source of traffic on Novell networks is SAP messages, which are generated by NetWare servers and our routers when they broadcast their available services. To control how SAP messages from network segments or specific servers are routed among IPX networks, perform the following steps:

Step 1 Create a SAP access list.

Step 2 Apply one or more filters to an interface.

To create a SAP access list, perform the following task in global configuration mode:

Task	Command
Create a SAP access list.	access-list <i>access-list-number</i> { deny permit } <i>network[.node] [network.node-mask] [service-type [server-name]]</i>

To apply SAP filters to an interface, perform one or more of the following tasks in interface configuration mode:

Task	Command
Filter incoming service advertisements.	ipx input-sap-filter <i>access-list-number</i>
Filter outgoing service advertisements.	ipx output-sap-filter <i>access-list-number</i>
Filter service advertisements received from a particular router.	ipx router-sap-filter <i>access-list-number</i>

You can apply one of each SAP filter to each interface.

For examples of creating and applying SAP filters, see the sections “SAP Input Filter Example” and “SAP Output Filter Example” later in this chapter.

Create GNS Response Filters

To create filters for controlling which servers are included in the GNS responses sent by the router, perform the following tasks:

Step 1 Create a SAP access list.

Step 2 Apply a GNS filter to an interface.

To create a SAP access list, perform the following task in global configuration mode:

Task	Command
Create a SAP access list.	access-list <i>access-list-number</i> { deny permit } <i>network</i> [, <i>node</i>] [<i>network.node-mask</i>] [<i>service-type</i>] [<i>server-name</i>]

To apply a GNS filter to an interface, perform the following task in interface configuration mode:

Task	Command
Filter the list of servers in GNS response messages.	ipx output-gns-filter <i>access-list-number</i>

Create IPX NetBIOS Filters

Novell's IPX NetBIOS allows messages to be exchanged between nodes using alphanumeric names as well as node addresses. Therefore, the router lets you filter incoming and outgoing NetBIOS packets by the node name or by an arbitrary byte pattern (such as the node address) in the packet.

Note These filters apply to IPX NetBIOS packets only. They have no effect on LLC2 NetBIOS packets.

Keep the following in mind when configuring IPX NetBIOS access control:

- Host (node) names are case sensitive.
- Host and byte access lists can have the same names because the two types of lists are independent of each other.
- When filtering by node name, the names in the access lists are compared with the destination name field for IPX NetBIOS "find name" requests.
- Access filters that filter by byte offset can have a significant impact on the packet transmission rate because each packet must be examined. You should use these access lists only when absolutely necessary.
- If a node name is not found in an access list, the default action is to deny access.

To create filters for controlling IPX NetBIOS access, perform the following tasks:

Step 1 Create a NetBIOS access list.

Step 2 Apply the access list to an interface.

To create one or more NetBIOS access lists, perform one or both of the following tasks in global configuration mode:

Task	Command
Create an access list for filtering IPX NetBIOS packets by node name.	netbios access-list host <i>name</i> {deny permit} <i>string</i>
Create an access list for filtering IPX NetBIOS packets by arbitrary byte pattern.	netbios access-list bytes <i>name</i> {deny permit} <i>offset byte-pattern</i>

To apply a NetBIOS access list to an interface, perform one or more of the following tasks in interface configuration mode:

Task	Command
Filter incoming packets by node name.	ipx netbios input-access-filter host <i>name</i>
Filter incoming packets by byte pattern.	ipx netbios input-access-filter bytes <i>name</i>
Filter outgoing packets by node name.	ipx netbios output-access-filter host <i>name</i>
Filter outgoing packets by byte pattern.	ipx netbios output-access-filter bytes <i>name</i>

You can apply one of each of these four filters to each interface.

Create Broadcast Message Filters

Routers normally block all broadcast requests and do not forward them to other network segments. This is done to prevent the degradation of performance inherent in broadcast traffic over the entire network. You can define which broadcast messages get forwarded to other networks by applying a broadcast message filter to an interface.

To create filters for controlling broadcast messages, perform the following tasks:

Step 1 Create an access list.

Step 2 Apply a broadcast message filter to an interface.

To create an access list, perform one of the following tasks in global configuration mode:

Task	Command
Create a standard IPX access list.	access-list <i>access-list-number</i> {deny permit} <i>source-network</i> [<i>.source-node</i> [<i>source-node-mask</i>]] [<i>destination-network</i> [<i>.destination-node</i> [<i>destination-node-mask</i>]]]
Create an extended IPX access list.	access-list <i>access-list-number</i> {deny permit} <i>protocol</i> [<i>source-network</i> [<i>.source-node</i> [<i>source-network-mask</i> . <i>source-node-mask</i>]]] <i>source-socket</i> [<i>destination-network</i> [<i>.destination-node</i> [<i>destination-network-mask</i> . <i>destination-node-mask</i>]]] <i>destination-socket</i>]

To apply a broadcast message filter to an interface, perform the following tasks in interface configuration mode:

Task	Command
Step 1 Specify a helper address for forwarding broadcast messages.	ipx helper-address <i>network.node</i>
Step 2 Apply a broadcast message filter to an interface.	ipx helper-list <i>access-list-number</i>

Note A broadcast message filter has no effect unless you have issued an **ipx helper-address** and/or an **ipx type-20-propagation** command on the interface to enable and control the forwarding of broadcast messages. These commands are discussed later in this chapter.

For examples of creating and applying broadcast message filters, see the section “Helper Facilities to Control Broadcasts Examples” later in this chapter.

Tune IPX Network Performance

To tune IPX network performance, perform the tasks in one of more of the following sections:

- Control Novell IPX Compliance
- Configure Static Routes
- Adjust Routing Table Update Timers
- Configure Static SAP Table Entries
- Configure the Queue Length for SAP Requests
- Adjust SAP Update Timers
- Set Maximum Paths
- Control Responses to GNS Requests
- Use Helper Addresses to Forward Broadcast Messages
- Control the Forwarding of Type 20 Packets
- Disable IPX Fast Switching
- Disable IPX Fast Switching
- Enable Autonomous Switching
- Enable SSE Switching
- Pad Odd-Length Packets
- Repair Corrupted Network Numbers

Control Novell IPX Compliance

Cisco’s implementation of Novell’s IPX protocol has been certified as providing full IPX router functionality, as defined by Novell’s IPX Router Specification, Version 1.10, published November 17, 1992.

To control specific aspects of IPX compliance, you can use a combination of global configuration and interface configuration commands. You can perform one or more of the following tasks in global configuration mode:

Task	Command
Restrict the acceptance of IPX type 20 propagation packets.	ipx type-20-input-checks
Restrict the forwarding of IPX type 20 propagation packets.	ipx type-20-output-checks
Set the interpacket delay of multiple-packet routing updates sent on all interfaces.	ipx default-output-rip-delay <i>delay</i>
Set the interpacket delay of multiple-packet triggered routing updates sent on all interfaces.	ipx default-triggered-rip-delay <i>delay</i>
Set the interpacket delay of multiple-packet SAP updates sent on all interfaces.	ipx default-output-sap-delay <i>delay</i>
Set the interpacket delay of multiple-packet triggered SAP updates sent on all interfaces.	ipx default-triggered-sap-delay <i>delay</i>

You can perform one or more of the following tasks in interface configuration mode:

Task	Command
Set the tick count, which is used in the IPX Routing Information Protocol (RIP) delay field.	ipx delay <i>number</i>
Administratively bring down an IPX network on an interface. This removes the network from the interface.	ipx down <i>network</i>
Set the interpacket delay of multiple-packet routing updates sent on a single interface.	ipx output-rip-delay <i>delay</i>
Set the interpacket delay of multiple-packet triggered routing updates sent on a single interface.	ipx triggered-rip-delay <i>delay</i>
Set the interpacket delay of multiple-packet SAP updates sent on a single interface.	ipx output-sap-delay <i>delay</i>
Set the interpacket delay of multiple-packet triggered SAP updates sent on a single interface.	ipx triggered-sap-delay <i>delay</i>
Forward IPX type 20 propagation packets to other network segments.	ipx type-20-propagation

Note We recommend that you use an **ipx output-rip-delay** and **ipx output-sap-delay** on slower speed WAN interfaces.

To achieve full compliance, issue the following interface configuration commands on each interface configured for IPX:

Task	Command
Step 1 Set the interpacket delay of multiple-packet routing updates to 55 ms.	ipx output-rip-delay 55
Step 2 Set the interpacket delay of multiple-packet SAP updates to 55 ms.	ipx output-sap-delay 55
Step 3 Optionally enable type 20 packet propagation if you want to forward type 20 broadcast traffic across the router.	ipx type-20-propagation

You can also globally set interpacket delays for multiple-packet RIP and SAP updates to achieve full compliance, eliminating the need to set delays on each interface. To do so, issue the following commands from global configuration mode:

Task	Command
Step 1 Set the interpacket delay of multiple-packet routing updates sent on all interfaces to 55 ms.	ipx default-output-rip-delay 55
Step 2 Set the interpacket delay of multiple-packet SAP updates sent on all interfaces to 55 ms.	ipx default-output-sap-delay 55

Configure Static Routes

IPX uses the Routing Information Protocol (RIP) to determine the best path when several paths to a destination exist. RIP then dynamically updates the routing table. However, you might want to add static routes to the routing table to explicitly specify paths to certain destinations. Static routes always override any dynamically learned paths.

Be careful when assigning static routes. When links associated with static routes are lost, traffic may stop being forwarded or traffic may be forwarded to a nonexistent destination, even though an alternative path might be available.

To add a static route to the router's routing table, perform the following task in global configuration mode:

Task	Command
Add a static route to the routing table.	ipx route <i>network network.node</i>

Adjust Routing Table Update Timers

You can set the interval between IPX RIP updates on a per-interface basis. You can also specify the delay between the packets of a multiple-packet RIP update on a per-interface or global basis. Additionally, you can specify the delay between packets of a multiple-packet triggered RIP update on a per-interface or global basis.

You can set RIP update times only in a configuration in which all routers are our routers or in which the IPX routers allow configurable timers. The timers for all routers connected to the same network segment should be the same. The RIP update value you choose affects internal IPX timers as follows:

- IPX routes are marked invalid if no routing updates are heard within three times the value of the update interval ($3 * interval$) and are advertised with a metric of infinity.
- IPX routes are removed from the routing table if no routing updates are heard within four times the value of the update interval ($4 * interval$).
- If you define an update timer for more than one interface in a router, the granularity of the update timer is determined by the lowest value defined for one of the interfaces in the router. The router “wakes up” at this granularity interval and sends out updates as appropriate. For more information about granularity, see the “Novell IPX Commands” chapter in the *Router Products Command Reference* publication.

You might want to set a delay between the packets in a multiple-packet update if there are some slower PCs on the network or on slower speed interfaces.

To adjust RIP update times on a per-interface basis, perform any or all of the following tasks in interface configuration mode:

Task	Command
Adjust the RIP update interval.	ipx update-time <i>interval</i>
Adjust the delay between multiple-packet routing updates sent on a single interface.	ipx output-rip-delay <i>delay</i>
Adjust the delay between multiple-packet triggered routing updates sent on a single interface.	ipx triggered-rip-delay <i>delay</i>

To adjust RIP update times on a global basis, perform any or all of the following tasks in global configuration mode:

Task	Command
Adjust the delay between multiple-packet routing updates sent on all interfaces.	ipx default-output-rip-delay <i>delay</i>
Adjust the delay between multiple-packet triggered routing updates sent on all interfaces.	ipx default-triggered-rip-delay <i>delay</i>

Configure Static SAP Table Entries

Servers use SAP to advertise their services via broadcast packets. Routers store this information in the SAP table, also known as the Server Information Table (SIT). This table is updated dynamically. You might want to explicitly add an entry to the SIT so that clients always use the services of a particular server. Static SAP assignments always override any identical entries in the SAP table that are learned dynamically, regardless of hop count. If a dynamic route that is associated with a static SAP entry is lost or deleted, the router will not announce the static SAP entry until it relearns the route.

To add a static entry to the router’s SAP table, perform the following task in global configuration mode:

Task	Command
Specify a static SAP table entry.	ipx sap <i>service-type name network.node socket hop-count</i>

Configure the Queue Length for SAP Requests

The router maintains a list of SAP requests to process, including all pending GNS queries from clients attempting to reach servers. When the network is restarted, the router can be inundated with hundreds of requests for servers. Typically, many of these are repeated requests from the same clients. You can configure the maximum length allowed for the pending SAP requests queue. SAP requests received when the queue is full are dropped, and the client must resend them.

To set the queue length for SAP requests, perform the following task in global configuration mode:

Task	Command
Configure the maximum SAP queue length.	ipx sap-queue-maximum <i>number</i>

Adjust SAP Update Timers

You can adjust the interval at which SAP updates are sent. You can also set the delay between packets of a multiple-packet SAP update on a per-interface or global basis. Additionally, you can specify the delay between packets of a multiple-packet triggered SAP update on a per-interface or global basis.

Changing the interval at which SAP updates are sent is most useful on limited-bandwidth, point-to-point links or on X.25 and Frame Relay multipoint interfaces. You should ensure that all Novell servers and routers on a given network have the same SAP interval. Otherwise, they might decide that a server is down when it is really up.

Adjusting the delay between packets sent in a multiple-packet SAP update is useful when the IPX network has slow IPX servers and/or routers. Setting a delay between packets in a multiple-packet SAP update forces our router to slow its output of SAP packets.

To modify the SAP timers on a per-interface basis, perform any or all of the following tasks in interface configuration mode:

Task	Command
Adjust the interval at which SAP updates are sent.	ipx sap-interval <i>interval</i>
Adjust the interpacket delay of multiple-packet SAP updates sent on a single interface.	ipx output-sap-delay <i>delay</i>
Adjust the interpacket delay of multiple-packet triggered SAP updates sent on a single interface.	ipx triggered-sap-delay <i>delay</i>

To adjust SAP update times on a global basis (eliminating the need to configure delays on a per interface basis), perform any or all of the following tasks in global configuration mode:

Task	Command
Adjust the interpacket delay of multiple-packet SAP updates sent on all interfaces.	ipx default-output-sap-delay <i>delay</i>
Adjust the interpacket delay of multiple-packet triggered SAP updates sent on all interfaces.	ipx default-triggered-sap-delay <i>delay</i>

Set Maximum Paths

You can set the maximum number of equal-cost, parallel paths to a destination. (Note that when paths have differing costs, the router chooses lower-cost routes in preference to higher-cost routes.) The router then distributes output on a packet-by-packet basis in round-robin fashion. That is, the first packet is sent along the first path, the second packet along the second path, and so on. When the final path is reached, the next packet is sent to the first path, the next to the second path, and so on. This round-robin scheme is used whether or not fast switching is enabled.

The cost of a path is determined by ticks, with hop count used as a tie breaker.

Limiting the number of equal-cost paths can save memory on routers with limited memory or very large configurations. Additionally, in networks with a large number of multiple paths and systems with limited ability to cache out-of-sequence packets, performance might suffer when traffic is split between many paths.

To set the maximum number of paths on the router, perform the following task in global configuration mode:

Task	Command
Set the maximum number of equal-cost paths to a destination.	ipx maximum-paths <i>paths</i>

Control Responses to GNS Requests

You can set the method in which the router responds to SAP GNS requests, and you can set the delay time in responding to these requests.

The default method of responding to GNS requests is to respond with the server whose availability was learned most recently.

To control responses to GNS requests, perform one or both of the following tasks in global configuration mode:

Task	Command
Respond to GNS requests using a round-robin selection method.	ipx gns-round-robin
Set the delay when responding to GNS requests.	ipx gns-response-delay [<i>milliseconds</i>]

Use Helper Addresses to Forward Broadcast Messages

Routers normally block all broadcast requests and do not forward them to other network segments. This is done to prevent the degradation of performance over the entire network. You can enable the forwarding of broadcast messages (except type 20 broadcasts) to other networks and forward all other unrecognized broadcast messages. These are non-RIP and non-SAP packets that are not addressed to the local network. Forwarding broadcast messages is sometimes useful when a network segment does not have an end-host capable of servicing a particular type of broadcast request. You can specify the address of a server, network, or networks that can process the broadcast messages.

Our routers support all-networks flooded broadcasts (sometimes referred to as *all-nets flooding*). These are broadcast messages that are forwarded to all networks. Use all-nets flooding carefully and only when necessary, because the receiving networks may be overwhelmed to the point that no other traffic can traverse them.

Use the **ipx helper-list** command, described earlier in this chapter, to define access lists that control which broadcast packets get forwarded.

To specify a helper address for forwarding broadcast messages, perform the following task in interface configuration mode:

Task	Command
Specify a helper address for forwarding broadcast messages.	ipx helper-address <i>network.node</i>

You can specify multiple helper addresses on an interface.

For an example of using helper addresses to forward broadcast messages, see the section “Helper Facilities to Control Broadcasts Examples” later in this chapter.

Control the Forwarding of Type 20 Packets

NetBIOS over IPX uses type 20 propagation broadcast packets flooded to all networks to get information about the named nodes on the network. NetBIOS uses a broadcast mechanism to get this information, because it does not implement a network layer.

Routers normally block all broadcast requests. By enabling type 20 packet propagation, IPX interfaces on the router may accept and forward type 20 propagation packets. Before forwarding (flooding) the packets, the router performs loop detection as described by the IPX router specification.

You can configure the router to apply extra checks to type 20 propagation packets above and beyond the loop detection described in the IPX specification. These checks are the same ones that are applied to helpered all-nets broadcast packets. They can limit unnecessary duplication of type 20 broadcast packets. The extra helper checks are as follows:

- Accept type 20 propagation packets only on the primary network, which is the network that is the primary path back to the source network.
- Forward type 20 propagation packets only via networks that do not lead back to the source network.

While this extra checking increases the robustness of type 20 propagation packet handling by decreasing the amount of unnecessary packet replication, it has two side effects:

- If type 20 packet propagation is not configured on all interfaces, these packets might be blocked when the primary interface changes.
- It might be impossible to configure an arbitrary, manual spanning tree for type 20 packet propagation.

You can enable the forwarding of type 20 packets on individual interfaces, and you can restrict the acceptance and forwarding of type 20 packets. The tasks to do this are described in the following sections.

Enable the Forwarding of Type 20 Packets

By default, type 20 propagation packets are dropped by the router. You can configure the router to receive type 20 propagation broadcast packets and forward (flood) them to other network segments, subject to loop detection.

To enable the receipt and forwarding of type 20 packets, perform the following task in interface configuration mode:

Task	Command
Forward IPX type 20 propagation packet broadcasts to other network segments.	ipx type-20-propagation

Restrict the Acceptance of Incoming Type 20 Packets

For incoming type 20 propagation packets, the router is configured by default to accept packets on all interfaces enabled to receive type 20 propagation packets. You can configure the router to accept packets only from the single network that is the primary route back to the source network. This means that similar packets from the same source that are received via other networks will be dropped.

Checking of incoming type 20 propagation broadcast packets is done only if the interface is configured to receive and forward type 20 packets.

To impose restrictions on the receipt of incoming type 20 propagation packets in addition to the checks defined in the IPX specification, perform the following global configuration task:

Task	Command
Restrict the acceptance of IPX type 20 propagation packets.	ipx type-20-input-checks

Restrict the Forwarding of Outgoing Type 20 Packets

For outgoing type 20 propagation packets, the router is configured by default to send packets on all interfaces enabled to send type 20 propagation packets, subject to loop detection. You can configure the router to send these packets only to networks that are not routes back to the source network. (The router uses the current routing table to determine routes.)

Checking of outgoing type 20 propagation broadcast packets is done only if the interface is configured to receive and forward type 20 packets.

To impose restrictions on the transmission of type 20 propagation packets and to forward these packets to all networks using only the checks defined in the IPX specification, perform the following global configuration task:

Task	Command
Restrict the forwarding of IPX type 20 propagation packets.	ipx type-20-output-checks

Disable IPX Fast Switching

Fast switching allows higher throughput by switching a packet using a cache created by previous packets. Fast switching is enabled by default on all interfaces.

Packet transfer performance is generally better when fast switching is enabled. However, you might want to disable fast switching in order to save memory space on interface cards and to help avoid congestion when high-bandwidth interfaces are writing large amounts of information to low-bandwidth interfaces.

To disable IPX fast switching, perform the following task in interface configuration mode:

Task	Command
Disable IPX fast switching.	no ipx route-cache

Enable Autonomous Switching

Autonomous switching provides faster packet switching by allowing the ciscoBus controller to switch packets independently without having to interrupt the system processor. It is available only in Cisco 7000 systems and in AGS+ systems with high-speed network controller cards, such as the CSC-HSCI, CSC-MEC, CSC-FCI, CSC-C2FCIT, and CSC-C2CTR, and with a CSC-CCTL2 ciscoBus controller running Microcode Version 11.0 or later. Autonomous switching is disabled by default on all interfaces.

To enable autonomous switching, perform the following task in interface configuration mode:

Task	Command
Enable autonomous switching.	ipx route-cache cbus

Enable SSE Switching

The silicon switching engine (SSE) is on the Silicon Switch Processor (SSP) board in the Cisco 7000 series. SSE switching contributes to very fast packet processing by allowing the SSE to perform switching independently of the system processor.

To enable SSE switching, perform the following task in interface configuration mode:

Task	Command
Enable the SSE switching cache.	ipx route-cache sse

Pad Odd-Length Packets

Some IPX end hosts reject Ethernet packets that are not padded to be an even length. Certain topologies can result in such packets being forwarded onto a remote Ethernet network. Under specific conditions, you can use padding on intermediate media as a temporary workaround for this problem.

To enable the padding of odd-length packets, perform the following tasks in interface configuration mode:

Task	Command
Step 1 Disable fast switching.	no ipx route-cache
Step 2 Enable the padding of odd-length packets.	ipx pad-process-switched-packets

Repair Corrupted Network Numbers

To repair corrupted network numbers on an interface, perform the following tasks in interface configuration mode:

Task	Command
Step 1 Disable fast switching.	no ipx route-cache
Step 2 Repair corrupted network numbers.	ipx source-network-update



Caution The **ipx source-network-update** interface configuration command interferes with the proper working of OS/2 Requestors. Do not use this command in a network that has OS/2 Requestors.



Caution Do not use the **ipx source-network-update** interface configuration command on interfaces on which NetWare servers are using internal network numbers (that is, all 3.1x and 4.0 servers).

Configure IPX Accounting

IPX accounting allows you to collect information about IPX packets and the number of bytes that are switched through the router. You collect information based on the source and destination IPX address. Accounting tracks only IPX traffic that is routed through the router; it does not track traffic generated by or terminating at the router.

IPX accounting statistics are accurate even if IPX fast switching is enabled or if IPX access lists are being used. However, IPX accounting does not keep statistics if autonomous switching is enabled.

The router software maintains two accounting databases: an active database and a checkpointed database.

To enable IPX accounting, perform the following task in interface configuration mode:

Task	Command
Enable IPX accounting.	ipx accounting

To control IPX accounting on the router, perform one or more of the following tasks in global configuration mode:

Task	Command
Set the maximum number of accounting entries.	ipx accounting-threshold <i>threshold</i>
Set the maximum number of transit entries.	ipx accounting-transits <i>count</i>
Filter the networks for which IPX accounting information is kept.	ipx accounting-list <i>number mask</i>

Configure IPX Enhanced IGRP

Enhanced IGRP is an enhanced version of the Interior Gateway Routing Protocol (IGRP) developed by Cisco Systems, Inc. Enhanced IGRP uses the same distance vector algorithm and distance information as IGRP. However, the convergence properties and the operating efficiency of Enhanced IGRP have improved significantly over IGRP.

The convergence technology is based on research conducted at SRI International and employs an algorithm referred to as the Diffusing Update Algorithm (DUAL). This algorithm guarantees loop-free operation at every instant throughout a route computation and allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations. The convergence time with DUAL rivals that of any other existing routing protocol.

Enhanced IGRP offers the following features:

- Fast convergence. The DUAL algorithm allows routing information to converge as quickly as any currently available routing protocol.
- Partial updates. Enhanced IGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature minimizes the bandwidth required for Enhanced IGRP packets.
- Less CPU usage than IGRP. This occurs because full update packets do not have to be processed each time they are received.
- Neighbor discovery mechanism. This is a simple hello mechanism used to learn about neighboring routers. It is protocol-independent.
- Scaling. Enhanced IGRP scales to large networks.

Enhanced IGRP has four basic components:

- Neighbor discovery/recovery
- Reliable transport protocol
- DUAL finite state machine
- Protocol-dependent modules

Neighbor discovery/recovery is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery/recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, a router can determine that a neighbor is alive and functioning. Once this status is determined, the neighboring routers can exchange routing information.

The reliable transport protocol is responsible for guaranteed, ordered delivery of Enhanced IGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some Enhanced IGRP packets must be transmitted reliably and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities, such as Ethernet, it is not necessary to send hellos reliably to all neighbors individually. Therefore, Enhanced IGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets, such as updates, require acknowledgment, and this is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. Doing so helps ensure that convergence time remains low in the presence of varying speed links.

The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information, known as a metric, to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur. This is the process whereby a new successor is determined. The amount of time it takes to recompute the route affects the convergence time. Even though the recomputation is not processor intensive, it

is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, it will use any it finds in order to avoid unnecessary recomputation.

The protocol-dependent modules are responsible for network layer protocol-specific tasks. It is also responsible for parsing Enhanced IGRP packets and informing DUAL of the new information received. Enhanced IGRP asks DUAL to make routing decisions, but the results are stored in the IPX routing table. Also, Enhanced IGRP is responsible for redistributing routes learned by other IPX routing protocols.

IPX Enhanced IGRP Configuration Task List

tasks in the following section. At a minimum, you must enable Enhanced IGRP. The remaining tasks are optional.

- Enable IPX Enhanced IGRP
- Configure Miscellaneous Enhanced IGRP Parameters

Enable IPX Enhanced IGRP

To create an IPX Enhanced IGRP routing process, perform the following tasks:

Task	Command
Step 1 Enable an Enhanced IGRP routing process in global configuration mode.	ipx router eigrp <i>autonomous-system-number</i>
Step 2 Enable Enhanced IGRP on a network in IPX router configuration mode.	network { <i>network-number</i> all }

For an example of how to enable Enhanced IGRP, see the section Enabling IPX Enhanced IGRP Example.

To associate multiple networks with an Enhanced IGRP routing process, you can repeat step 2.

Configure Miscellaneous Enhanced IGRP Parameters

To configure the following miscellaneous Enhanced IGRP parameters, perform one or more of the following tasks:

- Redistribute Routing Information
- Adjust the Interval between Hello Packets and the Hold Time
- Disable Split Horizon
- Control SAP Updates
- Control the Advertising of Routes in Routing Updates
- Control the Processing of Routing Updates
- Query the Backup Server

Redistribute Routing Information

By default, the router redistributes IPX RIP routes into Enhanced IGRP, and vice versa. When routes are redistributed, a RIP route to a destination with a hop count of 1 is always preferred over an Enhanced IGRP route with a hop count of 1. This ensures that the router always believes a Novell IPX server over a Cisco router for internal IPX networks. The only exception to this rule is if both the RIP and Enhanced IGRP updates were received from the same router. In this case, and in the case of all other RIP metrics (2 through 15), the Enhanced IGRP route always is preferred over the RIP route when the hop counts are the same.

Internal Enhanced IGRP routes are always preferred over external Enhanced IGRP routes. This means that if there are two Enhanced IGRP paths to a destination, the path that originated within the Enhanced IGRP autonomous system will always be preferred over the Enhanced IGRP path that originated from outside of the autonomous system, regardless of the metric. Redistributed RIP routes are always advertised in Enhanced IGRP as external.

To disable route redistribution, perform the following task in IPX router configuration mode:

Task	Command
Disable redistribution of RIP routes into Enhanced IGRP and Enhanced IGRP routes into RIP.	no redistribute { rip eigrp <i>autonomous-system-number</i> connected static }

Adjust the Interval between Hello Packets and the Hold Time

You can adjust the interval between hello packets and the hold time.

Routers periodically send hello packets to each other to dynamically learn of other routers on their directly attached networks. Routers use this information to discover who their neighbors are and to discover when their neighbors become unreachable or inoperative. By default, hello packets are sent every 5 seconds.

You can configure the hold time, in seconds, on a specified interface for the Enhanced IGRP routing process designated by the autonomous system number. The hold time is advertised in hello packets and indicates to neighbors the length of time they should consider the sender valid. The default hold time is three times the hello interval, or 15 seconds.

To change the interval between hello packets, perform the following task in interface configuration mode:

Task	Command
Set the interval between hello packets.	ipx hello-interval eigrp <i>autonomous-system-number</i> <i>seconds</i>

On very congested and large networks, 15 seconds may not be sufficient time for all routers to receive hello packets from their neighbors. In this case, you may want to increase the hold time. To do this, perform the following task in interface configuration mode:

Task	Command
Set the hold time.	ipx hold-time eigrp <i>autonomous-system-number</i> <i>seconds</i>

Note Do not adjust the hold time without advising technical support.

Disable Split Horizon

Split horizon controls the sending of Enhanced IGRP update and query packets. If split horizon is enabled on an interface, these packets are not sent for destinations if this interface is the next hop to that destination.

By default, split horizon is enabled on all interfaces.

Split horizon blocks information about routes from being advertised by a router out any interface from which that information originated. This behavior usually optimizes communication among multiple routers, particularly when links are broken. However, with nonbroadcast networks, such as Frame Relay and SMDS, situations can arise for which this behavior is less than ideal. For these situations, you may wish to disable split horizon.

To disable split horizon, perform the following task in interface configuration mode:

Task	Command
Disable split horizon.	no ipx split-horizon eigrp <i>autonomous-system-number</i>

Control SAP Updates

If IPX Enhanced IGRP peers are found on an interface, you can configure the router to send SAP updates either periodically or when a change occurs in the SAP table. When no IPX Enhanced IGRP peer is present on the interface, periodic SAPs are always sent.

On serial lines, by default, if an Enhanced IGRP neighbor is present, the router sends SAP updates only when the SAP table changes. On Ethernet, Token Ring, and FDDI interfaces, by default, the router sends SAP updates periodically. To reduce the amount of bandwidth required to send SAP updates, you might want to disable the periodic sending of SAP updates on LAN interfaces. Do this only when all nodes out this interface are Enhanced IGRP peers; otherwise, loss of SAP information on the other nodes will result.

To send SAP updates only when a change occurs in the SAP table, perform the following task in interface configuration mode:

Task	Command
Send SAP updates only when a change in the SAP table occurs, and send SAP changes only.	ipx sap-incremental eigrp <i>autonomous-system-number</i> rsup-only

To send periodic SAP updates, perform the following task in interface configuration mode:

Task	Command
Send SAP updates periodically.	no ipx sap-incremental eigrp <i>autonomous-system-number</i>

For an example of how to configure SAP updates, see the section “Enhanced IGRP SAP Update Examples” later in this chapter.

Control the Advertising of Routes in Routing Updates

To control which routers learn about routes, you can control the advertising of routes in routing updates. To do this, perform the following task in router configuration mode:

Task	Command
Control the advertising of routes in routing updates.	distribute-list <i>access-list-number</i> out [<i>interface-name</i> <i>routing-process</i>]

Control the Processing of Routing Updates

To control the processing of routes listed in incoming updates, perform the following task in router configuration mode:

Task	Command
Control which incoming route updates are processed.	distribute-list <i>access-list-number</i> in [<i>interface-name</i>]

Query the Backup Server

The backup server table is a table kept for each Enhanced IGRP peer. It lists the IPX servers that have been advertised by that peer. If a server is removed from the main server table at any time and for any reason, the router examines the backup server table to see if this just-removed server is known by any of the Enhanced IGRP peers. If it is, the information from that peer is advertised back into the main server table just as if that peer had readvertised the server information to this router. Using this method to allow the router to keep the backup server table consistent with what is advertised by each peer means that only changes to the table need to be advertised between Enhanced IGRP routers; full periodic updates do not need to be sent.

By default, the router queries its own copy of each Enhanced IGRP neighbor's backup server table every 15 seconds. To change this interval, perform the following global configuration task:

Task	Command
Specify the minimum period of time between successive queries of a neighbor's backup server table.	ipx backup-server-query-interval <i>interval</i>

Shut Down an IPX Network

You can administratively shut down an IPX network in two ways. In the first way, the network still exists in the configuration, but is not active. When shutting down, the network sends out update packets informing its neighbors that it is shutting down. This allows the neighboring systems to update their routing, SAP, and other tables without having to wait for routes and services learned via this network to time out.

To shut down an IPX network such that the network still exists in the configuration, perform the following task in interface configuration mode:

Task	Command
Shut down an IPX network but have the network still exist in the configuration.	ipx down <i>network</i>

In the second way, you shut down an IPX network and remove it from the configuration. To do this, perform one of the following tasks in interface configuration mode:

Task	Command
Shut down an IPX network and remove it from the configuration.	no ipx network
When multiple networks are configured on an interface, shut down all networks and remove them from the interface.	no ipx network network (where <i>network</i> is 1, the primary interface)
When multiple networks are configured on an interface, shut down one of the secondary networks and remove it from the interface.	no ipx network network (where <i>network</i> is the number of the secondary interface [not 1])

When multiple networks are configured on an interface and you want shut down one of the secondary networks and remove it from the interface, perform the second task in the previous table specifying the network number of one of the secondary networks.

For an example of shutting down an IPX network, see the section “Enabling IPX Routing Example” later in this chapter.

Configure IPX over WANs

You can configure IPX over dial-on-demand routing (DDR), Frame Relay, Point-to-Point Protocol (PPP), Switched Multimegabit Data Service (SMDS), and X.25 networks. To do this, you configure address mappings as described in the appropriate chapter. You can also configure IPX over Point-to-Point Protocol (PPP); address maps are not necessary for this protocol. You can fast switch IPX serial interfaces configured for Frame Relay, and you can fast switch SNAP-encapsulated packets over interfaces configured for ATM.

Additionally, you can configure the IPXWAN protocol.

Configure IPX over DDR

IPX sends periodic watchdog (keepalive) packets. These are keepalive packets that are sent from servers to clients after a client session has been idle for approximately 5 minutes. On a DDR link, this means that a call would be made every 5 minutes, regardless of whether there were data packets to send. You can prevent these calls from being made by configuring the router to respond to the server’s watchdog packets on a remote client’s behalf. This is sometimes referred to as “spoofing the server.”

When configuring IPX over DDR, you might want to disable the generation of these packets so that a call is not made every 5 minutes. This is not an issue for the other WAN protocols, because they establish dedicated connections rather than establishing connections only as needed.

To keep the serial interface idle when only watchdog packets are being sent, refer to the tasks described in the “Configuring DDR” chapter. For an example of configuring IPX over DDR, see the section “IPX over DDR Example” later in this chapter.

Configure the IPXWAN Protocol

Our routers support the IPXWAN protocol, as defined in RFC 1362. IPXWAN allows two routers that are running IPX routing to connect via a serial link to another router, possibly from another manufacturer, that is also running IPX routing and using IPXWAN.

You can use the IPXWAN protocol over PPP. You can also use it over HDLC; however, the routers at both ends of the serial link must be our routers.

To configure IPXWAN, perform the following tasks in interface configuration mode on a serial interface:

Task	Command
Step 1 Ensure that you have not configured an IPX network number on the interface.	no ipx network <i>number</i>
Step 2 Enable PPP.	encapsulation ppp ¹
Step 3 Enable IPXWAN.	ipx ipxwan <i>local-node number local-server</i> [<i>retry-interval</i>] [<i>retry-attempts</i>]

1. This command is documented in the “Interface Commands” chapter of the *Router Products Command Reference* publication.

Monitor and Maintain the IPX Network

To monitor and maintain a Novell IPX network, perform one or more of the following tasks at the EXEC prompt:

Task	Command
Delete all entries in the IPX accounting or accounting checkpoint database.	clear ipx accounting [checkpoint]
Delete all entries in the IPX fast-switching cache.	clear ipx cache
Delete entries in the IPX routing table.	clear ipx route [<i>network</i> *]
Have the Cisco 7000 route processor recompute the IPX SSE fast-switching cache.	clear ipx sse
Reinitialize the route processor on the Cisco 7000.	clear sse
Diagnose basic IPX network connectivity (user-level command).	ping ipx { <i>host</i> <i>address</i> }
Diagnose basic IPX network connectivity (privileged command).	ping [ipx] [<i>host</i> <i>address</i>]
List the entries in the IPX accounting or accounting checkpoint database.	show ipx accounting [checkpoint]
List the entries in the IPX fast-switching cache.	show ipx cache
List the neighbors discovered by Enhanced IGRP.	show ipx eigrp neighbors [servers] [<i>autonomous-system-number</i> <i>type unit</i>]
Display the contents of the Enhanced IGRP topology table.	show ipx eigrp topology [<i>network-number</i>]
Display the status of the IPX interfaces configured in the router and the parameters configured on each interface.	show ipx interface [<i>type number</i>]
List the entries in the IPX routing table.	show ipx route [<i>network</i>]

Task	Command
List the servers discovered through SAP advertisements.	show ipx servers [sorted [name net type]]
Display information about the number and type of IPX packets transmitted and received.	show ipx traffic
Display a summary of SSP statistics	show sse summary

Configuration Examples

This section provides configuration examples for the following IPX configuration situations:

- Enabling IPX Routing Example
- Enabling and Disabling IPX Routing on Multiple Networks Example
- Enabling and Disabling IPX Routing Protocols Examples
- Enabling IPX over a WAN Interface Example
- IPX over DDR Example
- IPX Network Access Example
- SAP Input Filter Example
- SAP Output Filter Example
- IPX NetBIOS Filter Examples
- Helper Facilities to Control Broadcasts Examples
- IPX Accounting Example
- Enabling IPX Enhanced IGRP Example
- Enhanced IGRP SAP Update Examples

Enabling IPX Routing Example

The following configuration commands enable IPX routing, defaulting the IPX host address to that of the first IEEE-conformance interface (in this example, Ethernet 0). Routing is then enabled on Ethernet 0 and Ethernet 1 for IPX networks 2abc and 1def, respectively.

```
ipx routing
interface ethernet 0
ipx network 2abc
interface ethernet 1
ipx network 1def
```

Enabling and Disabling IPX Routing on Multiple Networks Example

The following example creates four networks on Ethernet interface 0:

```
interface ethernet 0
ipx network 1
ipx encapsulation novell-ether
ipx network 2 encapsulation snap secondary
ipx network 3 encapsulation arpa secondary
ipx network 4 encapsulation isol secondary
```

Any configuration parameters that you specify on this interface are applied to all the logical networks. For example, if you set the routing update timer to 120 seconds, this value is used on all four networks.

If you administratively bring down Ethernet interface 0 using the **shut** interface configuration command, all four networks are shut down. You cannot bring down each network independently using the **shut** command; however, you can do this using the **ipx down** command.

To bring down network 1, use the following command:

```
ipx down 1
```

To shut down all four networks on the interface and remove all the networks on the interface, use one of the following commands:

```
no ipx network
```

```
no ipx network 1
```

To remove one of the secondary networks on the interface (in this case, network 2), use this command:

```
no ipx network 2
```

Enabling and Disabling IPX Routing Protocols Examples

Two routing protocols can run over interfaces configured for IPX: RIP and Enhanced, IGRP. This section provides examples of how to enable and disable various combinations of these routing protocols.

When you enable IPX routing with the **ipx routing** global configuration command, the RIP routing protocol is automatically enabled. The following example enables RIP on networks 1 and 2:

```
ipx routing
!
interface ethernet 0
ipx network 1
!
interface ethernet 1
ipx network 2
```

The following example enables RIP on networks 1 and 2 and Enhanced IGRP on network 1:

```
ipx routing
!
interface ethernet 0
ipx network 1
!
interface ethernet 1
ipx network 2
!
ipx router eigrp 100
network 1
```

The following example enables RIP on network 2 and Enhanced IGRP on network 1:

```
ipx routing
!
interface ethernet 0
ipx network 1
!
interface ethernet 1
ipx network 2
```

```

!
ipx router eigrp 100
ipx network 1
!
ipx router rip
no ipx network 1

```

Enabling IPX over a WAN Interface Example

When you configure the router to transport IPX packets over a serial interface that is running a WAN protocol such as X.25 or PPP, you specify how the packet will be encapsulated for transport. This encapsulation is not the same as the encapsulation used on an IPX LAN interface. Figure 19-1 illustrates IPX over a WAN interface.




Figure 19-1 IPX over a WAN Interface

The following examples configure a serial interface for X.25 encapsulation and for several IPX subinterfaces used in a nonmeshed topology.

Configuration for Main Router

```

hostname Main
!
no ip routing
novell routing 0000.0c17.d726
!
interface Ethernet0
no ip address
Novell network 100
media-type 10BaseT
!
interface Serial0
no ip address
shutdown
!

```

```
interface Serial1
no ip address
encapsulation x25
x25 address 33333
x25 htc 28
!
interface Serial1.1 point-to-point
no ip address
novell network 2
x25 map novell 2.0000.0c03.a4ad 11111 BROADCAST
!
interface Serial1.2 point-to-point
no ip address
novell network 3
x25 map novell 3.0000.0c07.5e26 55555 BROADCAST
```

Configuration for Router 1

```
hostname Remotel
!
no ip routing
novell routing 0000.0c03.a4ad
!
interface Ethernet0
no ip address
novell network 1
!
interface Serial0
no ip address
encapsulation x25
novell network 2
x25 address 11111
x25 htc 28
x25 map novell 2.0000.0c17.d726 33333 BROADCAST
```


Configuration for Router 2

```
hostname Remote2
!
no ip routing
novell routing 0000.0c07.5e26
!
interface Ethernet0
no ip address
novell network 4
media-type 10BaseT
!
interface Serial0
no ip address
shutdown
!
interface Serial1
no ip address
encapsulation x25
novell network 3
x25 address 55555
x25 htc 28
x25 map novell 3.0000.0c17.d726 33333 BROADCAST
```

IPX over DDR Example

In the configuration shown in Figure 19-2, an IPX client is separated from its server by a DDR telephone line.

Figure 19-2 IPX over DDR Configuration



Routing and service information is sent every minute. The output RIP and SAP filters defined in this example filter these updates, preventing them from being sent between Routers A and B. If you were to forward these packets, the two routers would each have to telephone the other once a minute. On a serial link that charges based on the number of packets transmitted, this is generally not desirable. This might not be an issue on a dedicated serial line.

Once the server and client have established contact, the server will send keepalive (watchdog) packets regularly. The purpose of these packets is to ensure that the connection between the server and the client is still functional; these packets contain no other information. Servers send watchdog packets approximately every 5 minutes. If you were to allow Router B to forward the server's keepalive packets to Router A and the client, Router B would have to telephone Router A every 5 minutes just to send these packets. Again, on a serial link that charges based on the number of packets transmitted, this is generally not desirable. Instead of having Router B telephone Router A only to send keepalive packets, you can enable watchdog spoofing on Router B. This way, when the server connected to this router sends keepalive packets, Router B will respond on behalf of the remote client (the client connected to Router A).

Configuration for Router A

```

access-list 1000 permit -1 7
access-list 1000 deny -1
!
!configure the router to which the client is connected
ipx routing 0000.0c00.59e8
ipx sap 4 SERVER-B BBB01.0000.0000.0001.4212
!
interface ethernet 0
ipx network aaa
!
interface serial 0
no keepalive
dialer in-band
dialer string 8986
ipx network ccc
pulse-time 1
dialer-group 1
ipx output-sap-filter 1000
!
ipx route bbb ccc.0000.0c01.d877

```



```

ipx route bbb01 ccc.0000.0c01.d877
!
access-list 800 permit ffffffff bbb01.0000.0000.0001
access-list 800 deny -1
access-list 1000 permit ffffffff bbb01.0000.0000.0001
access-list 1000 deny -1
dialer-list 1 list 800

```

Configuration for Router B

```

!configure the router to which the server is attached
ipx routing 0000.0x01.d877
!
interface ethernet 0
ipx network bbb
!
interface serial 1
no ip address
bandwidth 56
no keepalive
ipx output-sap-filter 1000
dialer in-band
ipx network bbb
pulse-time 1
no ipx route-cache
no ipx-route-cache cbus
!enable watchdog spoofing on the server's router
ipx watchdog-spoof
!
ipx route aaa ccc.0000.0c00.59e8
access-list 1000 permit 4 bbb01.0000.0001
access-list 1000 deny -1

```

IPX Network Access Example

Using access lists to manage traffic routing can be a powerful tool in overall network control. However, it requires a certain amount of planning and the appropriate application of several related commands. Figure 19-3 illustrates a network featuring two routers on two network segments.

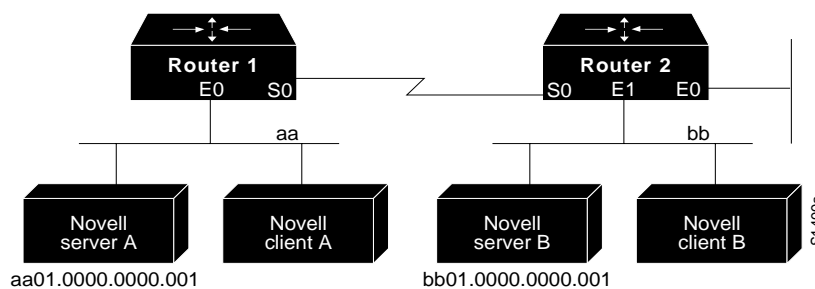


Figure 19-3 Novell IPX Servers Requiring Access Control

Suppose you want to prevent clients and servers on Network aa from using the services on Network bb, but you want to allow the clients and servers on Network bb to use the services on Network aa. To do this, you would need an access list on Ethernet interface 1 on Router 2 that blocks all packets coming from Network aa and destined for Network bb. You would not need any access list on Ethernet interface 0 on Router 1.

You would configure serial interface 0 on Router 2 with the following commands:

```

ipx routing
access-list 800 deny aa bb01
access-list 800 permit -1 -1
interface serial 0
ipx network bb
ipx access-group 800
    
```

SAP Input Filter Example

SAP input filters allow a router to determine whether or not to accept information about a service. Router C1, illustrated in Figure 19-4, will not accept and, consequently not advertise, any information about Novell server F. However, Router C1 will accept information about all other servers on the network 3c. Router C2 receives information about servers D and B.

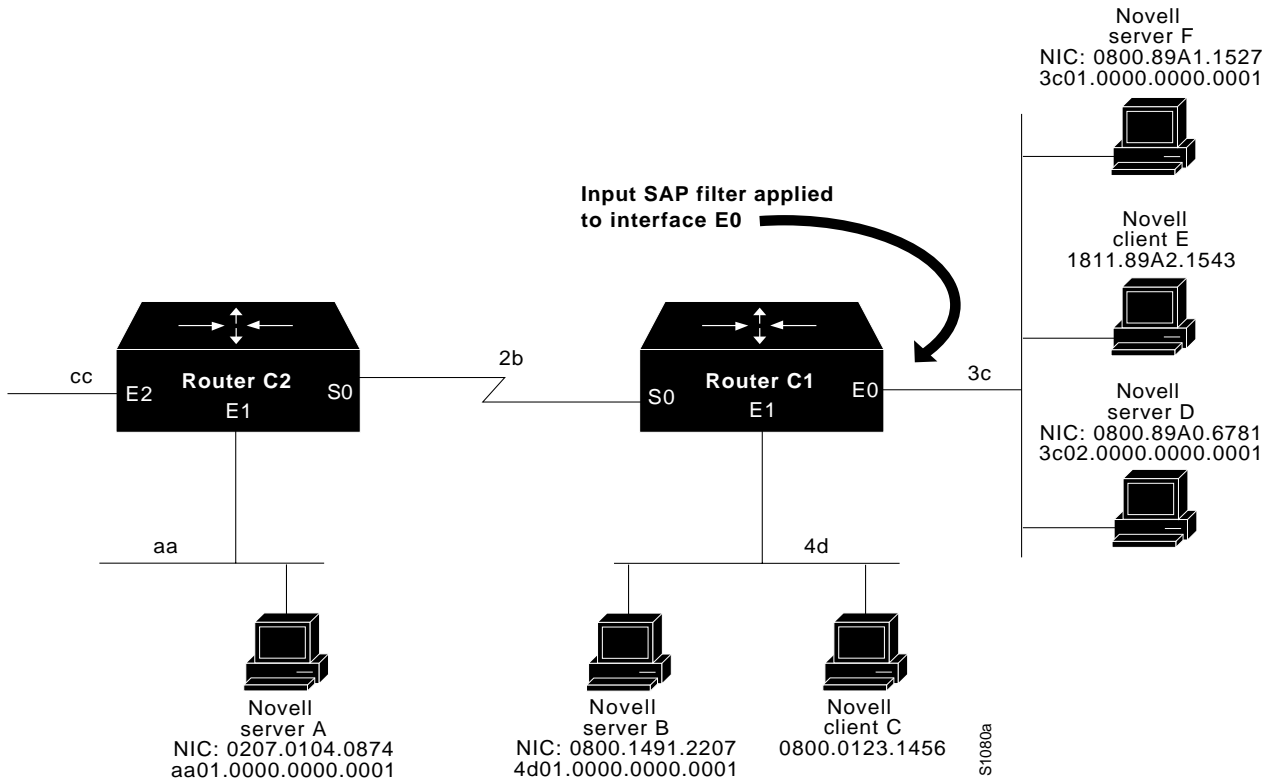


Figure 19-4SAP Input Filter

The following example configures Router C1. The first line denies server F, and the second line accepts all other servers.

```

access-list 1000 deny 3c01.0000.0000.0001
access-list 1000 permit -1
interface ethernet 0
ipx network 3c
ipx input-sap-filter 1000
interface ethernet 1
ipx network 4d
interface serial 0
ipx network 2b
    
```

Note NetWare Versions 3.11 and later use an internal network and node number as their address for access list commands (the first configuration command in this example).

SAP Output Filter Example

SAP output filters are applied prior to the router sending information out a specific interface. In the example that follows, Router C1 (illustrated in Figure 19-5) is prevented from advertising information about Novell server A out interface Ethernet 1, but can advertise server A on network 3c.

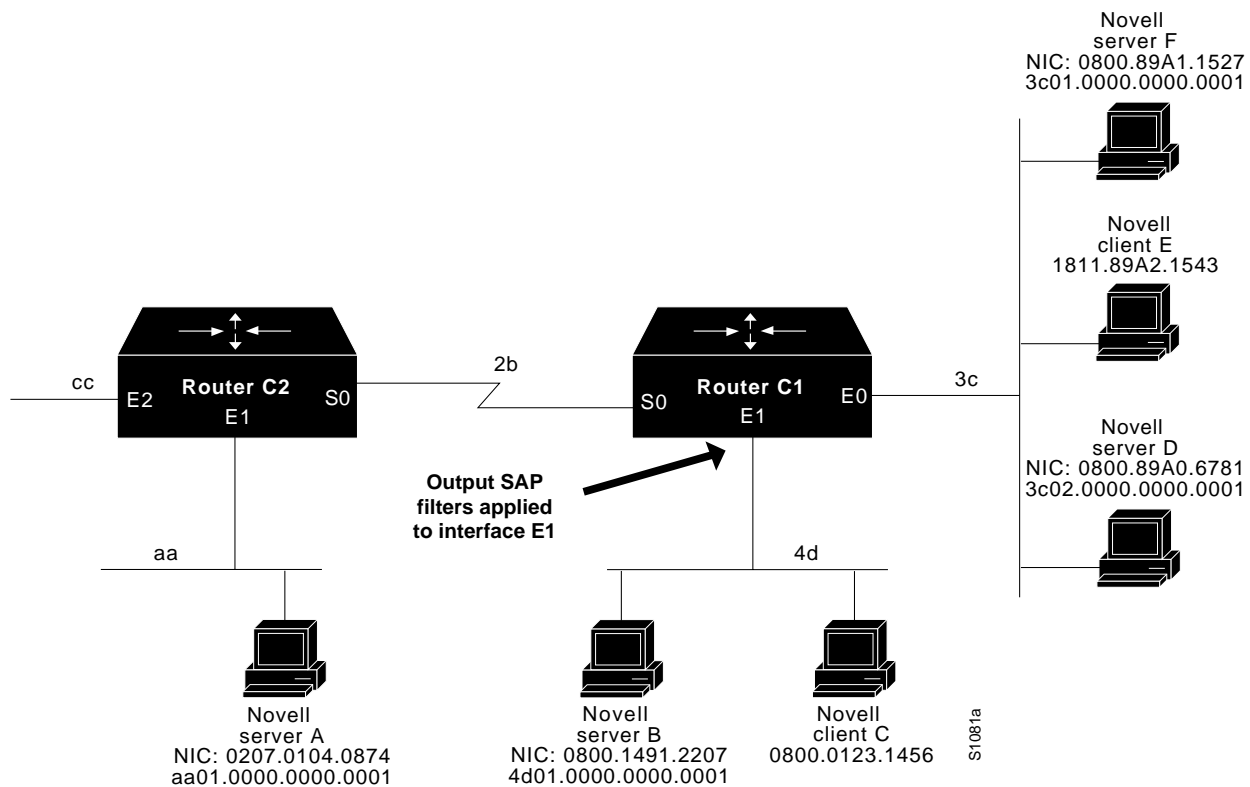


Figure 19-5 SAP Output Filter

The following example refers to Router C1. The first line denies server A. All other servers are permitted.

```
access-list 1000 deny aa01.0000.0000.0001
access-list 1000 permit -1
interface ethernet 0
novell net 3c
interface ethernet 1
ipx network 4d
ipx output-sap-filter 1000
interface serial 0
ipx network 2b
```

IPX NetBIOS Filter Examples

The following is an example of using a NetBIOS host name to filter IPX NetBIOS frames. The example denies all outgoing IPX NetBIOS frames with a NetBIOS host name of Boojum on Ethernet interface 0:

```
netbios access-list host token deny Boojum
netbios access-list host token permit *
!
ipx routing 0000.0c17.d45d
!
interface ethernet 0
ipx network 155 encapsulation ARPA
ipx output-rip-delay 60
ipx triggered-rip-delay 30
ipx output-sap-delay 60
ipx triggered-sap-delay 30
ipx type-20-propagation
ipx netbios output-access-filter host token
no mop enabled
!
interface ethernet 1
no ip address
ipx network 105
!
interface fddi 0
no ip address
no keepalive
ipx network 305 encapsulation SAP
!
interface serial 0
no ip address
shutdown
!
interface serial 1
no ip address
no keepalive
ipx network 600
ipx output-rip-delay 100
ipx triggered-rip-delay 60
ipx output-sap-delay 100
ipx triggered-sap-delay 60
ipx type-20-propagation
```

The following is an example of using a byte pattern to filter IPX NetBIOS frames. This example permits IPX NetBIOS frames from IPX network numbers that end in 05. This means that all IPX NetBIOS frames from Ethernet interface 1 (network 105) and FDDI interface 0 (network 305) will be forwarded by serial interface 0, but this interface will filter out and not forward all frames from Ethernet interface 0 (network 155).

```

netbios access-list bytes finigan permit 2 **05
!
ipx routing 0000.0c17.d45d
!
ipx default-output-rip-delay 1000
ipx default-triggered-rip-delay 100
ipx default-output-sap-delay 1000
ipx default-triggered-sap-delay 100
!
interface ethernet 0
 ipx network 155 encapsulation ARPA
 ipx output-rip-delay 55
 ipx triggered-rip-delay 55
 ipx output-sap-delay 55
 ipx triggered-sap-delay 55
 ipx type-20-propagation
 media-type 10BaseT
!
interface ethernet 1
 no ip address
 ipx network 105
 ipx output-rip-delay 55
 ipx triggered-rip-delay 55
 ipx output-sap-delay 55
 ipx triggered-sap-delay 55
 media-type 10BaseT
!
interface fddi 0
 no ip address
 no keepalive
 ipx network 305 encapsulation SAP
 ipx output-sap-delay 55
 ipx triggered-sap-delay 55
!
interface serial 0
 no ip address
 shutdown
!
interface serial 1
 no ip address
 no keepalive
 ipx network 600
 ipx type-20-propagation
 ipx netbios input-access-filter bytes finigan

```

Helper Facilities to Control Broadcasts Examples

The following examples illustrate how to control broadcast messages on IPX networks. Note that in the following examples, packet type 2 is used. This type has been chosen arbitrarily; the actual type to use depends on the specific application.

Forwarding to an Address Example

All broadcast packets are normally blocked by the router. However, type 20 propagation packets may be forwarded, subject to certain loop-prevention checks. Other broadcasts may be directed to a set of networks or a specific host (node) on a segment. The following examples illustrate these options.

Figure 19-6 shows a router (C1) connected to several Ethernet interfaces. In this environment, all IPX clients are attached to segment aa, while all servers are attached to segments bb and dd. In controlling broadcasts, the following conditions are to be applied:

- Only type 2 and type 20 broadcasts are to be forwarded.
- The IPX clients on network aa are allowed to broadcast via type 2 to any server on networks bb and dd.
- The IPX clients are allowed to broadcast via type 20 to any server on network dd.

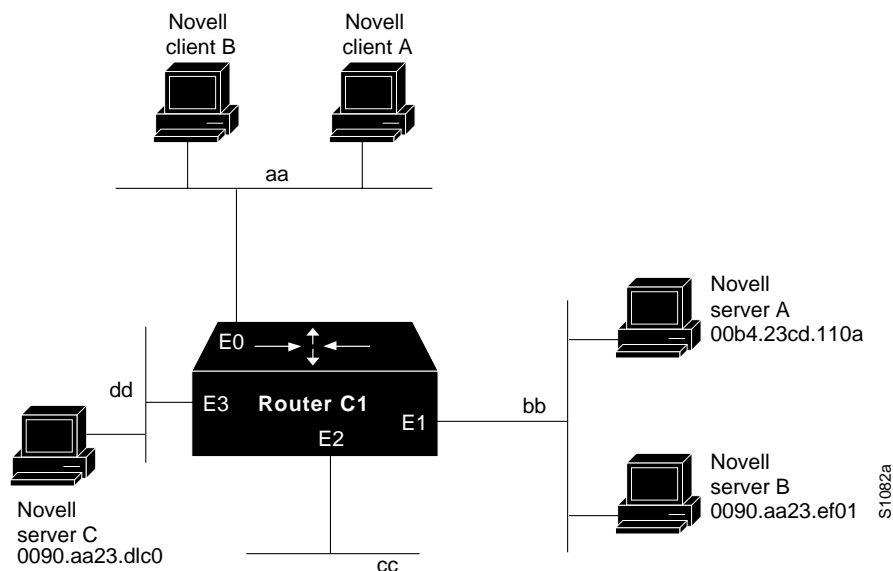


Figure 19-6 IPX Clients Requiring Server Access through a Router

The following example configures the router shown in Figure 19-6. The first line permits broadcast traffic of type 2 from network aa. The interface and network commands configure each specific interface. The **ipx helper-address** commands permit broadcast forwarding from network aa to bb and from network aa to dd. The helper list allows type 2 broadcasts to be forwarded. The **ipx type-20-propagation** command acts as a specific permission to allow type 20 broadcasts to be forwarded between networks aa and dd is also required.

```

access-list 900 permit 2 aa
interface ethernet 0
ipx network aa
ipx type-20-propagation
ipx helper-address bb.ffff.ffff.ffff
ipx helper-address dd.ffff.ffff.ffff
ipx helper-list 900
interface ethernet 1
ipx network bb
interface ethernet 3
ipx network dd
ipx type-20-propagation
    
```

This configuration means that any network that is downstream from network aa (for example, some arbitrary network aa1) will not be able to broadcast (type 2) to network bb through Router C1 unless the routers partitioning networks aa and aa1 are configured to forward these broadcasts with a series of configuration entries analogous to the example provided for Figure 19-5. These entries must be

applied to the input interface and be set to forward broadcasts between directly connected networks. In this way, such traffic can be passed along in a directed manner from network to network. A similar situation exists for type 20 packets.

The following example rewrites the **ipx helper-address** interface configuration command line to direct broadcasts to server A:

```
ipx helper-address bb.00b4.23cd.110a
! Permits node-specific broadcast forwarding to
! Server A at address 00b4.23cd.110a on network bb
```

Forwarding to All Networks Example

In some networks, it might be necessary to allow client nodes to broadcast to servers on multiple networks. If you configure your router to forward broadcasts to all attached networks, you are flooding the interfaces. In the environment illustrated in Figure 19-7, client nodes on network 2b1 must obtain services from IPX servers on networks 3c2, 4a1, and 5bb through Router C1. To support this requirement, use the flooding address (-1.ffff.ffff.ffff) in your **ipx helper-address** interface configuration command specifications.




Figure 19-7 Type 2 Broadcast Flooding

In the following example, the first line permits traffic of type 2 from network 2b1. Then the first interface is configured with a network number. The all-nets helper address is defined and the helper list limits forwarding to type 2 traffic. Type 2 broadcasts from network 2b1 are forwarded to all directly connected networks. All other broadcasts, including type 20, are blocked. To permit broadcasts, delete the **ipx helper-list** entry. To allow type 20 broadcast, enable the **ipx type-20-propagation** interface configuration command on all interfaces.

```
access-list 901 permit 2 2b1
interface ethernet 0
ipx network 2b1
ipx helper-address -1.ffff.ffff.ffff
ipx helper-list 901
interface ethernet 1
ipx network 3c2
interface ethernet 2
ipx network 4a1
interface ethernet 3
ipx network 5bb
```

All-Nets Flooded Broadcast Example

The following example configures all-nets flooding on an interface. As a result of this configuration, Ethernet interface 0 will forward all broadcast messages (except type 20) to all the networks it knows how to reach. This flooding of broadcast messages might overwhelm these networks with so much broadcast traffic that no other traffic may be able to pass on them.

```
interface ethernet 0
ipx network 23
ipx helper-address -1.FFFF.FFFF.FFFF
```

IPX Accounting Example

The following example configures two Ethernet network segments that are connected via a serial link. (See Figure 19-8.) On Router A, IPX accounting is enabled on both the input and output interfaces (that is, on Ethernet interface 0 and serial interface 0). This means that statistics are gathered for traffic traveling in both directions (that is, out to the Ethernet network and out the serial link). However, on Router B, IPX accounting is enabled only on the serial interface and not on the Ethernet interface. This means that statistics are gathered only for traffic that passes out the router on the serial link.




Figure 19-8 IPX Accounting Example

Configuration for Router A

```

ipx routing
interface ethernet 0
no ip address
ipx network C003
ipx accounting
interface serial 0
no ip address
ipx network 200
ipx accounting

```

Configuration for Router B

```

ipx routing
interface ethernet 1
no ip address
no keepalive
ipx network C001
no mop enabled
interface serial 1
no ip address
ipx network 200
ipx accounting

```

Enabling IPX Enhanced IGRP Example

The following example configures two interfaces for Enhanced IGRP routing in autonomous system 1:

```

ipx routing
!
interface ethernet 0
ipx network 10
!
interface serial 0
ipx network 20
!
ipx router eigrp 1
network 10
network 20

```

Enhanced IGRP SAP Update Examples

If an Ethernet interface has neighbors that are all configured for Enhanced IGRP, you might want to reduce the bandwidth used by SAP packets by sending SAP updates incrementally. To do this, you would configure the interface as follows:

```

ipx routing
!
interface ethernet 0
ipx network 10
ipx sap-incremental eigrp 1
!
interface serial 0
ipx network 20
!
ipx router eigrp 1
network 10
network 20

```

If you want to send periodic SAP updates on a serial line that is configured for Enhanced IGRP and that has an Enhanced IGRP peer on the other side, use the following commands:

```
ipx routing
!
interface ethernet 0
ipx network 10
!
interface serial 0
ipx network 20
no ipx sap-incremental eigrp 1
!
ipx router eigrp 1
network 10
network 20
```