CHAPTER 7

IPX Enhanced IGRP Commands

Use the commands in this chapter to configure and monitor IPX enhanced IGRP on Novell IPX networks. For IPX enhanced IGRP configuration information and examples, refer to the "Configuring IPX Enhanced IGRP" chapter in this publication. For a description of other IPX configuration commands, refer to the *Router Products Configuration Guide* and the *Router Products Command Reference* publications. For historical background and a technical overview of IPX, see the *Internetworking Technology Overview* publication.

distribute-list in

To filter networks received in updates, use the **distribute-list in** router configuration command. To change or cancel the filter, use the **no** form of this command.

distribute-list *access-list-number* **in** [*interface-name*]
no distribute-list *access-list-number* **in** [*interface-name*]

Syntax Description

access-list-number	Standard IPX access list number in the range 800 to 899. The list explicitly specifies which networks are to be received and which are to be suppressed.
in	Applies the access list to incoming routing updates.
interface-name	(Optional) Interface on which the access list should be applied to incoming updates. If no interface is specified, the access list is applied to all incoming updates.

Default

Disabled

Command Mode

Router configuration

Example

The following example causes only two networks—network 2 and network 3—to be accepted by an IP enhanced IGRP routing process:

```
access-list 800 permit 2
access-list 800 permit 3
access-list 800 deny -1
!
ipx router eigrp 100
network 3
distribute-list 800 in
```

Related Commands

A dagger (†) indicates that the command is documented in the *Router Products Command Reference* publication.

access-list †
distribute-list out
redistribute

distribute-list out

To suppress networks from being advertised in updates, use the **distribute-list out** router configuration command. To cancel this function, use the **no** form of this command.

distribute-list *access-list-number* **out** [*interface-name* | *routing-process*]
no distribute-list *access-list-number* **out** [*interface-name* | *routing-process*]

Syntax Description

access-list-number	Standard IPX access list number in the range 800 to 899. The list explicitly specifies which networks are to be sent and which are to be suppressed in routing updates.
out	Applies the access list to outgoing routing updates.
interface-name	(Optional) Interface on which the access list should be applied to outgoing updates. If no interface is specified, the access list is applied to all outgoing updates.
routing-process	(Optional) Name of a particular routing process (**rip** or **eigrp** *autonomous-system-number*).

Default

Disabled

Command Mode

Router configuration

Usage Guidelines

When redistributing networks, a routing process name can be specified as an optional trailing argument to the **distribute-list out** command. This causes the access list to be applied to only those routes derived from the specified routing process. After the process-specific access list is applied, any access list specified by a **distribute-list out** command without a process name argument is applied. Addresses not specified in the **distribute-list out** command are not advertised in outgoing routing updates.

Example

The following example causes only one network—network 3—to be advertised by an IPX enhanced IGRP routing process:

```
access-list 800 permit 3
access-list 800 deny -1
!
ipx router eigrp 100
network 3
distribute-list 800 out
```

Related Commands

A dagger (†) indicates that the command is documented in the *Router Products Command Reference* publication.

access-list †
distribute-list in
redistribute

ipx backup-server-query-interval

To change the time between successive queries of each IPX enhanced IGRP neighbor's backup server table, use the **ipx backup-server-query-interval** global configuration command. To restore the default time, use the **no** form of this command.

ipx backup-server-query-interval *interval*
no ipx backup-server-query-interval

Syntax Description

interval	Minimum time, in seconds, between successive queries of each enhanced IGRP neighbor's backup server table. The default is 15 seconds.

Default

15 seconds

Command Mode

Global configuration

Usage Guidelines

A lower interval may use more CPU resources, but may cause lost server information to be retrieved from other servers' tables sooner.

Example

The following example changes the server query time to 5 seconds:

```
ipx backup-server-query-interval 5
```

ipx hello-interval eigrp

To configure the interval between IPX enhanced IGRP hello packets, use the **ipx hello-interval eigrp** interface configuration command. To restore the default interval, use the **no** form of this command.

ipx hello-interval eigrp *autonomous-system-number seconds*
no ipx hello-interval eigrp *autonomous-system-number seconds*

Syntax Description

autonomous-system-number	IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.
seconds	Interval between hello packets, in seconds. The default interval is 5 seconds, which is one-third of the default hold time.

Default

5 seconds

Command Mode

Interface configuration

Usage Guidelines

If the current value for the hold time is less than two times the interval between hello packets, the hold time will be reset.

Example

The following example changes the hello interval to 10 seconds:

```
interface ethernet 0
ipx network 10
ipx hello-interval eigrp 4 10
```

Related Command

ipx hold-time eigrp

ipx hold-time eigrp

To specify the length of time a neighbor should consider IPX enhanced IGRP hello packets valid, use the **ipx hold-time eigrp** interface configuration command. To restore the default time, use the **no** form of this command.

ipx hold-time eigrp *autonomous-system-number seconds*
no ipx hold-time eigrp *autonomous-system-number seconds*

Syntax Description

autonomous-system-number	IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.
seconds	Hold time, in seconds. 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 15 seconds, which is three times the hello interval.

Default

15 seconds

Command Mode

Interface configuration

Usage Guidelines

If the current value for the hold time is less than two times the interval between hello packets, the hold time will be reset to three times the hello interval.

If a router does not receive a hello packet within the specified hold time, routes through the router are considered available.

Increasing the hold time delays route convergence across the network.

Example

The following example changes the hold time to 45 seconds:

```
interface ethernet 0
ipx network 10
ipx hold-time eigrp 4 45
```

Related Command

ipx hello-interval eigrp

ipx router

To specify the routing protocol to use, use the **ipx router** global configuration command.

ipx router {**eigrp** *autonomous-system-number* | **rip**}

Syntax Description

eigrp *autonomous-system-number*	Specifies the IPX enhanced IGRP routing protocol. The argument *autonomous-system-number* is the IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.
rip	Specifies the RIP routing protocol. It is on by default.

Default

RIP

Command Mode

Global configuration

Usage Guidelines

You can configure multiple IPX enhanced IGRP processes on a router. To do so, assign each a different autonomous system number.

Example

The following example enables IPX enhanced IGRP on the router:

```
ipx router eigrp 4
```

Related Commands

network
redistribute

ipx sap-incremental eigrp

To send SAP updates only when a change occurs in the SAP table, use the **ipx sap-incremental eigrp** interface configuration command. To send periodic SAP updates, use the **no** form of this command.

ipx sap-incremental eigrp *autonomous-system-number* [**rsup-only**]
no ipx sap-incremental eigrp *autonomous-system-number* [**rsup-only**]

Syntax Description

autonomous-system-number	IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.
rsup-only	(Optional) Indicates that the system uses enhanced IGRP on this interface to carry reliable SAP update information only.

Default

Enabled on serial interfaces
Disabled on LAN media (Ethernet, Token Ring, FDDI)

Command Mode

Interface configuration

Usage Guidelines

With this functionality enabled, if an IPX enhanced IGRP peer is found on the interface, SAP updates will be sent only when a change occurs in the SAP table. Periodic SAP updates are not sent. When no IPX enhanced IGRP peer is present on the interface, periodic SAPs are always sent regardless of how this command is set.

If you configure the local router to send incremental SAP updates on an Ethernet, and if the local router has at least one IPX enhanced IGRP neighbor and any servers, clients, or routers that do not have IPX enhanced IGRP configured on the Ethernet interface, these devices will not receive complete SAP information from the local router.

If the incremental sending of SAP updates on an interface is configured and no IPX enhanced IGRP peer is found, SAP updates will be sent periodically until a peer is found. Then, updates will be sent only when changes occur in the SAP table.

To reduce SAP traffic by sending partial SAP updates, specify the **rsup-only** keyword. SAP udpates are then sent only when changes occur, and only changes are sent. This feature works with existing IPX RIP networks and IPX enhanced IGRP networks.

Example

The following example sends SAP updates on Ethernet interface 0 only when there is a change in the SAP table:

```
interface ethernet 0
ipx sap-incremental eigrp 200
```

ipx split-horizon eigrp

To configure split horizon, use the **ipx split-horizon eigrp** interface configuration command. To disable split horizon, use the **no** form of this command.

ipx split-horizon eigrp *autonomous-system-number*
no ipx split-horizon eigrp *autonomous-system-number*

Syntax Description

autonomous-system-number	IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.

Default

Enabled

Command Mode

Interface configuration

Usage Guidelines

When split horizon is enabled, IPX enhanced IGRP update and query packets are not sent for destinations that have next hops on this interface. This reduces the number of enhanced IGRP packets on the network.

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.

Example

The following example disables split horizon on serial interface 0:

```
interface serial 0
no ipx split-horizon eigrp 200
```

network

To enable IPX enhanced IGRP on the router, use the **network** IPX-router configuration command. To disable IPX enhanced IGRP on the router, use the **no** form of this command.

network {*network-number* | **all**}
no network {*network-number* | **all**}

Syntax Description

network-number	IPX network number
all	Enables the routing protocol for all IPX networks configured on the router

Default

Disabled

Command Mode

IPX-router configuration

Usage Guidelines

Use the **network** command to enable the routing protocol specified in the **ipx router** command on each network.

Example

The following commands disable RIP on network 10 and enable IPX enhanced IGRP on networks 10 and 20:

```
ipx router rip
no network 10
ipx router eigrp 12
network 10
network 20
```

Related Command

ipx router

redistribute

To redistribute from one routing domain into another, and vice versa, use the **redistribute** IPX-router configuration command. To disable this feature, use the **no** form of this command.

redistribute {**rip** | **eigrp** *autonomous-system-number* | **connected** | **static**}
no redistribute {**rip** | **eigrp** *autonomous-system-number* | **connected** | **static**}

Syntax Description

rip	Specifies the RIP protocol.
eigrp *autonomous-system-number*	Specifies the enhanced IGRP protocol and the IPX enhanced IGRP autonomous system number. It can be a decimal integer from 1 to 65535.
connected	Specifies connected routes.
static	Specifies static routes.

Default

Redistribution is enabled by default between all routing domains except between separate enhanced IGRP processes.

Command Mode

IPX-router configuration

Usage Guidelines

Redistribution provides for routing information generated by one protocol to be advertised in another.

The only connected routes affected by this redistribute command are the routes not specified by the **network** command.

Examples

In the following example, RIP routing information is not redistributed:

```
ipx router eigrp 222
no redistribute rip
```

In the following example, enhanced IGRP routes from autonomous system 100 are redistributed into enhanced IGRP autonomous sytem 300:

```
router eigrp 300
redistribute eigrp 100
```

show ipx eigrp neighbors

To display the neighbors discovered by enhanced IGRP, use the **show ipx eigrp neighbors** EXEC command.

show ipx eigrp neighbors [**servers**] [*autonomous-system-number* | *type unit*]

Syntax Description

servers	(Optional) Displays the server list advertised by each neighbor. This is displayed only if the **ipx sap incremental** command is enabled on the interface on which the neighbor resides.
autonomous-system-number	(Optional) Autonomous system number. It can be a decimal integer from 1 to 65535.
type unit	(Optional) Interface type and number about which to display neighbor information.

Command Mode

EXEC

Sample Display

The following is sample output from the **show ipx eigrp neighbors** command:

```
Router# show ipx eigrp neighbors
IPX EIGRP Neighbors for process 200
H  Address                 Interface     Hold   Uptime   Q    Seq    SRTT  RTO
                                         (secs) (h:m:s)  Cnt  Num    (ms)  (ms)
6  90.0000.0c02.096e       Tunnel44444   13     0:30:57  0    21     9     20
5  80.0000.0c02.34f2       Fddi0         12     0:31:17  0    62     14    28
4  83.5500.2000.a83c       TokenRing2    13     0:32:36  0    626    16    32
3  98.0000.3040.a6b0       TokenRing1    12     0:32:37  0    43     9     20
2  80.0000.0c08.cbf9       Fddi0         12     0:32:37  0    624    19    38
1  85.aa00.0400.153c       Ethernet2     12     0:32:37  0    627    15    30
0  82.0000.0c03.4d4b       Hssi0         12     0:32:38  0    629    12    24
```

Table 7-1 explains the fields in the output.

Table 7-1 Show IPX EIGRP Neighbors Field Descriptions

Field	Description
process 200	Autonomous system number specified in the **ipx router** configuration command.
H	Handle. An arbitrary and unique number inside this router that identifies the neighbor.

Field	Description
Address	IPX address of the enhanced IGRP peer.
Interface	Interface on which the router is receiving hello packets from the peer.
Holdtime	Length of time, in seconds, that the router will wait to hear from the peer before declaring it down. If the peer is using the default hold time, this number will be less than 15. If the peer configures a nondefault hold time, it will be reflected here.
Uptime	Elapsed time, in hours, minutes, and seconds, since the local router first heard from this neighbor.
Q Count	Number of IPX enhanced IGRP packets (Update, Query, and Reply) that the router is waiting to send.
Seq Num	Sequence number of the last Update, Query, or Reply packet that was received from this neighbor.
SRTT	Smooth round-trip time. This is the number of milliseconds it takes for an IPX enhanced IGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet.
RTO	Retransmission timeout, in milliseconds. This is the amount of time the router waits before retransmitting a packet from the retransmission queue to a neighbor.

show ipx eigrp topology

To display the IPX enhanced IGRP topology table, use the **show ipx eigrp topology** EXEC command.

show ipx eigrp topology [*network-number*]

Syntax Description

network-number	(Optional) IPX network number whose topology table entry to display

Command Mode

EXEC

Sample Display

The following is sample output from the **show ipx eigrp topology** command:

```
Router# show ipx eigrp topology
IPX EIGRP Topology Table for process 109
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - Reply status
P 42, 1 successors, FD is 0
     via 160.0000.0c00.8ea9 (345088/319488), Ethernet0
P 160, 1 successor via Connected, Ethernet
     via 160.0000.0c00.8ea9 (307200/281600), Ethernet0
P 165, 1 successors, FD is 307200
     via Redistributed (287744/0)
     via 160.0000.0c00.8ea9 (313344/287744), Ethernet0
P 164, 1 successors, flags: U, FD is 200
     via 160.0000.0c00.8ea9 (307200/281600), Ethernet1
     via 160.0000.0c01.2b71 (332800/307200), Ethernet1
P A112, 1 successors, FD is 0
     via Connected, Ethernet2
     via 160.0000.0c00.8ea9 (332800/307200), Ethernet0
P AAABBB, 1 successors, FD is 10003
     via Redistributed (287744/0),
     via 160.0000.0c00.8ea9 (313344/287744), Ethernet0
A A112, 0 successors, 1 replies, state: 0, FD is 0
     via 160.0000.0c01.2b71 (307200/281600), Ethernet1
     via 160.0000.0c00.8ea9 (332800/307200) , r, Ethernet1
```

Table 7-2 explains the fields in the output.

Table 7-2 Show IPX EIGRP Topology Field Descriptions

Field	Description
Codes	State of this topology table entry. Passive and Active refer to the enhanced IGRP state with respect to this destination; Update, Query, and Reply refer to the type of packet that is being sent.
P – Passive	No enhanced IGRP computations are being performed for this destination.
A – Active	Enhanced IGRP computations are being performed for this destination.
U – Update	Indicates that an update packet was sent to this destination.
Q – Query	Indicates that a query packet was sent to this destination.
R – Reply	Indicates that a reply packet was sent to this destination.
r – Reply status	Flag that is set after the router has sent a query and is waiting for a reply.
42, 160, and so on	Destination IPX network number.
successors	Number of successors. This number corresponds to the number of next hops in the IPX routing table.
FD	Feasible distance. This value is used in the feasibility condition check. If the neighbor's reported distance (the metric after the slash) is less than the feasible distance, the feasibility condition is met and that path is a feasible successor. Once the router determines it has a feasible successor, it does not have to send a query for that destination.
replies	Number of replies that are still outstanding (have not been received) with respect to this destination. This information appears only when the destination is in Active state.
state	Exact enhanced IGRP state that this destination is in. It can be the number 0, 1, 2, or 3. This information appears only when the destination is Active.
via	IPX address of the peer who told the router about this destination. The first *n* of these entries, where *n* is the number of successors, are the current successors. The remaining entries on the list are feasible successors.
(345088/319488)	The first number is the enhanced IGRP metric that represents the cost to the destination. The second number is the enhanced IGRP metric that this peer advertised.
Ethernet0	Interface from which this information was learned.

The following is sample output from the **show ipx eigrp topology** command when you specify an IPX network number:

```
Router# show ipx eigrp topology 160
IPX-EIGRP topology entry for 160
State is Passive, Query origin flag is 1, 1 Successor(s)
Routing Descriptor Blocks:
  Next hop is Connected (Ethernet0), from 0.0000.0000.0000
  Composite metric is (0/0), Send flag is 0x0, Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 1000000 nanoseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 0
Next hop is 164.0000.0c00.8ea9 (Ethernet1), from 164.0000.0c00.8ea9
  Composite metric is (307200/281600), Send flag is 0x0, Route is External
  This is an ignored route
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2000000 nanoseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
  External data:
    Originating router is 0000.0c00.8ea9
    External protocol is RIP, metric is 1, delay 2
    Administrator tag is 0 (0x00000000)
    Flag is 0x00000000
```

Table 7-3 explains the fields in the output.

Table 7-3 Show IPX EIGRP Topology Field Descriptions for a Specified Network

Field	Description
160	IPX network number of the destination.
State is ...	State of this entry. It can be either Passive or Active. Passive means that no enhanced IGRP computations are being performed for this destination, and Active means that they are being performed.
Query origin flag	Exact enhanced IGRP state that this destination is in. It can be the number 0, 1, 2, or 3. This information appears only when the destination is Active.
Successors	Number of successors. This number corresponds to the number of next hops in the IPX routing table.
Next hop is ...	Indicates how this destination was learned. It can be one of the following: • Connected—The destination is on a network directly connected to this router. • Redistributed—The destination was learned via RIP or another enhanced IGRP process. • IPX host address—The destination was learned from that peer via this enhanced IGRP process.
Ethernet0	Interface from which this information was learned.
from	Peer from whom the information was learned. For connected and redistributed routers, this is 0.0000.0000.0000. For information learned via enhanced IGRP, this is the peer's address. Currently, for information learned via enhanced IGRP, the peer's IPX address always matches the address in the "Next hop is" field.

Field	Description
Composite metric is	Enhanced IGRP composite metric. The first number is this router's metric to the destination, and the second is the peer's metric to the destination.
Send flag	Numeric representation of the "flags" field described in Table 7-1. It is 0 when nothing is being sent, 1 when an Update is being sent, 3 when a Query is being sent, and 4 when a Reply is being sent. Currently, 2 is not used.
Route is ...	Type of router. It can be either internal or external. Internal routes are those that originated in an enhanced IGRP autonomous system, and external are routes that did not. Routes learned via RIP are always external.
This is an ignored route	Indicates that this path is being ignored because of filtering.
Vector metric:	This section describes the components of the enhanced IGRP metric.
Minimum bandwidth	Minimum bandwidth of the network used to reach the next hop.
Total delay	Delay time to reach the next hop.
Reliability	Reliability value used to reach the next hop.
Load	Load value used to reach the next hop.
Minimum MTU	Minimum MTU size of the network used to reach the next hop.
Hop count	Number of hops to the next hop.
External data	This section describes the original protocol from which this route was redistributed. It appears only for external routes.
Originating router	Network address of the router that first distributed this route into enhanced IGRP.
External protocol..metric..delay	External protocol from which this route was learned. The metric will match the external hop count displayed by the **show ipx route** command for this destination. The delay is the external delay.
Administrator tag	Not currently used.
Flag	Not currently used.

Related Command

show ipx route

show ipx route

To display the contents of the IPX routing table, use the **show ipx route** EXEC command.

show ipx route [*network-number*]

Syntax Description

network-number	(Optional) Number of the network whose routing table entry you want to display. This is an eight-digit hexadecimal number that uniquely identifies a network cable segment. It can be a number in the range 1 to FFFFFFFE. A network number of 0 matches the local network. A network number of –1 matches all networks. You do not need to specify leading zeros in the network number. For example, for the network number 000000AA, you can enter just AA.

Command Mode

EXEC

Sample Display

The following is sample output from the **show ipx route** command:

```
Router> show ipx route
Codes: R - RIP, E - EIGRP, C - connected, S - static, 5 Total IPX routes

No parallel paths allowed        Novell routing algorithm variant in use

E Net 42 [345088/2] via 160.0000.0c00.8ea9, age 2d21 1 uses, Ethernet0
C Net FFFF01 (NOVELL-ETHER), is directly connected, 33 uses, Ethernet3
R Net E04 [120/2] via 165.0000.0c01.d879,  17 sec, 1 uses, Ethernet0
R Net E02 [120/2] via 165.0000.0c01.d879,  17 sec, 2 uses, Ethernet0
R Net E03 [120/2] via 165.0000.0c01.d879,  17 sec, 57 uses, Ethernet0
C Net CC00C1 (HDLC), is directly connected, 0 uses, serial3, (down)
```

Table 7-4 describes the fields shown in the display.

Table 7-4 Show IPX Route Field Descriptions

Field	Description
Codes	Codes defining how the route was learned.
R	Route learned from a RIP update.
E	Route learned via enhanced IGRP.
C	Directly connected network.
S	Statically defined route via the **ipx route** command.
No parallel paths allowed	Maximum number of parallel paths for which the router has been configured with the **ipx maximum-paths** command.
Novell routing algorithm variant in use	Indicates whether the router is using the IPX-compliant routing algorithms (default).
sec	Number of seconds that have elapsed since information about this network was last received.
uses	Fair estimate of the number of times a route gets used. It actually indicates the number of times the route has been selected for use prior to operations such as access list filtering.
Ethernet0	Possible interface through which updates to the remote network will be sent.
(NOVELL-ETHER) (HDLC)	Encapsulation type. This is shown only for directly connected networks.
is directly connected	Indicates that the network is directly connected to the router.
Net E04	Network to which the route goes.
[345088/2]	The first number is the enhanced IGRP metric, and the second number is the external hop count (for routes learned from enhanced IGRP). The external hop count is the hop count that was advertised when this destination first entered the enhanced IGRP autonomous system. If the external hop count is 0, this is an internal enhanced IGRP route. When this destination is announced via RIP by this router, the router will add one to this value and use the result as the RIP hop count. The external hop count is not incremented when the destination is advertised via enhanced IGRP.
[120/2]	Ticks/metric (for routes learned from RIP). Ticks are the number of IBM clock ticks (1/18th second). When **ipx algorithms** is enabled, hop count is used as a tie-breaker between routes if the tick count is the same. The metric is the IPX metric used in making routing decisions.
age	(For routes learned from enhanced IGRP only.) Length of time since this route was first learned. The time is represented in the format *hh*:*mm*:*ss* until 24 hours have elapsed, after which it is represented *ddd*h*h* until one week has elapsed. After one week, the age is represented *wwwdd*.
via *network.node*	Address of a router that is the next hop to the remote network.

show ipx traffic

To display information about the number and type of IPX packets transmitted and received by the router, use the **show ipx traffic** EXEC command.

show ipx traffic

Syntax Description

This command has no arguments or keywords.

Command Mode

EXEC

Sample Display

The following is sample output from the **show ipx traffic** command:

```
Router> show ipx traffic
Rcvd:   644 total, 1705 format errors, 0 checksum errors, 0 bad hop count,
        0 packets pitched, 644 local destination, 0 multicast
Bcast:  589 received, 324 sent
Sent:   380 generated, 0 forwarded
        0 encapsulation failed, 4 no route
SAP:    1 SAP requests, 1 SAP replies
        61 SAP advertisements received, 120 sent
        0 SAP flash updates sent, 0 SAP poison sent
        0 SAP format errors
RIP:    0 RIP format errors
Echo:   Rcvd 55 requests, 0 replies
        Sent 0 requests, 55 replies
        0 unknown, 0 SAPs throttled, freed NDB len 0
Watchdog:
        0 packets received, 0 replies spoofed
Queue lengths: IPX input: 0, SAP 0, RIP 0, GNS 0
               Total length for SAP throttling purposes: 0/(no preset limit)
EIGRP:  Total received 0, sent 0
        Updates received 0, sent 0
        Queries received 0, sent 0
        Replies received 0, sent 0
        SAPs received 0, sent 0
```

Table 7-5 describes the fields that might be shown in the display.

Table 7-5 Show IPX Traffic Field Descriptions

Field	Description
Rcvd:	Description of the packets the router has received.
644 total	Total number of packets the router has received.
1705 format errors	Number of bad packets discarded (for example, packets with a corrupted header).
0 checksum errors	Number of packets containing a checksum error. This should always be 0, because IPX does not use a checksum.
0 bad hop count	Number of packets discarded because their hop count exceeded 16 (that is, the packets timed out).
0 packets pitched	Number of times the router received its own broadcast packet.
644 local destination	Number of packets sent to the local broadcast address or specifically to the router.
0 multicast	Number of packets received that were addressed to multiple destinations.
Bcast:	Description of the broadcast packets the router has received and sent.
589 received	Number of broadcast packets received.
324 sent	Number of broadcast packets sent. It includes broadcast packets the router is either forwarding or has generated.
Sent:	Description of those packets that the router generated and then sent, and also those the router has received and then routed to other destinations.
380 generated	Number of packets the router transmitted that it generated itself.
0 forwarded	Number of packets the router transmitted that it forwarded from other sources.
0 encapsulation failed	Number of packets the router was unable to encapsulate.
4 no route	Number of times the router could not locate a route to the destination in the routing table.
SAP:	Description of the SAP packets the router has sent and received.
1 SAP requests	Number of SAP requests the router has received.
1 SAP replies	Number of SAP replies the router has sent in response to SAP requests.
61 SAP advertisements received	Number of SAP advertisements the router has received from another router.
120 sent	Number of SAP advertisements the router has generated and then sent.
0 SAP flash updates sent	Number of SAP advertisements the router has generated and then sent as a result of a change in its routing table.
0 SAP poison sent	Number of times the router has generated an update indicating that a service is no longer reachable.
0 SAP format errors	Number of SAP advertisements that were incorrectly formatted.
RIP:	Description of the RIP packets the router has sent and received.
0 RIP format errors	Number of RIP packets that were incorrectly formatted.
freed NDB length	Number of Network Descriptor Blocks (NDBs) that have been removed from the network but still need to be removed from the router's routing table.
Watchdog:	Description of the watchdog packets the router has handled.
0 packets received	Number of watchdog packets the router has received from IPX servers on the local network.

Field	Description
0 replies spoofed	Number of times the router has responded to a watchdog packet on behalf of the remote client.
Echo:	Description of the ping replies and requests the router has sent and received.
Rcvd 55 request 0 replies	Number of ping requests and replies received by the router.
Sent 0 requests, 55 replies	Number of ping requests and replies sent by the router.
0 unknown	Number of incomprehensible ping packets received by the router.
0 SAPs throttled	Number of ping packets discarded because they exceeded buffer capacity.
Queue lengths	Description of outgoing packets currently in buffers that are waiting to be processed.
IPX input	Number of incoming packets waiting to be processed.
SAP	Number of outgoing SAP packets waiting to be processed.
RIP	Number of outgoing RIP packets waiting to be processed.
GNS	Number of outgoing GNS packets waiting to be processed.
Total length for SAP throttling purposes	Maximum number of outgoing SAP packets allowed in the buffer. Any packets received beyond this number are discarded.
EIGRP:	Description of the IPX enhanced IGRP packets the router has sent and received.
Updates	Number of IPX enhanced IGRP updates the router has sent and received.
Queries	Number of IPX enhanced IGRP queries the router has sent and received.
Replies	Number of IPX enhanced IGRP replies the router has sent and received.
SAPs	Number of SAP packets the router has sent to and received from IPX enhanced IGRP neighbors.
unknown counter	Number of packets the router was unable to forward, for example, because of a misconfigured helper address or because no route was available.