

# Understanding IPX–EIGRP

Document ID: 10579

## Contents

### Introduction

#### Before You Begin

- Conventions
- Prerequisites
- Components Used
- Background Information

#### EIGRP Components

#### IPX–EIGRP Features

#### IPX–EIGRP Internetworking Terms

#### Understanding Routing and Topology Tables

#### EIGRP Packet Format

- IPX–Specific TLVs
- IPX SAP Packets

#### IPX–EIGRP Configuration Commands

- Global IPX Commands
- Router Subcommands
- Interface Subcommands
- show Commands
- debug Commands
- Output of show Commands

#### Troubleshooting Neighbor Relationships

#### References

#### Related Information

## Introduction

Cisco Interior Gateway Routing Protocol (IGRP) is used in TCP/IP and Open System Interconnection (OSI) Internets. The original IP version was designed and deployed successfully in 1986. IGRP uses distance vector routing technology so that each router does not have to know all of the router/link relationships for the entire network. Each router advertises destinations with a corresponding distance. Each router, hearing the information, adjusts the distance and propagates it to neighboring routers.

## Before You Begin

### Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

### Prerequisites

There are no specific prerequisites for this document.

### Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

## Background Information

The distance information in IGRP is represented as a composite of available bandwidth, delay, load utilization, and link reliability. This allows fine tuning of link characteristics to achieve optimal paths.

EIGRP is Cisco's Enhanced version of IGRP, and has three versions: one for IP, one for Internetwork Packet Exchange (IPX), and one for AppleTalk. They each use the same Distributed Update Algorithm (DUAL). The same distance vector technology found in IGRP is also used in EIGRP, and the underlying distance information remains unchanged. The convergence properties and the operating efficiency of this protocol have improved significantly. This allows for an improved architecture while retaining existing investment in IGRP.

The convergence technology is based on research conducted at SRI International. The DUAL is used to obtain loop–freedom at every instant throughout a route computation. This 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 the recomputation. The convergence time with DUAL rivals that of any other existing routing protocol.

## EIGRP Components

EIGRP 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. This process 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 is determined, the neighboring routers can exchange routing information.
- **The reliable transport protocol** is responsible for the guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports inter–mixed transmission of multicast or unicast packets. Some EIGRP packets must be transmitted reliably; others are not. For efficiency, reliability is provided only when necessary. For example, on a multi–access network that has multicast capabilities such as Ethernet, it is not necessary to send hellos reliably to all neighbors individually. EIGRP will instead send 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; this is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending, which 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. The distance information, known as a metric, is used by DUAL 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 where 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 when unnecessary. When a topology change occurs, DUAL will test for feasible successors. If there are none, DUAL will use any it finds in order to avoid

unnecessary recomputation.

- **The protocol dependent modules** are responsible for network layer protocol specific requirements. For example, the IPX–EIGRP module is responsible for sending and receiving EIGRP packets that are encapsulated in IPX. IPX–EIGRP is responsible for passing EIGRP packets and informing DUAL of the new information received. IPX–EIGRP asks DUAL to make routing decisions, the results of which are stored in the IPX routing table.

## IPX–EIGRP Features

IPX–EIGRP provides the following features:

- **Automatic redistribution** – IPX–Routing Information Protocol (RIP) routes are automatically redistributed into EIGRP, and IPX–EIGRP routes are automatically redistributed into RIP, without any commands being entered by the user. Redistribution can be turned off with the use of the no redistribute router subcommand. Both IPX–RIP and IPX–EIGRP can be turned off completely on the router.
- **Increased network width** – With IPX–RIP, the largest possible width of your network is 15 hops. When IPX–EIGRP is enabled, the largest possible width is 224 hops. Since the EIGRP metric is sufficiently 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 only incrementing the transport control field when an IPX packet has traversed 15 routers and the next hop to the destination was learned via EIGRP. 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 until an EIGRP neighbor is found and thereafter only when there are changes to the SAP table. This works by taking advantage of EIGRP's reliable transport mechanism so an IPX–EIGRP peer must be present for incremental SAPs to be sent. If no peer exists on a particular interface, then periodic SAPs will be sent on that interface until a peer is found. This functionality is usually automatic on serial interfaces and may be configured on LAN media if desired.

## IPX–EIGRP Internetworking Terms

- **Active state** – A topology table entry is considered to be in the active state when a route recomputation occurs.
- **Autonomous system (AS)** – An autonomous system is a collection of networks under a common administration sharing a common routing strategy. An autonomous system may comprise one or many networks. All routers that belong to an autonomous system must be configured with the same autonomous system number.
- **DUAL** – A loop–free routing algorithm used with distance vectors or link states that provide a diffused computation of a routing table. DUAL was developed at SRI International by Dr. J.J. Garcia–Luna–Aceves.
- **External hop count** – The hop count to a destination that is advertised to the router in the protocol being redistributed. For example, if a router receives a RIP update advertising a destination as three hops away, when this RIP information is redistributed into EIGRP, the three hops will be stored as the external hop count and this information will be passed throughout the EIGRP autonomous system.
- **External routes** – A router considers an EIGRP route external if it does not originate in the same autonomous system as the router process that is receiving the route. RIP derived routes are always external, as are EIGRP routes redistributed from another autonomous system.
- **Feasible successor** – An attempt is made to move a destination entry from the topology table to the routing table when there is a feasible successor. All minimum cost paths to the destination form a set. From this set, the neighbors that have an advertised metric less than the current routing table metric are considered feasible successors. Feasible successors are viewed by a router as neighbors that are downstream with respect to the destination. These neighbors and the associated metrics are placed in

the forwarding table. When a neighbor changes the metric it has been advertising, or a topology change occurs in the network, the set of feasible successors may have to be re-evaluated. However, this is not categorized as a route recomputation.

- **Incremental SAP updates** – SAP updates that are only sent when a change occurs in the SAP information.
- **Infinity** – 4294967295 (–1 or 32 bits of all ones).
- **Internal route** – A router considers an EIGRP route internal if it originated in the same autonomous system as the router process that is receiving the route. Only networks that are directly connected to a Cisco router running EIGRP can be internal.
- **Neighbor (or Peer)** – Two routers that are connected to each other with a common network are known as adjacent neighbors. Neighbors dynamically discover each other and exchange EIGRP protocol messages. Each router keeps a topology table containing information learned from each of its neighbors.
- **Neighbor table** – Each router keeps state about adjacent neighbors. When newly discovered neighbors are learned, the address and interface of the neighbor is recorded. This information is stored in the neighbor data structure. The neighbor table holds these entries. There is one neighbor table for each protocol dependent module. When a neighbor sends a hello, it advertises a holdtime. The holdtime is the amount of time a router treats a neighbor as reachable and operational. If a hello packet isn't heard within the holdtime, then the holdtime will expire. When the holdtime expires, DUAL is informed of the topology change. The neighbor table entry also includes information required by the reliable transport mechanism. Sequence numbers are employed to match acknowledgments with data packets. The last sequence number received from the neighbor is recorded, so out of order packets can be detected. A transmission list is used to queue packets for possible retransmission on a per neighbor basis. Round trip timers are kept in the neighbor data structure to estimate an optimal retransmission interval.
- **Passive state** – A topology table entry is in passive state when the router is not performing a route recomputation for this destination.
- **Query** – A type of EIGRP packet that is sent to all EIGRP neighbors when a reroute recomputation commences. See the references for more information.
- **Redistribution** – In addition to running IPX–RIP and IPX–EIGRP simultaneously, the router can redistribute information from one routing protocol to another. The RIP metric does not directly translate into the IPX–EIGRP metric, and vice-versa, so an artificial metric is assigned to the redistributed route. The router uses the following artificial metrics in redistribution:
  - ◆ **RIP to EIGRP** – The reliability, load, and Maximum Transmission Unit (MTU) of the interface the RIP route was received on, plus the IPX ticks converted into tens of microseconds are used as the IPX–EIGRP metric. The RIP hop count and RIP ticks are preserved and passed with the IPX–EIGRP update throughout the network to be used in routing loop detection and for redistribution back into RIP.
  - ◆ **EIGRP to RIP** – The RIP hop count and ticks that were recorded when the route was first redistributed from RIP into EIGRP (see above) are incremented by one and announced in RIP. This causes an entire EIGRP autonomous system, no matter what its size, to appear as one RIP hop count away. To prevent a destination more than 223 hops away from being advertised into RIP, if the EIGRP hop count (which is incremented for each hop in the EIGRP autonomous system) plus the original RIP hop count exceeds 223, the destination is considered unreachable and not redistributed into RIP. Internal EIGRP routes are advertised with a RIP metric of one.
- **Reply** – A type of EIGRP packet that is sent in response to a query from a neighbor. See references.
- **Split horizon** – Normally, routers that are connected to broadcast-type IPX networks and that use distance vector routing protocols, employ the split horizon mechanism to prevent routing loops. Split horizon blocks information about routes from being advertised by a router out any interface from which that information originated. Since DUAL provides for loop freedom, split horizon isn't necessary, but may be tuned on or off on any interface. To save bandwidth, it is on by default. Customers with Frame Relay or Switched Multimegabit Data Service (SMDS) networks may wish to

turn it off on those interfaces.

- **Successor** – A neighboring router that has met the feasibility condition and has been selected as the next hop for forwarding packets.
- **Topology table** – The topology table is populated by the IPX routing process and acted upon by the DUAL finite state machine. It contains all destinations advertised by neighboring routers. Associated with each entry is the destination address and a list of neighbors that have advertised the destination. For each neighbor, the advertised metric is recorded. This is the metric that the neighbor stores in its routing table. If the neighbor is advertising this destination, it must be using the route to forward packets. This is an important rule that distance vector protocols must follow. Also associated with the destination is the metric that the router uses to reach the destination. This is the sum of the best advertised metric from all neighbors plus the link cost to the best neighbor. This is the metric that the router uses in the routing table and to advertise to other routers.
- **Update** – A type of EIGRP packet that is sent containing EIGRP routing information. See references.

## Understanding Routing and Topology Tables

RIP routes are automatically redistributed into EIGRP, and EIGRP routes are automatically redistributed into RIP, without any redistribution commands being entered by the user. Redistribution between different EIGRP processes is not turned on by default.

EIGRP routes are preferred over RIP routes, except when the external hop count in the EIGRP advertisement is greater than the RIP hop count. The external hop count is the RIP hop count that was used to advertise this route when it originally entered the EIGRP autonomous system.

Internal EIGRP routes are always preferred over external EIGRP routes. This means that given two EIGRP paths to a destination, the path that was originated in the EIGRP autonomous system will always be preferred over the EIGRP path that did not originate in the autonomous system, regardless of metric. Redistributed RIP routes are always advertised in EIGRP as external.

All EIGRP routes that are received for a destination, and are determined to be feasible successors, are placed in the topology table. If a RIP route is the current preferred path to a destination, and that destination is also being advertised in EIGRP, then the RIP route will also appear in the topology table (it is denoted with the word redistributed in the via field). RIP routes that are not being used in the routing table will not appear in the topology table. EIGRP routes that are not being used in the routing table will appear in the topology table.

A route will be in the routing table, but not in the topology table when 1) it is attached, but not listed in the router subcommand network list and no neighbors are advertising it, or 2) it is a RIP route and we have no EIGRP neighbors that are announcing it and RIP redistribution is turned off.

A topology table entry will have zero successors when it is attached, but not in the router subcommand network list. The router has at least one neighbor announcing this network. This will usually be observed when the **no redistribute rip** command is issued.

In all other cases, routes in the routing table should be in the topology table and those entries should have a non-zero successor count.

## EIGRP Packet Format

IPX EIGRP packets are carried in an IPX packet that begins with a standard IPX header. A value of 0x85BE in the Socket field of the header, along with a value of 0 (unknown) in the Packet Type field, identifies an EIGRP packet. These packets consist of a standard EIGRP header, followed by a set of variable-length fields consisting of Type/Length/Value (TLV) triplets. The following table shows the format of an EIGRP packet header.

Field	Length, in Bytes	Description
Version	1	EIGRP version. There are two major revisions of EIGRP, versions 0 and 1. Cisco IOS <sup>®</sup> software versions earlier than 10.3(11), 11.0(8), and 11.1(3) run the earlier version of EIGRP.
Opcode	1	One of the following values: <ul style="list-style-type: none"> <li>• 1---Update</li> <li>• 3---Query</li> <li>• 4---Reply</li> <li>• 5---Hello</li> <li>• 6---IPX SAP</li> </ul>
Checksum	2	Standard IP checksum over the entire packet, including the EIGRP header. The IP header is not included.
Flags	4	One of the following values: <ul style="list-style-type: none"> <li>• 0x00000001---Init</li> <li>• 0x00000002---Conditional receive</li> </ul>
Sequence	4	32-bit sequence number.
Ack	4	32-bit sequence number. A hello packet with a nonzero ACK field should be decoded as an Acknowledgement (ACK) packet rather than a hello packet.
AS Number	4	Autonomous system number.

Following the EIGRP header is one or more TLV. The following table lists general and IPX-specific TLVs.

Number	Type
<b>General TLV Types</b>	
0x0001	Enhanced IGRP parameters
0x0003	Sequence
0x0004	Software Version
0x0005	Next multicast sequence
<b>IPX-Specific TLV Types</b>	
0x0302	IPX internal routes
0x0303	IPX external routes

## IPX-Specific TLVs

### IPX Internal Routes

The IPX internal routes TLV (TLV type 0x0302) consists of a header followed by one or more destination network addresses. The following table lists the fields in this header. Each network number is four bytes in length.

Field	Length, in Bytes	Description
Next hop network	4	Network that is the next hop.
Next hop host	6	Host that is the next hop.
Delay	4	In units of 10 msec/256. A delay of 0xFFFFFFFF indicates an unreachable route.
Bandwidth	4	In units of 2,560,000,000/kbps
MTU	3	Packet MTU size.
Hop count	1	Current hop count.
Reliability	1	A value of 255 indicates 100 percent reliability.
Load	1	A value of 255 indicates 100 percent load.
Reserved	2	Unused

### IPX External Routes

The IPX external routes TLV (TLV type 0x0303) consists of a header followed by one or more destination network addresses. The following table lists the fields in this header. Each network number is four bytes in length.

Unlike the internal routes TLV, the external routes TLV includes such fields as the AS number, the external metric, and the external delay.

Field	Length, in Bytes	Description
Next hop network	4	Network that is the next hop.
Next hop host	6	Host that is the next hop.
Router ID	6	Router ID of originating router.
AS number	4	Identification number of the EIGRP domain.
Arbitrary tag	4	May be used to carry a tag set by route maps.
Protocol ID	1	One of the following values:

		<ul style="list-style-type: none"> <li>• 1---Enhanced IGRP</li> <li>• 2---Static</li> <li>• 3---RIP</li> <li>• 4---Connected</li> <li>• 5---IS-IS</li> <li>• 6---NetWare Link Services Protocol (NLSP)</li> <li>• 7---Internal</li> </ul>
Reserved	1	Unused
External metric	2	Hop count of a redistributed RIP route. IPX RIP routes are redistributed automatically into IPX EIGRP as external routes. The IPX RIP metric is copied into the external data portion of the EIGRP route. When an IPX EIGRP route is redistributed back into IPX RIP, the RIP hop count is set to the RIP hop count in the original redistribution point, incremented by one.
External delay	2	Delay value of a redistributed route. When an IPX EIGRP route is redistributed back into IPX RIP, the IPX delay field of the RIP route is set to the IPX delay value in the external metric field.
Delay	4	In units of 10 msec/256. A delay of 0xFFFFFFFF indicates an unreachable route.
Bandwidth	4	In units of 2,560,000,000/kbps
MTU	3	Packet MTU size.
Hop count	1	Current hop count.
Reliability	1	A value of 255 indicates 100 percent reliability.
Load	1	A value of 255 indicates 100 percent load.
Reserved	2	Unused.

## IPX SAP Packets

When they are carried inside EIGRP packets, IPX SAP packets consist of a standard EIGRP header with an Opcode value of 6 (refer to the first table of this section), followed by the standard payload of a standard IPX SAP packet without the original IPX header. Each IPX SAP packet generated by a Cisco router can carry up to seven 64-byte SAP entries plus 32 bytes of IPX overhead (for a total of 480 bytes), plus the media encapsulation overhead.

## IPX-EIGRP Configuration Commands

### Global IPX Commands

<b>[no] ipx</b>	To enable IPX routing, use the <b>ipx routing</b> global
-----------------	--



<b>routing</b> [ <i>node</i> ]	configuration command. If you omit <i>node</i> , the Cisco IOS software uses the hardware MAC address currently assigned to it as its node address. This is the MAC address of the first Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI) card. If no satisfactory interfaces are present in the router (such as only serial interfaces), you must specify <i>node</i> . The <b>ipx routing</b> command enables IPX-RIP and SAP services.
<b>ipx router</b> { <b>eigrp</b> <b>AS-number</b>   <b>nlsp</b> [ <i>tag</i> ]   <b>rip</b> }	Enables EIGRP. The argument <i>autonomous-system-number</i> is the EIGRP autonomous system number. It can be a number

from one to 65535.

## Router Subcommands

<b>[no] network</b> {< <i>network-number</i> > <b>all</b> }	Use the <b>network</b> command to enable the routing protocol specified in the <b>ipx router</b> command on each network.
<b>[no] redistribute</b> { <b>rip</b>   <b>igrp</b> < <i>as-number</i> >}	Configures redistribution of one protocol into another. This command is enabled by default.

The **no** form is used to disable redistribution.

**Note:** If you want to run EIGRP or RIP on many, but not all interfaces, enter the **all** form of this command followed by **no network <network-number>**, where <network-number> is the network you do not want to run the routing protocol on.

## Interface Subcommands

<b>[no] ipx sap-incremental eigrp</b> < <i>as-number</i> > [ <b>rsup-only</b> ]	To send SAP updates only when a change occurs in the SAP table, use the <b>ipx sap-incremental</b> interface configuration command. To send periodic SAP updates, use the <b>no</b> form of this command. The <b>rsup-only</b> option indicates that the system uses EIGRP on the interface to carry reliable SAP update information only. RIP routing updates are used, and EIGRP routing updates are ignored.
<b>[No] ipx hello-interval eigrp</b> < <i>as-number</i> >< <i>value</i> >	Configures the hello-interval in seconds

	<p>on the interface for the IPX–EIGRP routing process designated by. The default value is five seconds. This value may set the holding time advertised in hello packets. The holding time is three times the hello–interval. If the current value for the holding time is less than two times the hello–interval, the holding time will be reset. The default holding time is 15 seconds.</p>
<p><b>[No] ipx hold–time eigrp</b>  <i>&lt;as–number&gt; &lt;value&gt;</i></p>	<p>Configures the hold–time in seconds on the interface for the IPX–EIGRP routing process designated. The holding time is advertised in hello packets and indicates to neighbors the length of time they should consider the sender valid. The default holding time is three times the hello–interval. The default holding time is 15 seconds.</p>

## show Commands

<p><b>show ipx route</b>  <i>[network] [default]</i>  <b>[detailed]</b></p>	<p>To display the contents of the IPX routing table, use the <b>show ipx route</b> user EXEC command. <b>default</b> option displays the default route. <b>detailed</b> option displays detailed route information.</p>
<p><b>show ipx eigrp neighbors</b> <b>[servers]</b>  <i>[as–number   interface] [regex name]</i></p>	<p>To display the neighbors discovered by EIGRP, use the <b>show ipx eigrp neighbors</b> EXEC command. <b>servers</b> option displays the server list advertised by each neighbor. <b>regex</b> name option displays the IPX servers whose names</p>
<p><b>show ipx eigrp topology</b>  <i>[network–number]</i></p>	<p>To display the EIGRP topology table, use the <b>show ipx eigrp topology</b> EXEC command. <b>network–number</b> displays</p>

topology table of the IPX network number entered.
---

## debug Commands

<b>[no] debug eigrp packets</b>	Use the <b>debug eigrp packet</b> EXEC command to display general debugging information. The <b>no</b> form of <b>debug eigrp packets</b> disables debugging output.
<b>[no] debug eigrp fsm</b>	Use the <b>debug eigrp fsm</b> EXEC command to display debugging information about EIGRP feasible successor metrics (FSM). The <b>no</b> form of this command disables debugging output.

These configuration examples were tested on Cisco 2500 series routers with IOS version 12.0(4).

In the following example, we have configured interfaces Ethernet0 and Serial0 for IPX–EIGRP routing in an autonomous system number 100:

```
!  
ipx routing 0000.0c5c.ec39
```

**Note:** By default, the IPX process takes the MAC address of the first active Ethernet, Token Ring, or FDDI interface when IPX routing is enabled.

```
!  
interface Ethernet0  
ipx network AA  
!  
interface Serial0  
ipx network 10  
!ipx router eigrp 100  
network AA  
network 10  
!  
!  
no ipx router rip  
!
```

**Note:** IPX–RIP is disabled by using the command **no ipx router rip** (IPX–RIP is enabled by default when IPX routing is configured). If there is a non–Cisco device, such as Novell server, connected to the LAN segment, then RIP (or NLSP) must be running on the LAN interface for the router to see it. Be aware that NLSP is not redistributed into EIGRP by default.

When EIGRP is enabled, by default, SAPs are sent periodically on Ethernet interfaces and incrementally on serial interfaces. If Ethernet0 has only IPX–EIGRP peers present, you may want to reduce bandwidth use and only send SAPs incrementally. To do this, use the following commands:

```
!  
ipx routing 0000.0c5c.ec39  
!  
interface ethernet0  
ipx network AA  
ipx sap-incremental eigrp 100  
!  
interface serial0  
ipx network 10  
!
```

```

ipx router eigrp 100
network AA
network 10
!
no ipx router rip
!

```

**Note:** If the **ipx sap-incremental eigrp 100** command is configured on the Ethernet interface and no IPX-EIGRP peer is found, then SAP updates will be sent periodically. When a peer is found, then updates will be sent incrementally as intended (that is, when changes occur in the SAP table). Any router interfaces configured for periodic SAPs that instead receive incremental SAPs will not have complete SAP information from this router. Thus, when any two routers are enabled for SAP incremental, all other routers on that network segment must also be configured for SAP incremental.

If you wish to send periodic SAP updates on a serial interface which has an IPX-EIGRP peer on the other side, use the following commands to disable incremental SAP and enable periodic SAP updates:

```

!
ipx routing 0000.0c5c.ec39
!
interface ethernet0
ipx network AA
!interface serial0
ipx network 10
no ipx sap-incremental eigrp 100
!
ipx router eigrp 100
network AA
network 10
!
no ipx router rip
!

```

In most networks, one configures RIP on LAN interfaces and EIGRP on WAN interfaces. This is to avoid bandwidth hungry periodic RIP and SAP updates traversing bandwidth sensitive WAN interfaces. When configured as such, the Cisco router automatically redistributes IPX-RIP routes into EIGRP, and vice versa. Below, we have enabled IPX-RIP on an Ethernet interface, and IPX-EIGRP on a serial interface:

```

!
ipx routing 0000.0c5c.ec39
!
interface Ethernet0
ipx network AA
!
interface Serial0
ipx network 10
!
ipx router eigrp 100
network 10
!

```

**Note:** Here, IPX-RIP is enabled on Ethernet interface even though it isn't displayed in the running configuration. This is because IPX-RIP is enabled by default on all interfaces when IPX routing is enabled and any parameter enabled by default does not appear in the running configuration.

It is also possible to have periodic RIP and incremental SAP on a serial interface to reduce the SAP traffic. To do this, use the **rsup-only** option with the **ipx sap-incremental** command:

```

!
ipx routing 0000.0c5c.ec39
!

```

```

interface Ethernet0
ipx network AA
!
interface Serial0
ipx network 10
ipx sap-incremental eigrp 100 rsup-only
!
ipx router eigrp 100
network 10
!

```

**Note:** With the **rsup-only** option, RIPs are instead sent periodically; SAPs continue to be sent incrementally.

On very congested large networks, the default 15 seconds holdtime may not be sufficient for all routers to receive hello packets from their neighbors. In this case, you may want to increase the holdtime. In this example, we have increased the holdtime to 45 seconds:

```

!
ipx routing 0000.0c5c.ec39
!
interface ethernet 0
ipx network AA
!
interface serial 0
ipx network 10
ipx hold-time eigrp 100 45
!
ipx router eigrp 100
network AA
network 10
!

```

## Output of show Commands

R1#

show ipx route

Codes:

C - Connected primary network, c - Connected secondary network  
S - Static, F - Floating static, L - Local (internal), W - IPXWAN  
R - RIP, E - EIGRP, N - NLSP, X - External, A - Aggregate  
s - seconds, u - uses, U - Per-user static

5 Total IPX routes. Up to 1 parallel paths and 16 hops allowed.

No default route known.

```

C                10(HDLC)                Se0
C                AA (NOVELL-ETHER)       Et0
E                20 [41024000/0]via      10.0000.0c3b.ed69,
  age 00:26:43, 1u, Se0
E                BB [40537600/0]via      10.0000.0c3b.ed69,
  age 00:26:44, 1u, Se0
E                CC [41049600/0]via      10.0000.0c3b.ed69,
  age 00:26:44, 1u, Se0

```

R1#

**Note:** A value of EH for the route source indicates that the IPX EIGRP route is an active state while the local router is waiting for all relevant neighbors to respond to a query. Thus, this value should be a temporary state only.

R1#

show ipx eigrp neighbors

```

IPX EIGRP Neighbors for process 100
H Address          Interface      Hold Uptime  SRTT      RTO      Q Seq
                   (sec)         (ms)      Cnt      Num
0 10.0000.0c3b.ed69 Se0           12 00:28:10   30     2280    0   51
R1#

R1#
show ipx eigrp topology
IPX EIGRP Topology Table for process 100
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply, r - Reply status
P 10, 1 successors, FD is 40512000 via Connected, Serial0
P 20, 1 successors, FD is 41024000 via 10.0000.0c3b.ed69 (41024000/2169856), Serial0
P AA, 1 successors, FD is 281600 via Connected, Ethernet0
P BB, 1 successors, FD is 40537600 via 10.0000.0c3b.ed69 (40537600/281600), Serial0
P CC, 1 successors, FD is 41049600 via 10.0000.0c3b.ed69 (41049600/2195456), Serial0
R1#

R1#
show ipx eigrp traffic
IP-EIGRP Traffic Statistics for process 10
Hellos sent/received: 3900/3012
Updates sent/received: 23/16
Queries sent/received: 9/8
Replies sent/received: 8/9
Acks sent/received: 24/29
Input queue high water mark 2, 0 drops
R1#

```

## Troubleshooting Neighbor Relationships

Routers running EIGRP maintain state information about adjacent neighbors in a neighbor table. When a neighbor sends a hello, it advertises a hold time, which defines the duration over which the neighbor is considered reachable and operational. If a new hello packet is not received within the hold time, EIGRP declares the neighbor unreachable and begins to update its topology table. Both IP and IPX EIGRP use a default hello interval of 5 seconds for all interfaces other than non-broadcast multi-access networks with speeds of T1 or less, which use a default hello time of 60 seconds. By default, the hold timer is three times the value of the hello interval. For more information, refer to the command reference discussion of the **ipx hello-interval eigrp** command.

The EIGRP neighbor table also saves information required by the reliable transport mechanism. Sequence numbers are employed to match acknowledgments with data packets. The last sequence number received from the neighbor is recorded so out-of-order packets can be detected. A transmission list is used to queue packets for possible retransmission on a per-neighbor basis.

If the uptime in the output from the **show ipx eigrp neighbor** command never gets above approximately 80 seconds, it may be that the local router is hearing the neighbor's hellos, but the neighbor is not hearing the local router's hellos. While Open Shortest Path First (OSPF) requires a two-way exchange of hellos before a neighbor is declared, EIGRP will try to form a relationship as soon as it receives a hello from an adjoining router. If you have a one-way link, the router hearing the hello puts the adjoining router into the neighbor table, but shortly thereafter it will reset the connection since the neighboring router will not reply with the necessary packets needed to complete the formation of the neighbor relationship. Symptoms of this problem include the following:

- The local router does not appear in the remote router's neighbor table.
- The remote router's entry in the local router's neighbor table has a Smoothed Round Trip Time (SRTT) of 0.

Start your troubleshooting of unexpected EIGRP neighbor loss by enabling logging of neighbor changes.

Issue the **log-neighbor-changes** command in config-ipx-router mode. This command logs neighbor adjacency changes to monitor the stability of the routing system and to help you detect problems. By default, adjacency changes are not logged.

The following table lists sample output and explains how to interpret the output.

Log Message	Explanation
<pre>%DUAL-5-NBRCHANGE: IPX-EIGRP 2047: Neighbor x.y (Serial1/1/0.4) is up: new adjacency</pre>	<p>A hello has been received from an adjoining router, and the router is viewing this neighbor as brand new, although it may have known about it previously.</p>
<pre>%DUAL-5-NBRCHANGE: IPX-EIGRP 2047: Neighbor x.y (Serial1/1/0.6) is down: stuck in INIT state</pre>	<p>After receiving a hello, a router responds by sending an update packet with the initialization bit set. This packet prompts the adjacent router to queue its best entry for each network for transmission. If the adjacent router never responds, it appears as stuck in the INIT state in the local router's neighbor table. This problem typically is seen on a one-way link.</p>
<pre>%DUAL-5-NBRCHANGE: IPX-EIGRP 2047: Neighbor x.y (Serial1/1/0.1) is down: retry limit exceeded</pre>	<p>The local router sent an update, query, or reply, but did not receive an acknowledgment. Check Layer 1 (L1) and Layer 2 (L2) connectivity.</p>
<pre>%DUAL-5-NBRCHANGE: IPX-EIGRP 2047: Neighbor x.y (Serial1/1/0.4) is down: peer restarted</pre>	<p>The neighbor went down for an</p>

	<p>unknown reason and was detected when the local router received a hello or an update with the INIT flag set. To determine which router — local or remote — terminated the relationship, start by issuing the <b>show ipx eigrp neighbor</b> command. Look at the uptime and the Q Cnt values. The uptime value indicates how long since the neighbor relationship was last reset. The Q Cnt shows the number of packets waiting to be sent to the neighbor or that have been sent and have not been acknowledged. If the Q Cnt does not go to zero, the two EIGRP neighbors will not converge.</p>
<pre>%DUAL-5-NBRCHANGE: IPX-EIGRP 2047: Neighbor x.y (Serial1/1/0.4) is down: holding time expired</pre>	<p>If no hellos are received within the hold time, which is 15 seconds by default on most links, the router informs the neighbor that the neighbor relationship has been torn down and logs a syslog message.</p>

If you need more information beyond the above messages, try enabling specific IPX debugs. Ensure you understand the impact of debugs before enabling them.



- **debug eigrp packets** – May produce a large number of messages. Use with caution.
- **Debug eigrp packets terse** – Does not display EIGRP hellos.
- **Debug ipx eigrp events**
- **debug ipx eigrp** and also **debug ipx eigrp neigh** limit debugging information to a specific neighbor.

To minimize the impact of debug messages on the router, it is suggested that you disable the console logging and enabling buffered logging by issuing the **logging buffered** global configuration mode command.

The following are other points of consideration for troubleshooting IPX EIGRP neighbor relationships. After gathering answers to these questions, you should be able to narrow the fault domain for quicker resolution. For example, you should be able to isolate the problem to a particular router or a particular router's interface or packet queue.

- Did multiple neighbors on the same device bounce at the same time?
- What do the remote neighbors see?
- Which side initiated the tear down — the local router or the remote router?
- Is the interface congested? Is there tremendous delay in queuing the hello packets?
- If you are running IPX EIGRP over a low-speed link such as Frame Relay, look for drops in the interface broadcast queue. If you still are running RIP over the link even though you do not need it (since it is enabled by default when you enable IPX routing), try disabling RIP with the **no network {network number}** command in the router-rip configuration mode.

```
%DUAL-5-NBRCHANGE: IPX-EIGRP 1: Neighbor 95081004.0060.3e00.4000
  (Serial0.801) is down:
%DUAL-5-NBRCHANGE: IPX-EIGRP 1: Neighbor 95081004.0060.3e00.4000
  (Serial0.801) is up: new adjacency
```

## References

[1] A Unified Approach to Loop-Free Routing using Distance Vectors or Link States, J.J. Garcia-Luna-Aceves, 1989 ACM 089791-332-9/89/0009/0212, pages 212-223.

[2] Loop-Free Routing using Diffusing Computations, J.J. Garcia-Luna-Aceves, Network Information Center, SRI International, IEEE/ACM Transactions on Networking, Vol. 1, No. 1, 1993.

## Related Information

- **Switches Product Support**
- **LAN Switching Technology Support**
- **Technical Support & Documentation – Cisco Systems**

---

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Oct 04, 2005

Document ID: 10579

---