

Integrating Enhanced IGRP into Existing Networks

The Enhanced Interior Gateway Routing Protocol (IGRP) combines the ease of use of traditional routing protocols with the fast rerouting capabilities of link-state protocols, providing advanced capabilities for fast convergence and partial updates. When a network topology change occurs, the Diffusing Algorithm (DUAL) used with Enhanced IGRP provides convergence in less than five seconds in most cases. This is equivalent to the convergence achieved by link-state protocols such as Open Shortest Path First (OSPF), Novell Link Services Protocol (NLSP), and Intermediate System-to-Intermediate System (IS-IS). In addition, Enhanced IGRP sends routing update information only when changes occur, and only the changed information is sent to affected routers.

Enhanced IGRP supports three network level protocols: IP, AppleTalk, and Novell Internetwork Packet Exchange (IPX). Each of these has protocol-specific, value-added functionality. IP Enhanced IGRP supports variable-length subnet masks (VLSMs). IPX Novell Enhanced IGRP supports incremental Service Advertisement Protocol (SAP) updates, removes the Routing Information Protocol (RIP) limitation of 15 hop counts, and provides optimal path use. A router running AppleTalk Enhanced IGRP supports partial, bounded routing updates and provides load sharing and optimal path use.

The case study provided here discusses the benefits and considerations involved in integrating Enhanced IGRP into the following types of internetworks:

- *IP*—The existing IP network is running IGRP
- *Novell IPX*—The existing IPX network is running RIP and SAP
- *AppleTalk*—The existing AppleTalk network is running the Routing Table Maintenance Protocol (RTMP)

When integrating Enhanced IGRP into existing networks, plan a phased implementation. Add Enhanced IGRP at the periphery of the network by configuring Enhanced IGRP on a boundary router on the backbone off the core network. Then integrate Enhanced IGRP into the core network.

Note For a discussion of Enhanced IGRP network design considerations and details on DUAL convergence, see the *Internetwork Design Guide*.



Caution If you are using *candidate default route* in IP Enhanced IGRP and have installed multiple releases of Cisco router software within your internetwork that include any versions prior to September 1994, contact your Cisco technical support representative for version compatibility and software upgrade information. Refer to your software release notes for details. If you plan to implement Enhanced IGRP over a Frame Relay network, you should ensure that your network is hierarchical in design and adheres to sound design principles.

IP Network

This case study illustrates the integration of Enhanced IGRP into an IGRP internetwork in two phases: configuring an IGRP network and adding Enhanced IGRP to the network. The key considerations for integrating Enhanced IGRP into an IP network running IGRP are as follows:

- Route selection
- Metric handling
- Redistribution from IGRP to Enhanced IGRP and vice versa
- Route summarization

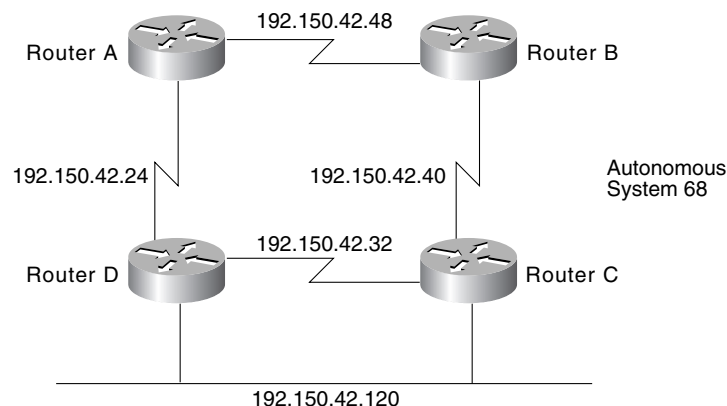
Configuring an IGRP Network

IGRP is a dynamic distance vector routing protocol designed by Cisco Systems in the mid-1980s for routing in an autonomous system (AS) containing large, arbitrarily complex networks with diverse media.

An autonomous system is a collection of interconnected routers under common management control, or with similar routing policies and requirements. Typically, an autonomous system consists of routers connecting multiple IP network numbers. Routes originating from one autonomous system that need to be advertised into other autonomous systems must be redistributed.

In Figure 4-1, Routers A, B, C, and D are configured to run IGRP in autonomous system 68.

Figure 4-1 Configuring an IGRP network.



The configuration commands to enable IGRP routing for Routers A, B, C, and D are as follows:

```
router igrp 68
network 192.150.42.0
```

Adding Enhanced IGRP to IGRP Networks

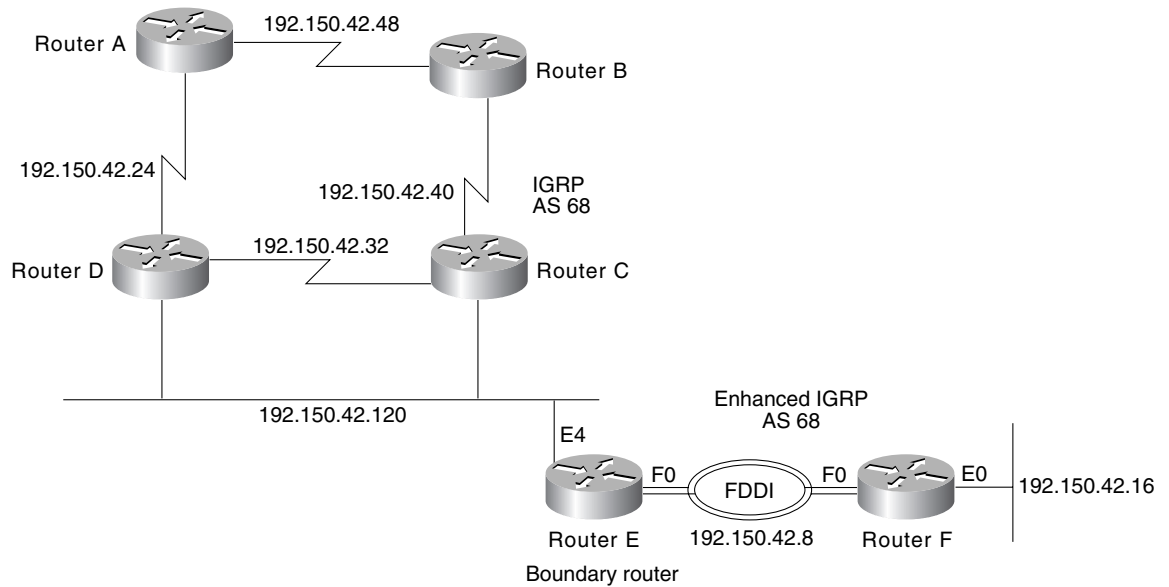
This section provides two examples of adding Enhanced IGRP to IGRP networks:

- Adding Enhanced IGRP to a Single IGRP Network
- Adding Enhanced IGRP to Multiple IGRP Networks

Adding Enhanced IGRP to a Single IGRP Network

In Figure 4-2, Router E acts as the boundary router, running both IGRP and Enhanced IGRP, and redistributing information between IGRP autonomous system 68 into the Enhanced IGRP autonomous system 68.

Figure 4-2 Adding Enhanced IGRP to a single IGRP network.



Router E, the boundary router, is configured to run both IGRP and Enhanced IGRP as follows:

```
router igrp 68
network 192.150.42.0
router eigrp 68
network 192.150.42.0
```

Note Redistribution is automatic because the autonomous system number for IGRP and Enhanced IGRP are the same.

Router F runs Enhanced IGRP only:

```
router eigrp 68
network 192.150.42.0
```

A **show ip route** command on Router E shows networks that are directly connected (C), routes learned from IGRP (I), and routes learned from Enhanced IGRP (D):

```
192.150.42.0 is subnetted (mask is 255.255.255.248), 7 subnets
C    192.150.42.120 is directly connected, Ethernet4
I    192.150.42.48 [100/2860] via 192.150.42.123, 0:00:08, Ethernet4
I    192.150.42.40 [100/2850] via 192.150.42.121, 0:00:08, Ethernet4
I    192.150.42.32 [100/2850] via 192.150.42.121, 0:00:08, Ethernet4
I    192.150.42.24 [100/2760] via 192.150.42.123, 0:00:08, Ethernet4
D    192.150.42.16 [90/30720] via 192.150.42.10, 0:00:38, Fddi0
C    192.150.42.8 is directly connected, Fddi0
```

A **show ip route** command on Router F shows that all routes are learned via enhanced IGRP (D) or are directly connected (C):

```
192.150.42.0 is subnetted (mask is 255.255.255.248), 7 subnets
D      192.150.42.120 [90/729600] via 192.150.42.9, 0:01:16, Fddi0
D EX   192.150.42.48 [170/757760] via 192.150.42.9, 0:01:16, Fddi0
D EX   192.150.42.40 [170/755200] via 192.150.42.9, 0:01:16, Fddi0
D EX   192.150.42.32 [170/755200] via 192.150.42.9, 0:01:16, Fddi0
D EX   192.150.42.24 [170/732160] via 192.150.42.9, 0:01:16, Fddi0
C      192.150.42.16 is directly connected, Ethernet0
C      192.150.42.8  is directly connected, Fddi0
```

Subnetwork 120 is seen as an internal route. All other routes are external (EX) because they were learned via IGRP in Router E and redistributed into Enhanced IGRP.


A **show ip eigrp topology** command on Router F shows that the state of each of the networks is passive (P) and that each network has one successor and lists the feasible distance (FD) of each successor via a neighbor to the destination. The computed/advertised metric is listed. Then the interface through which the neighbor network is available is provided.

```
IP-EIGRP Topology Table for process 68
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
P 192.150.42.120 255.255.255.248, 1 successors, FD is 2172416
   via 192.150.42.9 (2172416/2169856), Fddi0
P 192.150.42.8 255.255.255.248, 1 successors, FD is 28160
   via Connected, Fddi0
P 192.150.42.48 255.255.255.248, 1 successors, FD is 2560515840
   via 192.150.42.9 (2560515840/2560513280), Fddi0
P 192.150.42.16 255.255.255.248, 1 successors, FD is 281600
   via Connected, Ethernet0
P 192.150.42.40 255.255.255.248, 1 successors, FD is 2560026880
   via 192.150.42.9 (2560026880/2560001280), Fddi0
P 192.150.42.32 255.255.255.248, 1 successors, FD is 2560026880
   via 192.150.42.9 (2560026880/2560001280), Fddi0
```

Adding Enhanced IGRP to Multiple IGRP Networks

In Figure 4-3, Routers A, B, and C are connected to each other through several different networks. Routers A, B, and C are configured to run IGRP only within IGRP autonomous system (AS) 68. Router A redistributes static routes for subnetworks of network 9.0.0.0 (not shown). Assume that the IGRP AS continues at network 10.0.0.0.

Figure 4-3 Adding Enhanced IGRP to multiple IGRP networks.



The configuration for Router A is as follows:

```
router igrp 68
network 10.0.0.0
network 11.0.0.0
default-metric 1000 100 1 1 1500
redistribute static
ip route 9.1.0.0 255.255.0.0 e0
ip route 9.2.0.0 255.255.0.0 e1
```

The configuration for Router B is as follows:

```
router igrp 68
network 11.0.0.0
```

The configuration for Router C is as follows:

```
router igrp 68
network 11.0.0.0
network 12.0.0.0
```

This example takes you through the steps to add Enhanced IGRP to the internetwork one router at a time:

Step 1 Configure Enhanced IGRP for Router C as follows:

```
router eigrp 68
network 11.0.0.0
network 12.0.0.0
```

Because they are directly connected networks, Router C automatically summarizes networks 11.0.0.0 and 12.0.0.0 in its routing updates. Router C learns about networks 9.0.0.0 and 10.0.0.0 through IGRP. Networks 9.0.0.0 and 10.0.0.0 are already IGRP-summarized by Router A before they reach Router C.

Step 2 Configure Router A to run Enhanced IGRP as follows:

```
router eigrp 68
network 10.0.0.0
network 11.0.0.0
default-metric 1000 100 1 1 1500
redistribute static
```

Router A now automatically summarizes networks 10.0.0.0 and 11.0.0.0 in its Enhanced IGRP routing updates. It also continues to summarize these networks in its IGRP routing updates. However, automatic summarization of network 9.0.0.0 through Enhanced IGRP is not performed.

Router C now learns Enhanced IGRP routes for specific subnetworks of network 9.0.0.0 from Router A. At the same time, Router C continues to receive a summary route for network 9.0.0.0 through IGRP from Router A. The summary route for network 10.0.0.0, which Router C had previously learned through IGRP from Router A, is replaced with an Enhanced IGRP route in Router C's routing table.

- Step 3** Configure Router A to ensure that Router C does not unnecessarily learn about specific subnetworks of network 9.0.0.0. The following commands enable summarization of network 9.0.0.0 at Router A:

```
interface serial 1
ip summary-address eigrp 68 9.0.0.0 255.0.0.0
```

With this configuration on Router A, Router C's IGRP summary route for network 9.0.0.0 is replaced with an Enhanced IGRP summary route, and the more specific subnetworks of network 9.0.0.0 are no longer known by Router C.

- Step 4** Enable Enhanced IGRP on Router B as follows:

```
router eigrp 68
network 11.0.0.0
```

- Step 5** Ensure that Router B does not unnecessarily learn about specific subnetworks of network 9.0.0.0. Therefore, configure summarization of network 9.0.0.0 at Router A as follows:

```
interface serial 0
ip summary-address eigrp 68 9.0.0.0 255.0.0.0
```

With this configuration on Router A, Router B learns a summary route for network 12.0.0.0 through Enhanced IGRP from Router C. Router B learns summary routes for networks 9.0.0.0 and 10.0.0.0 through Enhanced IGRP from Router A.

- Step 6** Now that both of the next hop routers (Routers B and C) are running Enhanced IGRP, it is no longer necessary for these routers to run IGRP. Disable IGRP on Routers B and C with the following command:

```
no router igrp 68
```

Router A continues to run both IGRP and Enhanced IGRP and redistribute static routes.

If there were more routers on the network, you could continue deployment of Enhanced IGRP throughout network 10.0.0.0 one router at a time.

Route Selection

Enhanced IGRP uses three kinds of routes: internal, external, and summary. Internal routes are routes that are learned from Enhanced IGRP. External routes are routes that are learned from another protocol and then redistributed into Enhanced IGRP. Summary routes are routes that Enhanced IGRP may dynamically create due to auto summarization, or due to an explicit summary route configuration. Route selection is based on administrative distance. The default administrative distance for Enhanced IGRP is 90 (internal), 170 (external), or 5 (summary). For IGRP, the default administrative distance is 100 because internal Enhanced IGRP routes take precedence over IGRP routes, and IGRP routes are preferred to external Enhanced IGRP routes.

Metric Handling

The metric calculation and default metric value for IGRP and Enhanced IGRP are the same. By default, the composite metric is the sum of the segment delays and the lowest segment bandwidth (scaled and inverted) for a given route. Although you can adjust the default value with the **metric weights** command, the defaults were carefully selected to provide excellent operation in most networks.


Redistribution

Enhanced IGRP can be added to an IGRP network in two ways: using the same IGRP AS number or using a new AS number. If Enhanced IGRP uses the same AS number as IGRP, redistribution of IGRP into Enhanced IGRP and redistribution of Enhanced IGRP into IGRP occurs. If Enhanced IGRP uses a different AS number, the network administrator needs to configure redistribution manually with the **redistribute** command. For redistributing information from Enhanced IGRP into other dynamic routing protocols besides IGRP and vice versa, the designer must use the **redistribute** and **default-metric** commands. IGRP routes redistributed into Enhanced IGRP are marked as external.

Route Summarization

With IGRP, routing information advertised out an interface is often automatically summarized at major network number boundaries. Specifically, this automatic summarization occurs for those routes whose major network number differs from the major network number of the interface to which the advertisement is being sent. The remaining routes, which are part of the major network number of the interface, are advertised without summarization. For the following example, refer to Figure 4-4.

Figure 4-4 Route summarization.



In this example, Router A is directly connected to two different major networks and configured as follows:

```
interface ethernet 0
ip address 128.105.1.1 255.255.255.0
interface fddi 1
ip address 128.105.2.1 255.255.255.0
interface ethernet 2
ip address 128.106.1.1 255.255.255.0
router igrp 5
network 128.105.0.0
network 128.106.0.0
```

When advertising routing information out Ethernet interface 0, IGRP will summarize network 128.106.0.0 and will not summarize network 128.105.0.0. Therefore, IGRP will advertise routes for 128.106.0.0 with a network mask of 255.255.0.0 and routes for 128.105.2.1 with a network mask of 255.255.255.0.

Because it provides automatic route summarization, Enhanced IGRP will advertise the same routing information in the previous IGRP example. However, in the Enhanced IGRP example that follows, the previous configuration is modified so that it allows redistribution of routing information that is not summarized:

```
ip route 128.107.1.0 255.255.255.0 128.106.1.2
router eigrp 5
 redistribute static
 network 128.105.0.0
 network 128.106.0.0
router igrp 5
 redistribute static
```


At this point, there is a third subnetted major network in the IP routing table. When advertising out Ethernet interface 0, IGRP will summarize the route for 128.107.1.0 as 128.107.0.0 with a network mask of 255.255.0.0. However, Enhanced IGRP will not summarize network 128.107.0.0. It will advertise 128.107.1.0 with network mask 255.255.255.0. Enhanced IGRP's automatic summarization only applies to networks that are directly connected, not redistributed. For Enhanced IGRP, you can explicitly cause network 128.107.0.0 to be summarized out all three interfaces as shown in the following example:

```
interface ethernet 0
 ip summary-address eigrp 5 128.107.0.0 255.255.0.0
interface fddi 1
 ip summary-address eigrp 5 128.107.0.0 255.255.0.0
interface ethernet 2
 ip summary-address eigrp 5 128.107.0.0 255.255.0.0
```

Redistribution between Enhanced IGRP and RIP

Figure 4-5 shows a router that connects two networks; one network uses RIP and the other network uses Enhanced IGRP. The goal for the router is to advertise RIP routes in the Enhanced IGRP network and to advertise Enhanced IGRP routes in the RIP network, while preventing the occurrence of route feedback. (That is, the router must be configured so that Enhanced IGRP does not send routes learned from RIP back into the RIP network and so that RIP does not send routes learned from Enhanced IGRP back into the Enhanced IGRP network.)

Figure 4-5 Redistributing RIP routes into Enhanced IGRP.



The RIP portion of the configuration for Router A is as follows:

```
router rip
network 171.108.0.0
redistribute eigrp 90
default-metric 2
passive-interface serial 0
```

The **router rip** global configuration command starts a RIP process.

The **network** router configuration command specifies that the RIP process is to send RIP updates out on the interfaces that are directly connected to network number 171.108.0.0. In this case, the RIP process will send updates out on Ethernet interface 0 and not on serial interface 0 because of the **passive-interface** command applied to serial interface 0.

The **redistribute eigrp** router configuration command specifies that routing information derived from Enhanced IGRP be advertised in RIP routing updates.

The **default-metric** router configuration command causes RIP to use the same metric value (in this case, a hop count of 2) for all routes obtained from Enhanced IGRP. A default metric helps solve the problem of redistributing routes that have incompatible metrics. Whenever metrics do not convert, using a default metric provides a reasonable substitute and enables the redistribution to proceed.

The **passive-interface** router configuration command disables the sending of routing updates on serial interface 0. In this case, the **passive-interface** command is used with RIP, which means the router does not send out any updates on a passive interface, but the router still processes updates that it receives on that interface. The result is that the router still learns of networks that are behind a passive interface. (The same is true when the **passive-interface** command is used with IGRP.)

The Enhanced IGRP portion of the configuration for Router A is as follows:

```
router eigrp 90
network 171.108.0.0
redistribute rip
default-metric 1544 100 255 1 1500
distribute-list 1 in
passive interface ethernet 0
access-list 1 permit ip 171.108.1.0 255.255.255.0
access-list 1 permit ip 171.108.2.0 255.255.255.0
access-list 1 permit ip 171.108.3.0 255.255.255.0
access-list 1 permit ip 171.108.4.0 255.255.255.0
access-list 1 permit ip 171.108.5.0 255.255.255.0
access-list 1 permit ip 171.108.6.0 255.255.255.0
access-list 1 permit ip 171.108.7.0 255.255.255.0
access-list 1 permit ip 171.108.8.0 255.255.255.0
access-list 1 permit ip 171.108.9.0 255.255.255.0
access-list 1 permit ip 171.108.10.0 255.255.255.0
access-list 1 deny ip
```

The **router eigrp** global configuration command starts an Enhanced IGRP process and assigns to it autonomous system number 90.

The **network** router configuration command specifies that the Enhanced IGRP process is to send Enhanced IGRP updates to the interfaces that are directly connected to network number 171.108.0.0. In this case, the Enhanced IGRP process will send updates out on serial interface 0 and not on Ethernet interface 0 because of the **passive-interface** command applied to Ethernet interface 0.

The **redistribute eigrp** router configuration command specifies that routing information derived from RIP be advertised in Enhanced IGRP routing updates.

The **default-metric** router configuration command assigns an Enhanced IGRP metric to all RIP-derived routes. The first value (1544) specifies a minimum bandwidth of 1544 kilobits per second. The second value (100) specifies a route delay in tens of microseconds. The third value (255)

specifies the connection is guaranteed to be 100 percent reliable. The fourth value (1) specifies the effective bandwidth of the route. The fifth value (1500) specifies in bytes the maximum transmission unit (MTU) of the route.

The **distribute-list in** router configuration command causes the router to use access list 1 to filter networks learned from RIP and allows only those networks that match the list to be redistributed into Enhanced IGRP. This prevents route feedback loops from occurring.

When used with Enhanced IGRP, the **passive-interface** router configuration command has a different effect than it has when used with RIP or IGRP. When the **passive-interface** command is used with Enhanced IGRP, the router does not send out any updates—including hello messages—on the interface. Because hello messages are not sent, the router cannot discover any neighbors on that interface, which means that the router does not learn about networks that are behind a passive interface.

Access list 1 permits subnetworks 1 through 10 and denies all other networks. Although ten statements have been used, this particular access list could be written with four **access-list** commands if the address space had been divided efficiently. This example illustrates the need to think carefully about how to divide an address space. For example, if the RIP AS had been subnets 0 through 7, a single access list statement would have covered all of the subnetworks. The implication is that, when using a protocol that can summarize, summarization can be achieved much more efficiently when the IP address space is divided optimally. For information about dividing an IP address space optimally, see Appendix A, “Subnetting an IP Address Space.”


Novell IPX Network

This case study illustrates the integration of Enhanced IGRP into a Novell IPX internetwork in two phases: configuring an IPX network and adding Enhanced IGRP to the IPX network. The key considerations for integrating Enhanced IGRP into an IPX network running RIP and SAP are as follows:

- Route selection
- Redistribution metric handling
- Redistribution from IPX RIP to Enhanced IGRP and vice versa
- Reducing SAP traffic

Configuring a Novell IPX Network

Cisco’s implementation of Novell’s IPX protocol provides all the functions of a Novell router. In this case study, routers are configured to run Novell IPX. (See Figure 4-6.)

Figure 4-6 Configuring a Novell IPX network.

The configuration commands to enable IPX routing for Router A are as follows:

```

ipx routing
interface ethernet 0
ipx network 2ad
interface ethernet 1
ipx network 3bc

```

Note In Software Release 9.21 and later, the command to enable Novell IPX routing is **ipx** rather than **novell**.

Adding Enhanced IGRP to a Novell IPX Network


Enhanced IGRP for a Novell IPX network has the same fast rerouting and partial update capabilities as Enhanced IGRP for IP. In addition, Enhanced IGRP has several capabilities that are designed to facilitate the building of large, robust Novell IPX networks.

The first capability is support for incremental SAP updates. Novell IPX RIP routers send out large RIP and SAP updates every 60 seconds. This can consume substantial amounts of bandwidth. Enhanced IGRP for IPX sends out SAP updates only when changes occur and sends only changed information.

The second capability that Enhanced IGRP adds to IPX networks is the ability to build large networks. IPX RIP networks have a diameter limit of 15 hops. Enhanced IGRP networks can have a diameter of 224 hops.

The third capability that Enhanced IGRP for Novell IPX provides is optimal path selection. The RIP metric for route determination is based on ticks with hop count used as a tie-breaker. If more than one route has the same value for the tick metric, the route with the least number of hops is preferred. Instead of ticks and hop count, IPX Enhanced IGRP uses a combination of these metrics: delay, bandwidth, reliability, and load. For an illustration of how IPX Enhanced IGRP provides optimal path selection, see Figure 4-7.

Figure 4-7 Enhanced IGRP Novell IPX optimal path utilization.




Both Ethernet and FDDI interfaces have a tick value of 1. If configured for Novell RIP, Router A will choose the Ethernet connection via network 4 to reach network 5 because Router D is only one hop away from Router A. However, the fastest path to network 5 is two hops away, via the FDDI rings. With IPX Enhanced IGRP configured, Router A will automatically take the optimal path through Routers B and C to reach network 5.

To add Enhanced IGRP to a Novell RIP and SAP network, configure Enhanced IGRP on the Cisco router interfaces that connect to other Cisco routers also running Enhanced IGRP. Configure RIP and SAP on the interfaces that connect to Novell hosts and or Novell routers that do not support Enhanced IGRP.

In Figure 4-8, Routers E, F, and G are running IPX Enhanced IGRP. Router E redistributes Enhanced IGRP route information via Network AA to Router D.

Figure 4-8 Adding Enhanced IGRP to a Novell IPX network.



The configuration for Router E is as follows:

```
ipx routing
interface ethernet 0
ipx network AA
interface serial 0
ipx network 20
interface serial 1
ipx network 30
ipx router eigrp 10
network 20
network 30
ipx router rip
no network 20
```

With Enhanced IGRP configured, periodic SAP updates are replaced with Enhanced IGRP incremental updates when an Enhanced IGRP peer is found. Unless RIP is explicitly disabled for an IPX network number, as shown for network 20, both RIP and Enhanced IGRP will be active on the interface associated with that network number. Based on the above configuration, and assuming an Enhanced IGRP peer on each Enhanced IGRP configured interface, RIP updates are sent on networks AA and 30, while Enhanced IGRP routing updates are sent on networks 20 and 30. Incremental SAP updates are sent on network 20 and network 30, and periodic SAP updates are sent on network AA.

The configuration for Router F is as follows:

```
ipx routing
interface ethernet 0
ipx network 45
interface serial 0
ipx network 30
ipx router eigrp 10
network 30
network 45
```

Partial output for a **show ipx route** command on Router E indicates that network 45 was discovered using Enhanced IGRP (E), whereas network BB was discovered via a RIP (R) update:

```
R Net 3bc
R Net 2ad
C Net 20 (HDLC), is directly connected, 66 uses, Serial0
C Net 30 (HDLC), is directly connected, 73 uses, Serial1
E Net 45 [2195456/0] via 30.0000.0c00.c47e, age 0:01:23, 1 uses, Serial1
C Net AA (NOVELL-ETHER), is directly connected, 3 uses, Ethernet0
R Net BB [1/1] via AA.0000.0c03.8b25, 48 sec, 87 uses, Ethernet0
```

Partial output for a **show ipx route** command on Router F indicates that networks 20, AA, and BB were discovered using Enhanced IGRP (E):

```
E Net 20 [2681856/0] via 30.0000.0c01.f0ed, age 0:02:57, 1 uses, Serial0
C Net 30 (HDLC), is directly connected, 47 uses, Serial0
C Net 45 (NOVELL-ETHER), is directly connected, 45 uses, Ethernet0
E Net AA [267008000/0] via 30.0000.0c01.f0ed, age 0:02:57, 1 uses, Serial0
E Net BB [268416000/2] via 30.0000.0c01.f0ed, age 0:02:57, 11 uses, Serial0
```

A **show ipx servers** command on Router E shows that server information was learned via periodic (P) SAP updates:

```
Codes: S - Static, I - Incremental, P - Periodic, H - Holddown
5 Total IPX Servers
Table ordering is based on routing and server info
Type Name           Net Address      Port      RouteHopsItf
P    4 Networkers   100.0000.0000.0001:0666  2/022Et1
P    5 Chicago      100.0000.0000.0001:0234  2/022Et1
P    7 Michigan     100.0000.0000.0001:0123  2/022Et1
P    8 NetTest1     200.0000.0000.0001:0345  2/022Et1
P    8 NetTest      200.0000.0000.0001:0456  2/022Et1
```

A **show ipx servers** command on Router F shows that server information was learned via incremental SAP (I) updates allowed with Enhanced IGRP:

```
Codes: S - Static, I - Incremental, P - Periodic, H - Holddown
5 Total IPX Servers
Table ordering is based on routing and server info
Type Name           Net Address      Port      RouteHopsItf
I    4 Networkers   100.0000.0000.0001:0666  268416000/033Se0
I    5 Chicago      100.0000.0000.0001:0234  268416000/033Se0
I    7 Michigan     100.0000.0000.0001:0123  268416000/033Se0
I    8 NetTest1     200.0000.0000.0001:0345  268416000/033Se0
I    8 NetTest      200.0000.0000.0001:0456  268416000/033Se0
```

A **show ipx eigrp topology** command on Router E shows that the state of the networks is passive (P) and that each network provides one successor, and it lists the feasible distance (FD) of each successor via a neighbor to the destination. For example, for network 45, the neighbor is located at address 0000.0c00.c47e and the computed/advertised cost metric for that neighbor to the destination is 2195456/281600:

```
IPX EIGRP Topology Table for process 10
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
P 20, 1 successors, FD is 1
    via Connected, Serial0
P 30, 1 successors, FD is 1
    via Connected, Serial1
P 45, 1 successors, FD is 2195456
    via 30.0000.0c00.c47e (2195456/281600), Serial1
P AA, 1 successors, FD is 266496000
    via Redistributed (266496000/0),
P BB, 1 successors, FD is 267904000
    via Redistributed (267904000/0),
```

The output for a **show ipx eigrp topology** command on Router F lists the following information:

```
IPX EIGRP Topology Table for process 10
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
P 20, 1 successors, FD is 2681856
    via 30.0000.0c01.f0ed (2681856/2169856), Serial0
P 30, 1 successors, FD is 1
    via Connected, Serial0
P 45, 1 successors, FD is 1
    via Connected, Ethernet0
P AA, 1 successors, FD is 267008000
    via 30.0000.0c01.f0ed (267008000/266496000), Serial0
P BB, 1 successors, FD is 268416000
    via 30.0000.0c01.f0ed (268416000/267904000), Serial0
```

Route Selection

IPX Enhanced IGRP routes are automatically preferred over RIP routes regardless of metrics unless a RIP route has a hop count less than the external hop count carried in the Enhanced IGRP update, for example, a server advertising its own internal network.

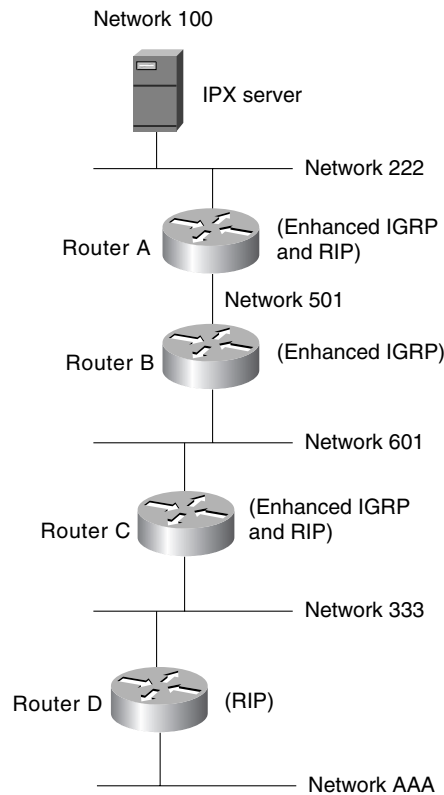
Redistribution and Metric Handling

Redistribution is automatic between RIP and Enhanced IGRP, and vice versa. Automatic redistribution can be turned off using the **no redistribute** command. Redistribution is not automatic between different Enhanced IGRP autonomous systems.

The metric handling for integrating RIP into Enhanced IGRP is bandwidth plus delay, left shifted by 8 bits. The metric handling for Enhanced IGRP to RIP is the external metric plus 1. An IPX Enhanced IGRP router that is redistributing RIP into Enhanced IGRP takes the RIP metric associated with each RIP route, increments it, and stores that metric in the Enhanced IGRP routing table as the external metric.

In Figure 4-9, a Novell IPX server with an internal network number of 100 advertises this network number using RIP on network 222. Router A hears this advertisement and installs it in its routing table as being 1 hop and 1 tick away. Router A then announces this network to Router B on network 501 using Enhanced IGRP.

Figure 4-9 IPX metric handling example.



The configuration for Router A is as follows:

```
ipx routing
!
interface ethernet 0
ipx network 222
!
interface serial 0
ipx network 501
!
ipx router eigrp 9000
network 222
network 501
!
!The following commands turn off IPX RIP on the serial interface:
!
ipx router rip
no network 501
```

The configuration for Router B is as follows:

```
ipx routing
!
interface ethernet 0
ipx network 601
!
interface serial 0
ipx network 501

ipx router eigrp 9000
network 501
network 601
!
!The following command turns off IPX RIP on this router:
!
no ipx router rip
```

The configuration for Router C is as follows:

```
ipx routing
!
interface ethernet 0
ipx network 333
!
interface ethernet 1
ipx network 601
!
ipx router eigrp 9000
network 333
network 601
!
!The following commands turn off IPX RIP on ethernet 1:
!
ipx router rip
no network 601
```

The configuration for Router D is as follows:

```
ipx routing
!
interface ethernet 0
ipx network 333
!
interface ethernet 1
ipx network AAA
```

The output from a **show ipx route** command on Router A is as follows:

```
R Net 100 [1/1] via 222.0260.8c4c.4f22, 59 sec, 1 uses, Ethernet0
C Net 222 (ARPA), is directly connected, 1252 uses, Ethernet0
E Net 333 [46277376/0] via 501.0000.0c05.84bc, age 0:04:07, 1 uses, Serial0
C Net 501 (HDLC), is directly connected, 3908 uses, Serial0
E Net 601 [46251776/0] via 501.0000.0c05.84bc, age 5:21:38, 1 uses, Serial0
E Net AAA [268441600/2] via 501.0000.0c05.84bc, age 0:16:23, 1 uses, Serial0
```

The output from a **show ipx route** command on Router B is as follows:

```
E Net 100 [268416000/2] via 501.0000.0c05.84b4, age 0:07:30, 2 uses, Serial0
E Net 222 [267008000/0] via 501.0000.0c05.84b4, age 0:07:30, 1 uses, Serial0
E Net 333 [307200/0] via 601.0000.0c05.84d3, age 0:07:30, 1 uses, Ethernet0
C Net 501 (HDLC), is directly connected, 4934 uses, Serial0
C Net 601 (NOVELL-ETHER), is directly connected, 16304 uses, Ethernet0
E Net AAA [267929600/2] via 601.0000.0c05.84d3, age 0:14:40, 1 uses, Ethernet0
```

The output from a **show ipx route** command on Router C is as follows:

```
E Net 100 [268441600/2] via 601.0000.0c05.84bf, age 0:07:33, 1 uses, Ethernet1
E Net 222 [267033600/0] via 601.0000.0c05.84bf, age 0:07:34, 1 uses, Ethernet1
C Net 333 (NOVELL-ETHER), is directly connected, 15121 uses, Ethernet0
E Net 501 [46251776/0] via 601.0000.0c05.84bf, age 0:07:32, 9 uses, Ethernet1
C Net 601 (NOVELL-ETHER), is directly connected, 1346 uses, Ethernet1
R Net AAA [1/1] via 333.0000.0c05.8b25, 35 sec, 1 uses, Ethernet0
```

The output from a **show ipx route** command on Router D is as follows:

```
R Net 100 [8/2] via 333.0000.0c05.84d1, 18 sec, 1 uses, Ethernet0
R Net 222 [6/1] via 333.0000.0c05.84d1, 18 sec, 1 uses, Ethernet0
R Net 333 [1/1] via 333.0000.0c05.84d1, 18 sec, 1 uses, Ethernet0
R Net 501 [3/1] via 333.0000.0c05.84d1, 17 sec, 3 uses, Ethernet0
R Net 601 [1/1] via 333.0000.0c05.84d1, 18 sec, 1 uses, Ethernet0
C Net AAA (SNAP), is directly connected, 20 uses, Ethernet1
```

The Enhanced IGRP metric is created using the RIP ticks for the delay vector. The hop count is incremented and stored as the external metric. The external delay is also stored. Router B computes the metric to network 100 given the information received from Router A and installs this in its routing table. In this case, the tick value for network 100 is 8.

The “2” after the slash in the routing entry for network 100 is the external metric. This number does not increase again while the route is in the Enhanced IGRP autonomous system. Router C computes the metric to network 100 through Router B and stores it in its routing table. Finally, Router C redistributes this information back into RIP with a hop count of 2 (the external metric) and a tick value derived from the original tick value of the RIP route (1) plus the Enhanced IGRP delay through the autonomous system converted to ticks.

Reducing SAP Traffic

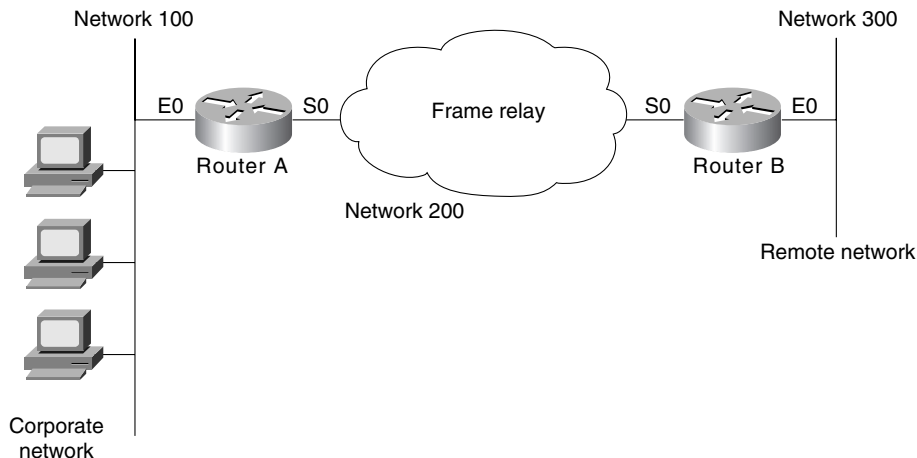
Novell IPX RIP routers send out large RIP and SAP updates every 60 seconds regardless of whether a change has occurred. These updates can consume a substantial amount of bandwidth. You can reduce SAP update traffic by configuring Enhanced IGRP to do incremental SAP updates. When Enhanced IGRP is configured for incremental SAP updates, the updates consist only of information that has changed and the updates are sent out only when a change occurs, thus saving bandwidth.

When you configure Enhanced IGRP for incremental SAP updates, you can do the following:

- Retain RIP, in which case only the reliable transport of Enhanced IGRP is used for sending incremental SAP updates. (This is the preferred configuration over bandwidth-sensitive connections.)
- Turn off RIP, in which case Enhanced IGRP replaces RIP as the routing protocol.

Figure 4-10 shows a bandwidth-sensitive topology in which configuring incremental SAP updates is especially useful. The topology consists of a corporate network that uses a 56-Kbps Frame Relay connection to communicate with a remote branch office. The corporate network has several Novell servers, each of which advertises many services. Depending on the number of servers and the number of advertised services, a large portion of the available bandwidth could easily be consumed by SAP updates.

Figure 4-10 Example of incremental SAP updates.



Router A is configured as follows:

```

ipx routing
!
interface ethernet 0
ipx network 100
!
interface serial 0
encapsulation frame-relay
!
interface serial 0.1 point-to-point
ipx network 200
ipx sap-incremental eigrp 90 rsup-only
frame-relay interface-dlci 101
!
ipx router eigrp 90
network 200
    
```

The **ipx routing** global configuration command enables IPX routing on the router.

The **ipx network** interface configuration command enables IPX routing on Ethernet interface 0 for network 100.

For serial interface 0, the **encapsulation frame-relay** interface configuration command establishes Frame Relay encapsulation using Cisco’s own encapsulation, which is a 4-byte header, with 2 bytes to identify the DLCI and 2 bytes to identify the packet type.

The **interface serial** global configuration command establishes a point-to-point subinterface (**0.1**). Subinterfaces are logical interfaces associated with a physical interface. Using subinterfaces allows Router A to receive multiple simultaneous connections over a single Frame Relay interface.

The **ipx network** interface configuration command enables IPX routing on subinterface serial interface 0.1 for network 200.

The **ipx sap-incremental** interface configuration command enables the incremental SAP feature. The required **eigrp** keyword enables Enhanced IGRP and its transport mechanism and, in this case, specifies an autonomous system number of 90. Because this command uses the **rsup-only** keyword, the router sends incremental SAP updates on this link.

The **frame-relay interface-dlci** interface configuration command associates data link connection identifier (DLCI) 101 with subinterface serial interface 0.1.

The **ipx router eigrp** global configuration command starts an Enhanced IGRP process and assigns to it autonomous system number 90.

The **network** IPX-router configuration command enables Enhanced IGRP for network 200.

Router B is configured as follows:

```
ipx routing
!
interface ethernet 0
ipx network 300
!
interface serial 0
encapsulation frame-relay
ipx network 200
ipx sap-incremental eigrp 90 rsup-only
!
ipx router eigrp 90
network 200
```

The **ipx routing** global configuration command enables IPX routing on the router.

The **ipx network** interface configuration command enables IPX routing on Ethernet interface 0 for network 300.

On serial interface 0, the **encapsulation frame-relay** interface configuration command establishes Frame Relay encapsulation using Cisco's own encapsulation, which is a 4-byte header, with 2 bytes to identify the DLCI and 2 bytes to identify the packet type.

The **ipx network** interface configuration command enables IPX routing on subinterface serial 0 for network 200.

The **ipx sap-incremental** interface configuration command enables the incremental SAP feature. The required **eigrp** keyword enables Enhanced IGRP and its transport mechanism and, in this case, specifies an autonomous system number of 90. Because this command uses the **rsup-only** keyword, the router sends incremental SAP updates on this link.

The **ipx router eigrp** global configuration command starts an Enhanced IGRP process and assigns to it autonomous system number 90.

The **network** IPX-router configuration command enables Enhanced IGRP for network 200.

Note The absence of the **ipx router rip** command means the IPX RIP is still being used for IPX routing, and the use of the **rsup-only** keyword means that the router is sending incremental SAP updates over the Frame Relay link.

AppleTalk Network


This case study illustrates the integration of Enhanced IGRP into an existing AppleTalk internetwork in two phases: configuring an AppleTalk network and adding Enhanced IGRP to an AppleTalk network. The key considerations for integrating Enhanced IGRP into an AppleTalk network are as follows:

- Route selection
- Metric handling
- Redistribution from AppleTalk to Enhanced IGRP and vice versa

Configuring an AppleTalk Network

Cisco routers support AppleTalk Phase 1 and AppleTalk Phase 2. For AppleTalk Phase 2, Cisco routers support both extended and nonextended networks. In this case study, Routers A, B, and C are running AppleTalk, as illustrated in Figure 4-11.

Figure 4-11 Configuring an AppleTalk network.



The configuration for Router A is as follows:

```

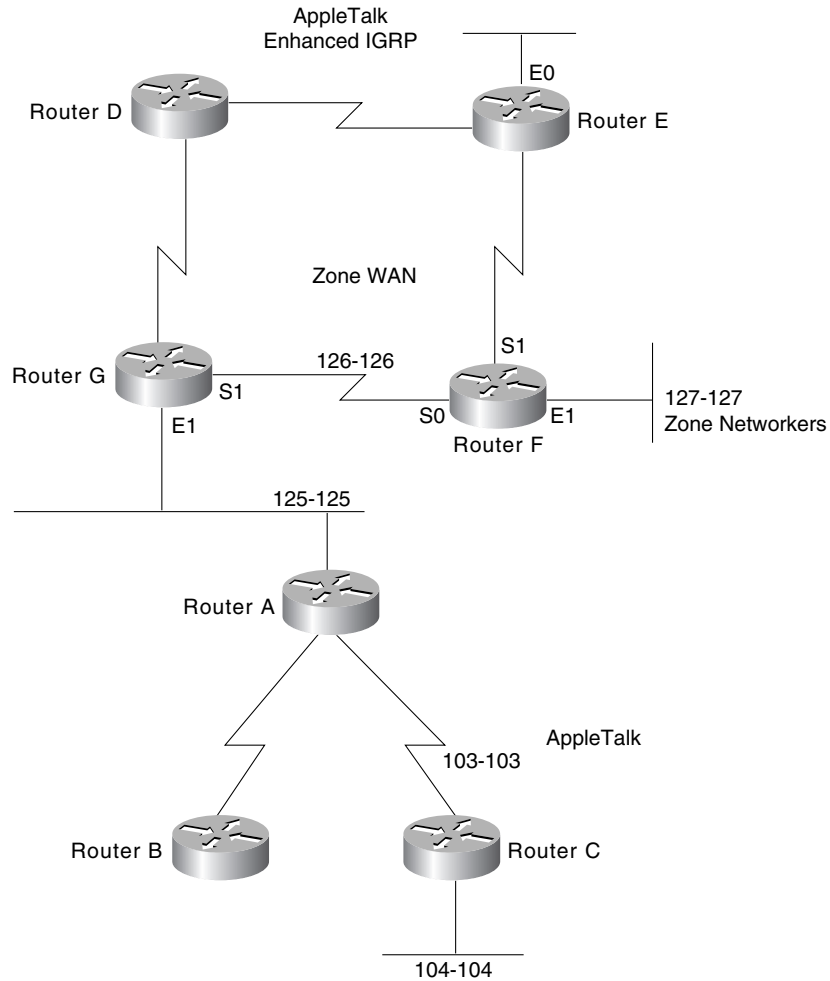
appletalk routing
interface ethernet 0
appletalk cable-range 10-10
appletalk zone casestudy
interface serial 0
appletalk cable-range 50-50
appletalk zone casestudy
  
```

Adding Enhanced IGRP to an AppleTalk Network

To add Enhanced IGRP to an AppleTalk network, configure Enhanced IGRP on the interface that connects to the routers. Do not disable RTMP on the interfaces that connect to AppleTalk hosts or that connect to AppleTalk routers that do not support Enhanced IGRP. RTMP is enabled by default when AppleTalk routing is enabled and when an interface is assigned an AppleTalk cable range.

In this case study, Routers D and E are running AppleTalk Enhanced IGRP. Routers F and G run both AppleTalk and AppleTalk Enhanced IGRP. Router G redistributes the routes from the AppleTalk network to the AppleTalk Enhanced IGRP network, and vice versa. (See Figure 4-12.)

Figure 4-12 Example of adding Enhanced IGRP to an AppleTalk network.



The configuration for Router G is as follows:

```

appletalk routing eigrp 1
interface ethernet 1
appletalk cable-range 125-125
appletalk zone Marketing Lab
appletalk protocol eigrp
interface serial 1
appletalk cable-range 126-126
appletalk zone WAN
appletalk protocol eigrp
no appletalk protocol rtmp
    
```

The configuration for Router F is as follows:

```

appletalk routing eigrp 2
interface serial 0
appletalk cable-range 126-126
appletalk zone WAN
appletalk protocol eigrp
no appletalk protocol rtmp
    
```

A **show appletalk route** command on Router G shows that the first set of routes is learned from an RTMP update, that the second set of routes is directly connected, and that the last route is learned by AppleTalk Enhanced IGRP via serial interface 1:

```
R Net 103-103 [1/G] via 125.220, 0 sec, Ethernet1, zone Marketing Lab
R Net 104-104 [1/G] via 125.220, 1 sec, Ethernet1, zone Marketing Lab
R Net 105-105 [1/G] via 125.220, 1 sec, Ethernet1, zone Marketing Lab
R Net 108-108 [1/G] via 125.220, 1 sec, Ethernet1, zone Marketing Lab
C Net 125-125 directly connected, Ethernet1, zone Marketing Lab
C Net 126-126 directly connected, Serial1, zone Wan
E Net 127-127 [1/G] via 126.201, 114 sec, Serial1, zone Networkers
```

A **show appletalk route** command on Router F shows that routes are learned from AppleTalk Enhanced IGRP:

```
E Net 103-103 [2/G] via 126.220, 519 sec, Serial0, zone Marketing Lab
E Net 104-104 [2/G] via 126.220, 520 sec, Serial0, zone Marketing Lab
E Net 105-105 [2/G] via 126.220, 520 sec, Serial0, zone Marketing Lab
E Net 108-108 [2/G] via 126.220, 520 sec, Serial0, zone Marketing Lab
E Net 125-125 [1/G] via 126.220, 520 sec, Serial0, zone Marketing Lab
C Net 126-126 directly connected, Serial0, zone Wan
C Net 127-127 directly connected, Ethernet1, zone Networkers
```

Route Selection

AppleTalk Enhanced IGRP routes are automatically preferred over Routing Table Maintenance Protocol (RTMP) routes. Whereas the AppleTalk metric for route determination is based on hop count only, AppleTalk Enhanced IGRP uses a combination of these configurable metrics: delay, bandwidth, reliability, and load.

Metric Handling

The formula for converting RTMP metrics to AppleTalk Enhanced IGRP metrics is hop count multiplied by 252524800. This is a constant based on the bandwidth for a 9.6-Kbps serial line and includes an RTMP factor. An RTMP hop distributed into Enhanced IGRP appears as a slightly worse path than an Enhanced IGRP-native, 9.6-Kbps serial link. The formula for converting Enhanced IGRP to RTMP is the value of the Enhanced IGRP external metric plus 1.

Redistribution

Redistribution between AppleTalk and Enhanced IGRP and vice versa is automatic by default. Redistribution involves converting the Enhanced IGRP metric back into an RTMP hop count metric. In reality, there is no conversion of an Enhanced IGRP composite metric into a RTMP metric. Because a hop count is carried in an Enhanced IGRP metric tuple as the Enhanced IGRP route spreads through the network, 1 is added to the hop-count carried in the Enhanced IGRP metric blocks through the network and put into any RTMP routing tuple generated.

There is no conversion of an Enhanced IGRP metric back into an RTMP metric because, in reality, what RTMP uses as a metric (the hop count) is carried along the Enhanced IGRP metric all the way through the network. This is true of Enhanced IGRP-derived routes and routes propagated through the network that were originally derived from an RTMP route.

Summary

This case study illustrates the integration of Enhanced IGRP in graduated steps, starting at the periphery of the network before adding Enhanced IGRP into the core network. With Enhanced IGRP for IP networks, route summarization and redistribution of routing updates are key considerations. To add Enhanced IGRP to IPX networks, it is critical to configure RIP and SAP on interfaces connecting to Novell hosts or routers that do not support Enhanced IGRP. When adding Enhanced IGRP to AppleTalk networks, turn off RTMP on the interfaces that are configured to support Enhanced IGRP.

