7. RIP and SAP

RIP is the traditional routing protocol used with the IPX™ protocol. Like IPX itself, RIP is derived from the Xerox Network System specifications. SAP is a protocol used with IPX for a server to advertise availability of a service, and for clients to find that service (that is, to determine its IPX network address). Both use a broadcast mechanism. The two are usually implemented together.

As described in Section 2, the NLSP™ specification includes RIP/SAP as its end node/router protocol. End nodes include

a) Clients that want to find services and routes

b) Servers that want to advertise their existence using SAP, and to find routes

This part of RIP/SAP support is an indispensable part of NLSP operation. It must always be active for all circuits. It includes

- Responding to RIP requests for route information
- Responding to SAP requests for service information
- Absorbing SAP information about services that reside on the same system as the router

The router-to-router part of RIP/SAP support is also included in the NLSP design, so that customers can deploy NLSP routers in the same internetwork with RIP routers. It includes

- Absorbing RIP broadcasts from RIP routers
- Absorbing SAP broadcasts transmitted by servers
- Generating RIP and SAP broadcasts for consumption by RIP routers and NetWare® servers

The compatibility part of RIP support is operational conditionally, on a per-circuit basis. The end user chooses between three alternatives, for each router attached to each link:

a) "On"..... cause RIP packets to be generated and absorbed on the link.

b) "Off"..... inhibit RIP from being generated or absorbed on the link.

c) "Auto".. If a RIP router is detected on the link, generate and absorb RIP packets; otherwise, don't.

The default choice is "Auto."

The basic idea of "Auto" mode is

- Run RIP on a circuit where there are non-NLSP routers, but
- Deactivate RIP once all routers are running NLSP, except to respond to client (workstation) requests.

If an "Auto" circuit has RIP deactivated, receipt of certain RIP packets reactivates RIP (see page 7-14). In the absence of these packets, RIP is deactivated after (ripUpdate × ripAgeMultiplier).

Likewise, the compatibility part of SAP support is operational only conditionally, on a per-circuit basis. Its activation or deactivation is independent of RIP's. The end user chooses between the same three alternatives, with the same "Auto" default. If an "Auto" circuit has SAP

deactivated, receipt of certain SAP packets reactivates SAP. In the absence of these packets, SAP is deactivated after (sapUpdate × sapAgeMultiplier).

The remainder of this chapter should be read in conjunction with Reference [Nov92].

In the diagrams of this section, square-shaped symbols represent RIP routers. Diamond-shaped symbols continue to represent NLSP routers, including their RIP/SAP-emulation roles.

7.1. Maintaining RIP and SAP Information

NLSP routers keep track of specific information to fulfill their RIP/SAP compatibility requirements.

7.1.1. XRoutes and Services Defined

RIP information consists of *external routes*, abbreviated *XRoutes* in this chapter. Each XRoute summarizes the distance to a particular IPX network from the router reporting the XRoute. Specifically, each XRoute contains the following values:

- Network Number: The IPX network number to which this XRoute refers.
- Hops: The number of hops to reach that network. Each router traversed is counted as one hop.
- RIP Delay, or Ticks: The number of RIP timer ticks to reach that network. There are 18.21 ticks per second. The value is the time to deliver a 576-byte packet one way.

Each RIP packet sent or received conveys information about routes. When an NLSP router receives a RIP packet, the route information from the packet is absorbed into the Link State database as XRoutes, subject to the procedures specified in this chapter. When an NLSP router sends a RIP packet, it reports not only XRoutes, but also routes reachable by NLSP: the network numbers of routers and pseudonodes in the Link State database.

SAP information consists of *Services*. Each Service identifies a potentially reachable application service that is accessible using the IPX Network layer protocol. Each Service record contains the following values:

- Service Name: The name of the service. This is a text string up to 47 bytes. It identifies the service uniquely in the internetwork. Routers need not be concerned with the naming conventions.
- Service Type: The type of service offered. Novell assigns these values; for example, NetWare servers use the value 4. See Reference [Nov92] for a (partial) list of defined types and the contact point for registering new service types. Routers need not be concerned with the semantics of the Service Type value.
- IPX Network-Layer Address. The address at which the service resides.
- Hops: The number of hops to reach that service. Each router traversed is counted as one hop.

The Name/Type combination identifies a service. There can be several distinct Services having the same Name if their Types are different. If a Service appears with the same Name/Type combination as one previously known but at a different Address, it is an indication that the Service has changed location.

Each SAP packet sent or received is, in essence, a list of zero or more Services. The router does not transmit a SAP reporting a Service unless that router has a route to the network number on which the service resides.

7.1.2. Relation to Link State Database of Receiving RIP/SAP

RIP routes are absorbed by NLSP routers as they receive RIP packets from other RIP routers on directly connected networks. Those RIP routes that are kept as XRoutes are included in the pseudonode LSPs for the network on which they were received.

On each LAN, one of the NLSP routers is the Designated Router, chosen by a dynamic election procedure. The Designated Router for the LAN maintains the pseudonode LSP. The Designated Router is responsible for flooding the pseudonode LSP to the other NLSP routers on behalf of that network segment. Only the Designated Router propagates received RIP information in LSPs. Other NLSP routers on the network also retain received RIP information, maintaining a *potential pseudonode*. The potential pseudonode is information that a router keeps in anticipation of possibly becoming the Designated Router and being required to produce an actual pseudonode.

For each WAN circuit connecting an NLSP router with a RIP router, the NLSP router has the Designated Router responsibilities.

A router can be maintaining more than one pseudonode (or potential pseudonode)—different ones for different directly connected circuits.

After determining that a RIP route received is not already reachable as an NLSP destination, the router must determine whether to keep that route as an XRoute in the pseudonode (or potential pseudonode) for that circuit. The router only absorbs an XRoute if it is the best XRoute (or tied for best) known in the area. The criteria for deciding the best is to consider ticks first. In case of a tie consider the hops. Only the ticks and hops reported by RIP are considered; the NLSP portion of the path is ignored for this purpose. If another router later advertises the same network number in a better XRoute, delete the local XRoute. Likewise, delete it if the network number becomes reachable as an NLSP route. A result of these rules is that any particular XRoute is advertised by, at most, one NLSP router in an area (except during a transitional period). Ties are an exception. If both ticks and hops are tied, there can be several routers in an area reporting the same XRoute.

XRoutes are not constrained by area addresses. When a router receiving a RIP packet absorbs a network number into a pseudonode (or potential pseudonode) LSP as an XRoute, it does not compare that network number with the router's own manualAreaAddresses. Neither does it compare the network number with other area addresses in the routing area.

Other NLSP routers—those not acting as Designated Routers—also gather the RIP information, even though they do not propagate the information in LSPs they originate. First, they do so to be prepared to become the Designated Router should the need arise by failure or resignation of the current Designated Router. Second, the RIP packets are needed for the immediate address to use when forwarding data packets destined for the network number reported in the RIP packet. The XRoute in the pseudonode LSP does not contain the immediate address, so information from directly received RIP packets is indispensable. (Data packets are not forwarded to the Designated Router to be forwarded over the same LAN again.) Consider Figure 7-1. Network numbers E, F, and G are all attached to the pseudonode for N, which is generated by Designated Router B. However, the pseudonode does not indicate which of the RIP routers (R or S) provides access to each of the three network numbers. So when A receives

data traffic from its WAN link destined for E, for example, the pseudonode itself does not contain enough information to forward that traffic. To supplement the pseudonode, A keeps information it receives by RIP over the LAN. Specifically, it records that R is the path to E, that S is the path to G, and whichever of R and S is the better path to F.

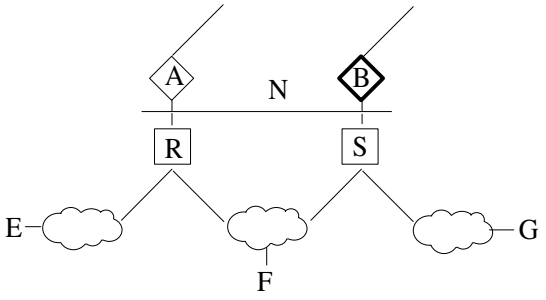

A and B are NLSP routers.
B is Designated Router for N.

R and S are RIP routers.
R reports a route to networks E and F.
S reports a route to networks F and G.

Figure 7-1: Routing between NLSP and RIP

When an NLSP router is connected by a WAN circuit to a RIP router, the NLSP router constructs a WAN pseudonode (numbered or unnumbered) for the circuit and acts as the WAN Designated Router for that pseudonode. There is no election procedure. RIP and SAP packets it receives over the circuit are advertised as XRoutes and Services in the pseudonode LSP.

Incoming RIP packets from NLSP routers are ignored. Those packets are presumed to have been sent for backward compatibility. An NLSP router is recognized by its data-link address matching that of a router for which an NLSP adjacency exists in either the “Up” or the “Initializing” state.

From the Link State database, each router builds an exactly matching Link State graph, which includes the XRoutes. Each XRoute has forwarding information (hops and ticks) relative to the network where its existence was discovered. For a WAN pseudonode, the Designated Router assigns the NLSP Delay to be 15 milliseconds and the NLSP Throughput is calculated from the measurements of the Timer Request/Response exchange in Section 3:

```
start_time = time when Timer Request packet is sent, 1/18 second

end_time = time that matching Timer Response packet is received, 1/18 second

Throughput = 8 × 1,000 × 576 /
                [ { 55 × MAX ( 1, end_time − start_time) / 2 } − 15 ]
```

SAP packets and Services are like RIP packets and XRoutes. Received SAP packets report network Services that are absorbed into Services Information options of pseudonode LSPs, flooded, and kept synchronized. The rules for absorbing SAPs are the same as for RIP: check the network number against manualAreaAddresses, and check that the sender is not an NLSP router. Each Service has forwarding information (hops) relative to the network where its existence was discovered.

A system that is Designated Router on two circuits need only accept SAP packets received on one of them. However, if a system connected to two circuits is Designated Router for the first but not the second, it must accept SAP packets received on the first. Otherwise, a SAP might fail to be included in any LSP.

7.1.3. Relation to Link State Database of Sending RIP/SAP

When sending RIP and SAP packets on a LAN, a router's behavior is the same whether or not it is the Designated Router for the LAN.

RIP routers include in the ticks value the cost of the network onto which they are broadcasting a RIP packet. This value is included in LSPs. This must be taken into account when reporting ticks in outgoing RIP packets: the value must be adjusted back to duplicating that portion in the report. For example, consider Figure 7-2. When RIP router C reports about LAN P by broadcasting a RIP packet on LAN M, it reports two ticks: one for P and one for M. The pseudonode LSP for M includes an XRoute for P showing two ticks. Then, when NLSP router B broadcasts a RIP on LAN N describing P, it reports three ticks, not four. The three includes provision for N itself.

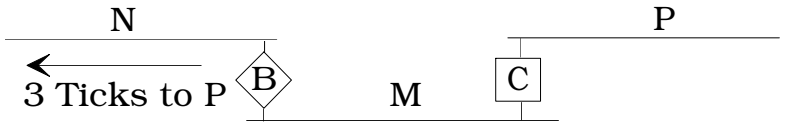

Figure 7-2: Adjusting the Ticks in RIP

There is an additional provision for the internal network numbers of routers (either RIP or NLSP). An additional tick is added to traverse the (conceptual) gap from the router entity within the system to the internal network number within the same system. There is no corresponding additional hop. An internal network number is considered to be 1 tick and 0 hops away from the router entity within the same system. In Figure 7-3, suppose NLSP router A has internal network number AA, NLSP router B has internal network number BB, and RIP router R has internal network number RR. When A sends a RIP on LAN N advertising AA, it reports 1 hop and 2 ticks. When it sends a RIP on LAN N advertising BB, it reports 2 hops and 3 ticks. The 3 are as follows: one for the network being transmitted onto, one for the AB link, and one additional tick from router B to its internal network BB. When A sends a RIP on LAN N advertising RR, it bases the ticks on the XRoute in the pseudonode representing LAN M. The tick value in the pseudonode is 2, and already contains the additional tick within R. So, A adds one for the AP link and one for N, making 4.

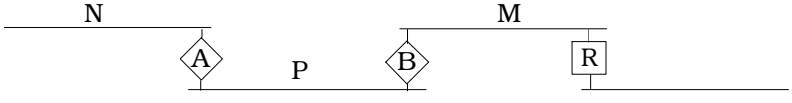

Figure 7-3: Additional Tick for Internal Networks

Counting hops is easier than dealing with ticks. The adjustments described for ticks do not apply. Each traversal of a router is one hop. The total hop count is reported in both RIP and SAP broadcasts. For example, when sending a SAP onto network N, router A reports 1 hop to a Service residing on its internal network AA, two hops to a Service on BB, three hops to a Service on network M, and three hops to a Service on network RR. In the latter case, the

pseudonode for M shows one hop to a Service on RR, because R sends a SAP for the Service having hop count 1.

Each NLSP router continues to receive RIP/SAP broadcasts and to keep the information current in its local database, in its potential pseudonode. However, RIP/SAP broadcasts sent by the NLSP router are based on the XRoute information in the Link State graph built by the Decision Process. This means that RIP/SAP broadcasts by NLSP routers reflect only the Designated Router's view of that particular LAN. The broadcasts change in content only when the Designated Router detects a change. (By the Split Horizon rule described later, the broadcasts are being sent on a different LAN than the one where the information was learned.) In Figure 7-3, for example, suppose A is the Designated Router for LAN P. When B sends RIP/SAP information about P onto M, it uses information from the pseudonode built by A representing P; it does not use information from RIP/SAP packets B itself received over P (even though it absorbs that information for other purposes).

7.1.4. XRoutes, NLSP Routes, Services, and the Decision Process

The Decision Process constructs a Link State graph of the routing area. An *NLSP route* is a path in the Link State graph using only network numbers that appear in the Management Information field of LSPs. As will be seen in the following discussion, NLSP routes are preferred over other routes (which include External Route fields of LSPs).

The Decision Process determines the best path for inclusion in the Forwarding database, based only on the ticks and hops reported in XRoutes known to the router. This means that only the portion of a packet's path in the RIP domain is taken into account. The NLSP portion of the path is ignored when choosing the best route. Consider Figure 7-4. Generally, only one route from C to E is in the Link State database (except for ties). In a transient situation, though, the pseudonodes for M and N might both contain XRoutes for E. When C determines its route to E, it compares

- The ticks and hops in the pseudonode for M

with

- The ticks and hops for E in the pseudonode for N.

Consider the ticks first. In case of a tie, consider the hops. Neither the path from C to A nor the path from C to B enters the decision to this point.

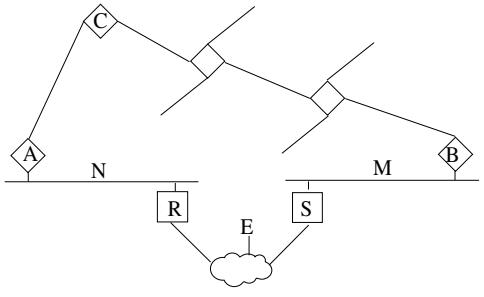

A, B, and C are NLSP routers.

R and S are RIP routers.
Both report paths to network E..

The pseudonodes of both network N
and network M report paths to E.

Figure 7-4: Multiple XRoutes for a Destination

If both ticks and hops are tied, break the tie by choosing the least NLSP cost. In Figure 7-4, this means comparing the NLSP path from C to A with the NLSP path from C to B. If there is

still a tie, apply the order of preference of routes listed in section 6, with (optional) load-splitting.

Likewise, the router discovering SAP information determines the smallest hop count to that Service, for each network address at which it is accessible, for inclusion in the pseudonode (or potential pseudonode). As with RIP, a Service appears only once in an area, except for ties and for transient conditions.

When XRoute or Service information changes, the link state method propagates the information throughout the routing area. This implies updating the Forwarding database of each router accordingly. However, the computational burden is comparatively small. The Decision Process actually consists of two parts.

a) Run Dijkstra's algorithm to build the Link State graph and spanning tree. This includes determining the costs to every NLSP destination.

b) Determine the best costs of XRoutes, based on the spanning tree and the RIP costs reported at nodes of the tree.

XRoutes and Services are like leaf nodes connected to various LSP nodes of the Link State graph. So, when an XRoute or Service appears, disappears, or changes in ticks or hops, there is no structural change to the graph. A router performs only part b). Because there are no structural changes to the graph, there is no need to run Dijkstra's algorithm.

7.1.5. Building a RIP Route from the Link State Graph

When an NLSP router transmits a RIP packet, it must insert values for ticks and hops. It combines values from the NLSP area with values appearing in the XRoute.

As part of the Decision Process, each router maintains hop count and RIP Delay values for each Link State graph node's *upstream NLSP route*; that is, the route from itself to that node. Call this the *Upstream RIP Delay*. The value is the sum of individual RIP Delay values of the links the route traverses. For each link, the individual RIP Delay is calculated from the NLSP Delay and Throughput values. For a LAN, the calculation is

```
MAX [ 1, ( 576 × 8 × 18  / Throughput) + (Delay × 18  / 1,000,000 )]
```

For a WAN, the calculation is

```
MAX [ 1, ( 576 × 64 × 18 / Throughput) + (2 × Delay × 18 / 1,000,000 ) ]
```

Note: When receiving LSPs, calculate the ticks immediately from throughput and delay. Do not wait until the Decision Process. The chances are that the Decision Process is run several times before that particular LSP (with its throughput and delay) is received again. You might as well calculate only once and save the result. There is less computing burden overall, especially because multiplication and division are involved. Also, you want to limit the duration of the Decision Process, because input of new LSPs is deferred during that interval.

When summing the individual RIP Delay values, special attention is needed when load-splitting is in force. It can happen that several paths have equal NLSP costs—so load-splitting is applied—but with different RIP Delay values. In this case, apply the highest of the several RIP Delay values to the path. Figure 7-5 shows an example.

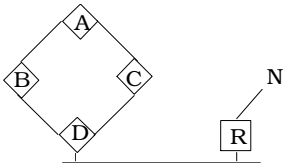

NLSP router A is load-splitting between B and C to reach D.

That is, AB+BD = AC+CD, with values in NLSP's cost.

The RIP Delay from A to D is MAX(AB+BD, AC+CD), where this time the values are the individual RIP Delays.

Figure 7-5: RIP Delay and Load-splitting

XRoutes maintain hop count and ticks that are relative to the LSP to which they are attached.

Using

a) The upstream RIP delays (defined previously) stored in the graph nodes

b) The values reported in XRoutes

this router can generate RIP route values to all reachable IPX networks, as follows:

- If a network is an internal route (that is, it appears in the graph as a non-XRoute), then
  a) Its RIP ticks value is the upstream RIP Delay
  b) Its RIP hop count is the number of links traversed by the downstream route (excluding WAN pseudonodes)
- If a network is an XRoute, then one starts with the downstream hop/tick values (as above) to the pseudonode where it is reported, and adds the respective hop/tick values reported in the External Routes field.

Likewise, the hops from this router to a Service is computed by adding

a) The number reported in the Services Information field

To

b) The number of graph links (excluding pseudonodes) to the pseudonode whose LSP includes the Services Information.

7.1.6. Aging XRoutes and Services

XRoutes and Services are aged in a way similar to the one RIP routers use to age entries.

An aging timer is applied to each XRoute and Service in the local potential and actual pseudonode LSPs maintained by this router. When a RIP Response arrives, each route reported has its timer reinitialized if it is already represented by an XRoute. Likewise for SAP and Services.

The interval at which RIP is activated for its periodic broadcast is ripUpdate. For SAP, the corresponding value is sapUpdate.

The aging timer after which a stale XRoute is purged from its pseudonode is ripAgeMultiplier × ripUpdate. The aging timer after which a stale Service is purged from its pseudonode is sapAgeMultiplier × sapUpdate.

If the timer expires, the XRoute or Service is removed from the LSP. If an XRoute is removed from a potential pseudonode LSP, all Services residing on that network number are also

removed from that potential pseudonode LSP. If this router is the Designated Router for that LSP, this change in the pseudonode LSP results in removal of the XRoute from the Link State graph. This causes a triggered update to reflect the changes to the XRoutes. As with other triggered RIP/SAP updates, this causes routers to update their respective Forwarding databases, but does not require Dijkstra's algorithm to be run.

7.2. Generating Periodic Updates

If a router (or link) were to stop operation without warning, a way is required for other routers to discover the loss of connectivity. This is accomplished by periodically broadcasting current RIP/SAP information and aging any database entries not refreshed by received broadcasts.

All NLSP routers generate RIP periodic update broadcasts every ripUpdate for its connected links supporting RIP compatibility. The default is 60 seconds, but the interval is configurable on a per-link basis.

> **Note:** The ability to configure the periodic broadcast interval and the timeout to age out RIP routes and SAP entries is a recent design, not supported by all installed systems. Older versions of IPX RIP/SAP routers use 60-second periodic timers and a four-minute timeout. All routers on any one network segment (LAN or WAN) must be configured with the same values.

The network numbers to report in these broadcasts are all the IPX network numbers that are reachable—link-state destinations (pseudonodes, internal networks) and XRoutes alike. In RIP, reporting a hop count of 16 means that the network is unreachable. Only destinations whose total hop count is less than 16 are reported in the periodic updates. (This is not affected by a possible configured node limit greater than 16.) When sending a RIP update reporting 16 hops, the ticks value is unspecified.

If more than 50 routes are to be reported, multiple RIP packets are used, each with 50 entries or fewer. If more than seven Services are to be reported, multiple SAP packets are used, each with seven entries or fewer. These packet size limitations can be relaxed by an overriding user configuration on a per-circuit basis. The value ripPacketSize determines the size, in bytes, that a RIP packet is allowed to reach. The corresponding value for SAP is sapPacketSize.

7.3. Split Horizon

RIP implements a split horizon heuristic to cut down on sending and processing redundant RIP traffic. Split horizon applies equally whether a route is learned by RIP or NLSP. There are two parts:

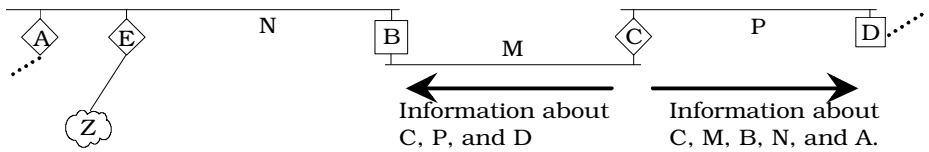

Figure 7-6: Split Horizon

a) A router about to broadcast onto a particular circuit, C, does not include any information about other networks for which another router on C has a better route. For example, in Figure 7-6, suppose B is a RIP router and C is an NLSP router. When router C broadcasts a RIP packet on LAN M, it does not include information received from B about LAN N. (Otherwise, a system on M might erroneously conclude that there are two paths to N.)

Moreover, if A is an NLSP router, it does not broadcast RIP packets onto N reporting routes in the network cloud Z (it does not matter whether E is a RIP or an NLSP router).

A way to implement this heuristic is to ask the following questions when preparing to broadcast information about network H onto circuit J: "Is J the best next hop to H? Is J one of the candidate equal-cost best hops to H?" If the answer to either of these questions is "yes," then refrain from broadcasting information about network H onto circuit J. In the equal-cost path case, it does not matter which of the paths are actually chosen for forwarding. Figure 7-7 illustrates this. In this example, the LANs have equal cost. When A runs the Decision Process, it notes that it can reach H through either B or C with equal cost. Suppose the maximal splitting degree is one, and A chooses B as the forwarding path to H. Nonetheless, A does not broadcast RIP information about H onto J, because C is one of the candidate equal-cost best hops from A to H.

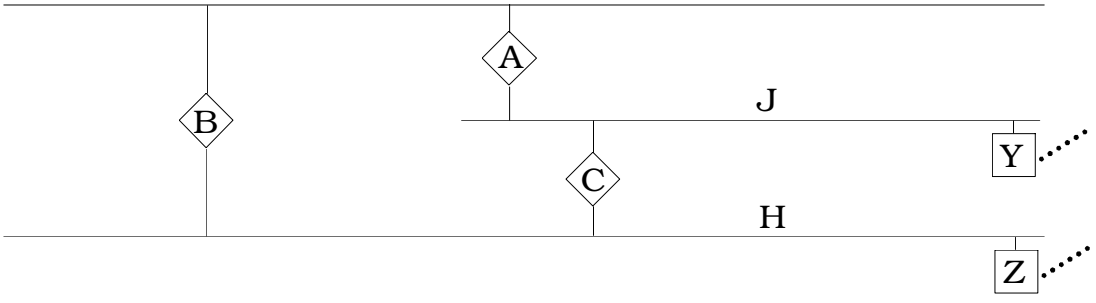

Figure 7-7: Split Horizon and Load-Splitting

b) A router about to broadcast on a particular segment does not include any information about the network on which the packet is to be sent. Router C does not include information about H in its RIP broadcast on H.

The RIP Delay values in outgoing RIP reports on a link include provision for that link itself. On the other hand, hops correspond to router traversals. For example, when C transmits a RIP on M describing P, the report indicates "1 hop, 2 ticks"—one tick transits M; another transits P.

SAP broadcasts also conform with the split horizon rule.

7.4. Generating Triggered Updates

When there are changes in the routes that have been broadcast as RIP routes, or when the router discovers changes to local links, or when internally provoked events occur (for example, activating or deactivating a circuit), triggered updates of RIP routes are broadcast on circuits supporting RIP compatibility.

The form and content of triggered updates are the same as those of periodic broadcasts.

7.4.1. Changes in XRoutes and Services

When an XRoute is added to or deleted from the database, a triggered update is generated to reflect the changes. One is also generated when an XRoute's forwarding information (hops, ticks) changes.

When a router detects that a route becomes unreachable, it generates a triggered update to convey this. The cause can be (a) the Decision Process finds a network to become unreachable, or (b) the hop count reaches 16 or more.

Likewise, SAP updates are generated when Services are added or deleted from the Link State database, or there is a change in hop count as reported in LSPs.

When the router detects that a network number becomes inaccessible, it marks Services on that network as inaccessible as well, and sends triggered SAP updates accordingly.

When a change is detected, there is a hold-down timer of 0.55 seconds before transmitting the resulting RIP or SAP packet. This allows multiple changes occurring close in time to be consolidated.

7.4.2. Changes in the Link State Graph

RIP triggered updates are generated when the following occurs:

- Nodes are added to the Link State graph
- Something changes in the Link State database that changes the Hops or Ticks (a link going up or down, a cost changing, or similar change)
- Any node's upstream route first hop changes, or its upstream hop count or IPX Delay changes
- Any node becomes unreachable

Only the new or changed routes are broadcast. Even if an XRoute does not itself change, the ticks and hops to it can change when there are changes to the Link State graph between this router and the pseudonode to which the XRoute is attached.

If, prior to some change, a node or XRoute was within the hop count range for RIP (15 hops), but after the change it was outside the RIP range, it is included in at least one update with its unreachable hop count, so that neighboring RIP routers delete that route.

Likewise, SAP triggered updates are generated when an LSP having Services is added to or deleted from the Link State database, or when a Link State graph change causes the hop count for a Service to cross the 16-hop threshold (in one direction or the other).

When the router detects that a network number becomes inaccessible, it marks Services on that network as inaccessible.

When a change is detected, there is a hold-down timer of 0.55 seconds before transmitting the resulting RIP or SAP packet. This allows multiple changes occurring close in time to be consolidated.

7.4.3. Circuit Activation and Deactivation

When a directly attached circuit becomes active, the router's actions depend on the configured mode for that circuit.

- If the RIP compatibility mode is "Off," there is no RIP-compatibility action to perform.
- If the RIP compatibility mode is "On,"
  a) Send the first NLSP Hello on the circuit.
  b) Transmit a complete RIP update on the circuit.

c) Start the timer that controls periodic broadcasts.

d) Broadcast a RIP "All Routes" request on the circuit.

- If the RIP compatibility mode is "Auto,"
  a) Send the first NLSP Hello on the circuit.
  b) Broadcast a RIP "All Routes" request on the circuit.
  c) If a RIP response is received from a non-NLSP router, the first time this happens
     i) Transmit a complete RIP update on the circuit.
     ii) Start the timer that controls periodic broadcasts.
     iii) Tag the circuit as having RIP activated.

Note: An NLSP router is recognized by its data-link address matching that of a router for which an NLSP adjacency exists in either the "Up" or the "Initializing" state.

The result of this is that with NLSP active and RIP in "Auto" mode, RIP periodic responses and triggered updates are sent only if there are other non-NLSP routers on a LAN. With RIP compatibility mode "On," RIP packets are always sent.

The same procedure is used for SAP, depending on the SAP support configuration choice for the circuit.

It is possible that another NLSP router responds to the RIP request and this router receives the response before it recognizes that the other router is an NLSP router. In this case, the adjacency is formed later, and the other router's LSPs begin to arrive. The RIP routes previously learned from that router duplicate the information in the now expanding Link State database. To prevent having that duplicate information in pseudonode LSPs, the router removes any XRoutes that were absorbed from RIP packets from that router. If the router is in "Auto" mode for that link, and if the RIPs just described had been the ones that activated RIP compatibility for the circuit, the RIP compatibility is deactivated when the XRoutes are removed.

When a directly attached circuit supporting SAP compatibility becomes active, the router broadcasts a SAP general request to discover Services accessible over the circuit. It then starts the timer for generating periodic updates. As soon as any SAP packet is received from a non-NLSP router, it acts as though the periodic interval has expired and broadcasts the Service information contained in its database (with split horizon applied).

When a directly attached circuit is about to become inactive, the router sends RIP/SAP updates on that circuit indicating that destinations reached through the router through that link are now unreachable. (There is no hold-down timer applied for this part, and no interpacket gap is required.) It then generates triggered RIP/SAP updates to all the other RIP/SAP-support circuits about the routes/Services that have become unreachable as a result of the deletion.

7.4.4. Router Activation and Deactivation

When a router begins operation,

a) It determines the network numbers of its directly attached RIP-support links.

b) If the system itself contains application Services, it records the Service names as Services in its database.

c) It performs the circuit activation procedures described on page 7-11 for each circuit.

When router operation is about to terminate, the system performs the circuit deactivation procedures described on page 7-12 for each circuit.

7.5. Receiving RIP and SAP Packets

When a valid RIP packet arrives on circuit *C*, the router takes the following steps:

a) If *C* is a LAN, and the packet's source MAC data-link address matches that of a router for which an NLSP adjacency exists, the packet is ignored.

> **Note:** The other router is an NLSP router doing RIP compatibility. Ignoring the packet prevents this router from having to process RIP routes for destinations it already knows about through LSPs.

b) If the packet is a RIP Request,

i) If the request is for a specific route (or routes), the response is returned to the requester's address. Information is returned no matter how large the hop count is.

> **Note:** This allows end nodes to access the larger networks possible through NLSP. With RIP/SAP alone, without NLSP, the maximum supported network diameter is 15 hops.

The split horizon rule is applied. In other words, the router responds with any information typically sent in a periodic broadcast packet that would be applicable to the request.

If the router does not have any record of the requested network number,

i.1) If the network number is outside the local area, and if an Exit Router (offering a path to other routing areas) is known to this router, the router responds to the request. The hops and ticks reported in the response are those of the path to the Exit Router.

i.2) Otherwise, the router does not reply regarding the requested network number.

One exception to the split horizon rule occurs if a specific RIP request is received on a network segment for the network number of the segment itself. The router responds with the requested information. Even though the information is not useful for routing, it allows a system to learn the node addresses of all routers on a local network segment.

ii) If the request is a general request (for "All Routes"), the response is returned to the requester's address only if the circuit is configured to be in "On" or "Auto" mode. If the mode is "Off," a general request is ignored. Responding to a general request is the same as generating a periodic update. Entries that would have a total reported hop count of 16 or greater are not included. This prevents other RIP routers from having to process entries that they automatically discard. Split horizon is applied.

c) If the packet is a RIP Response,

i) If the circuit has RIP compatibility configured to be "Off," the packet is ignored.

ii) The source and destination network numbers in the IPX header are compared. If they are not the same, the packet is ignored.

iii) If RIP is in "Auto" mode and deactivated, it is reactivated. If SAP is in "Auto" mode and deactivated, it is reactivated.

iv) For each route in the RIP Response, the Link State database is searched.

- If the destination network is reachable by an NLSP route (that is, by a path in the Link State graph using only network numbers that appear in the Management Information field of LSPs), the RIP route is ignored.
- If the network number is not within the ranges delineated by any one of the router's own manualAreaAddresses, the RIP route is ignored.
- If the destination is not a Link State destination, this RIP route information is searched for in the local existing database. That is, if this is the Designated Router for *C*, it is searched for in the pseudonode LSP; if it is not the Designated Router, it is searched for in the potential pseudonode.
  - If the route is absent, and this is the best RIP route known in the area (or tied for best), it is added.
  - If it is present from the same RIP router, changes in value are recorded and the timeout for the XRoute is refreshed.
  - If it is present from a different RIP router, and if the route received is better than the one in the database, the old information is replaced by the new.
  - If this is the Designated Router, any change is reflected in the pseudonode LSP, triggering an LSP regeneration.

Any time a received RIP packet indicates a hop count of 16, it means that the network number cannot be reached through the router sending the packet. If this is the only path to that network number, it (the network number) is removed from the pseudonode (or potential pseudonode) and a triggered update results.

The logic to process arriving SAP packets is the same as for RIP, with a few additional considerations.

As with RIP XRoutes, SAP Services discovered by SAPs arriving over directly attached circuits are maintained as part of the corresponding pseudonode LSP (or potential pseudonode). If this router is the Designated Router for that pseudonode, it (the router) floods the pseudonode LSP throughout the routing area. As with RIP, SAP packets arriving from another NLSP router (performing SAP compatibility) are ignored.

In the SAP counterpart of step c) part iii), it is not enough that a SAP response has arrived. Before reactivating RIP and SAP for an "Auto" circuit, check that the SAP packet contains at least one entry of type "File Server" (0x04). In the absence of such an entry, do not reactivate RIP and SAP. The purpose of the additional check is to avoid activating RIP and SAP just because there is, for example, a print server on the LAN, advertising its Service as reachable on the LAN itself.

If SAP is in "Off" mode for a LAN circuit, Services advertised as reachable on the LAN itself are absorbed into the locally generated pseudonode LSP for the circuit. The same is true if SAP is

in “Auto” mode and deactivated. The purpose of this provision is to be able to run in SAP “Off” mode and still advertise systems (such as the print server in the preceding paragraph) directly reachable on that LAN.

When receiving a SAP packet advertising that a Service is now reachable at a new address, a router accepts this Service and deletes the old Service if the new address is nearer than (or at the same distance as) the old address. (This is to deal with the situation where, for example, a server is brought up with a new internal network number.) To decide that the SAP describes the same Service, compare both the Service name and the Service type. SAPs with the same name but different types represent different Services.

As with RIP, each SAP Service is absorbed once, at most (except for ties). Each Service is represented in, at most, one locally generated LSP, even if its SAP is received on multiple circuits.

The router is responsible not only for absorbing the SAP packets as they are received, but also for representing “local” SAP Services detected within the router’s own system. Services advertised on the router’s own internal network are always accepted, and those Services are added to the local LSP, the one describing this NLSP router.

When responding to a “Get Nearest Server” SAP request, a system considers all known Services of the requested type, determines the IPX address of each, and responds with the one that is closest to itself. The closeness criterion is ticks, applying the same method as in building a RIP route. If there is a tie for best, hops are used to break the tie.

7.6. Maintaining a Proper Interpacket Gap

Certain nodes can have problems properly processing a sequence of RIP packets arriving too close together. Currently, the minimum interpacket gap time allowed for transmitting RIP packets is 55 milliseconds.

Similarly, the minimum interpacket gap time allowed for transmitting SAP packets is currently 55 milliseconds.

7.7. RIP and SAP Filters

An NLSP router can be configured on a per-link basis to filter incoming and outgoing reports of RIP routes (by network number or a pattern of numbers) and Services (by Service name or name pattern). The end user is responsible for maintaining uniformity in the configured filters. In Figure 7-8, for example, if A has different RIP/SAP filters on the LAN than B does, router C perceives radical changes in what is accessible through R if the Designated Router status changes on the LAN.

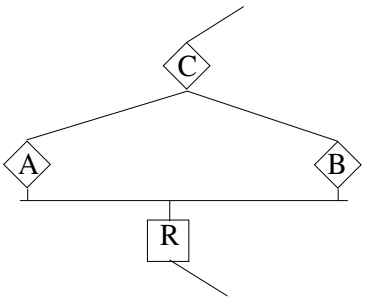

Figure 7-8: Filter Uniformity

Filtering applies to the transmission and absorption of information by the RIP/SAP protocols. It does not apply to exchange of XRoute and Service information in LSPs. Once information about a route or Service is in one router's LSP, that information is propagated throughout the routing area. This means that users can control the amount of RIP/SAP traffic, but cannot necessarily prevent RIP/SAP information from reaching another region of an internetwork, if NLSP routers provide connectivity to that region.

7.8. RIP/SAP Database

7.8.1. Configured Values per Circuit

ripState
: Mode of support for RIP on the circuit: "On," "Off," or "Auto." Default is "Auto."

ripUpdate
: The RIP periodic update interval, in seconds. Default is 60 seconds.

ripAgeMultiplier
: After this many ripUpdate intervals, unrefreshed RIP information is considered expired. Default is four; it ages out after three or four minutes.

ripPacketSize
: The maximum RIP packet size used on this circuit. Default is 432 bytes (header plus 50 entries).

sapState
: Mode of support for SAP on the circuit: "On," "Off," or "Auto." Default is "Auto."

sapUpdate
: The SAP periodic update interval, in seconds. Default is 60 seconds.

sapAgeMultiplier
: After this many sapUpdate intervals, unrefreshed SAP information is considered expired. Default is four; it ages out after three or four minutes.

sapPacketSize
: The maximum SAP packet size used on this circuit. Default is 480 bytes (IPX header plus seven entries).