A Routing Agent for TCP/IP: RFC 1058 Implemented for the KA9Q Internet Protocol Package

Albert G. Broscius, N3FCT

University of Pennsylvania
200 S. 33rd Street
Philadelphia, PA 19104-6319

ABSTRACT

The KA9Q Internet Protocol Package has introduced full TCP/IP internetworking to the Amateur packet community yet, until now, automatic routing has not been available within the package. This paper describes an implementation for the KA9Q package of the DDN Internet standard Routing Information Protocol as specified in RFC 1058. Proper usage and configuration of this routing code are explained.

Introduction

Routing in the Amateur Internet has been done so far on an ad-hoc basis with manual table manipulation for routes to both single hosts and to host clusters. The steady-state routing solution will see KA9Q interchanging routing information in a variety of forms such as RIP[1], EGP[2], HELLO[3], IGRP[4], and RSPF[5]. With this array of routing protocols available for use the Ampmet may interoperate effectively with commercial Internet gateway implementations[6].

In this context, RIP is the first of the Internet standard routing protocols to be implemented in the KA9Q package because it may easily interface a packet radio (PR) network to 4.3 BSD Unix-based networks which also use RIP. Guidelines for its use in this application are outlined briefly below in the Configuration section. As the first routing protocol available in KA9Q, RIP will most likely also see use as an inter-PR routing protocol even though its suitability for this purpose is questionable and its use discouraged.

The RIP Protocol

The Routing Information Protocol, as specified in RFC 1058, is a distance vector routing algorithm based on the 4.3 BSD Unix program, "routed"[1]. The basic distance vector routing algorithm consists of a periodic transmission by each IP switch of its internal routing table on the several attached network interfaces. A switch's routing table reflects the costs of sending packets to a set of destinations; each cost is termed a "metric" from the jargon of higher mathematics. The table which is broadcast to a switch's neighbors has its metrics incremented to account for the extra hop that packets would need to take to use the route. Upon reception of a neighbor's transmitted routing table, the metrics of the routes to each of the destinations listed is compared to those now held in the switch's routing table. Those destinations which can be routed through the neighbor more cheaply (lower metric) than the current route are entered into the switch's table and the current route is dropped. After some number of iterations the algorithm converges on a set of routing tables for the participating switches.

RFC 1058 RIP Modifications

The "routed" protocol has been augmented with two policies designed to speed the algorithm's convergence in response to changes in topology. The first of these policies, Split Horizon with Poisoned Reverse (SHPR), attempts to prevent two party routing loops from forming when a destination known to both switches becomes unreachable. SHPR assumes that the subnetworks connected to a switch are themselves fully connected. This implies that any packet arriving at a given network interface should not be forwarded through that interface since this would be a waste of network bandwidth; the proper route would not need to involve the forwarding node but would take the packet directly across the subnet. Since UI frame AX.25 subnets obviously do not qualify as fully connected, this feature of the protocol should be disabled on these subnet interfaces. However, the assumption of full subnetwork connectivity is valid for SLIP, Ethernet, and (ideally) Net/Rom subnetwork interfaces so SHPR should be enabled in these cases. To accomplish SHPR, an IP switch discriminates in its route update messages, telling its neighbors on SHPR subnetworks that all routes currently using this interface are inaccessible through this IP switch. This inaccessibility translates to an infinity metric for all such routes (in this case, infinity is defined to be 16). In this way, IP switch X on a SHPR subnetwork which loses adjacency to a destination will not attempt to route to that destination through neighbor switches whose only route is through switch X.

The second new policy is called Triggered Updates. When an IP switch makes a RIP-induced change to its routing tables, it sends a route update with only the affected routes very soon afterward to all of its neighbors. This update is unconditional and must be performed regardless of the status of the pending periodic route update. A small, random delay is inserted before the triggered update is performed to prevent a small network pileup. Effectively, these triggered updates force rapid convergence of the various IP switches' routing tables as the cascade of updates only traverses the tree of nodes dependent on the original affected route. This protocol feature may also be disabled on a per-interface basis if desired, Note that it is recommended that this feature be enabled for all subnetworks known at this time.

Optional RIP Additions Implemented for the KA9Q Package

Self Announcement

In the land-line Internet the main use of RIP is the announcement of subnetwork reachability by gateways attached directly to those subnetworks. Each subnetwork consists uniformly of hosts and gateways having a subnetwork IP address with a fixed prefix of a given length. Gateways belonging to several subnetworks are given an IP address for each of their attached interfaces. In this system, routes to individual hosts are the exception and subnetwork routes predominate in the routing tables.

Where the wired intemet has a clearly defined topology, the Amateur PR internet (amprnet) has few limitations on, and correspondingly few expectations of, the realized network topology.. After much discussion [7], it was decided that the amprnet IP address assignment should not reflect any correlation between IP address and subnetwork medium/frequency/modulation technique/etc. The IP addresses should reflect only the region-of-issuance of the station address as assigned by the regional IP address coordinator [8]. This flat IP address space approach has been used successfully in the NSFNET Backbone implemented with the Distributed Computer Network (DCN) architecture [9].

Since the subnetwork grouping of the land-line Internet may not be used in the amprnet, the RIP processes have the option of advertising a route to their own IP address. This option is configured on a per-interface basis. It is suggested that any Amateur constructing a personal, wired IP subnetwork obtain a Class C network address from the DDN management at SRI-NIC[lO] to lessen the routing space requirements of the amprnet. Since all of the hosts on the wired subnet will be routed as a single entity, less network overhead is needed for this group as a Class C network.

Private Routes

Provision has also been made for the construction of private routes which are used as normal by the IP router module but are not advertised to neighbors in route update messages. These routes could be used for intentional partitions of the network where organizational boundaries necessitate it. For example, a university-to-PR gateway may know the route to net 10, the ARPAnet, but it may not wish to tell other PR hosts that this route exists. It then uses a private route to facilitate its own communication without disclosing to others the route's existence.

Neighbor Refusal

Occasionally, there arises undesirable one-way or unpredictable propagation paths which introduce routes of dubious value into the network. These routes may be suppressed with the judicious use of neighbor refusals. A neighbor refusal list is maintained which is compared against incoming RIP datagrams. If a match is found the datagram and its contents are discarded.

KA9Q Package RIP Configuration

For each network interface which is to be used to transmit route update messages an ARP entry must exist in the appropriate table. In the case of subnetwork broadcasts, this entry must be installed manually in the autoexec.net file to associate the hardware broadcast address with the IP subnetwork broadcast address. For example, to participate in routing on the local Ethernet here at UPenn I must add:

```
arp add 128.91.0.0 ether ff:ff:ff:ff:ff:ff
```

to enable subnet broadcasts onto the Ethernet and also

```
route addprivate 128.91/16 ec0
```

to have my IP switch route local network traffic onto the Ethernet.

To start the routing agent itself, the following command is also added:

```
rip init 128.91 .O.O
```

which opens the RIP UDP socket locally to listen for route updates. It also broadcasts a request for the default router on this network. If no default router is available/desired, this command may be abbreviated to

```
rip init
```

which will send no broadcast request. At this point the routing agent is still silent, only listening for route updates from the routers conversing on the local network. To enable the routing agent to participate in the conversation, the “rip add" command is used. The general form of this command is:

```
rip add <destination-address> <iface> <interval> <flags>
```

where <destination-address> is the broadcast or host address of the neighbor which should receive route updates, <iface> is the interface over which we will receive route updates from this neighbor, <interval> is the amount of time in seonds we should wait between route update messages to this peer, and <flags> is a set of three bits which enable/disable SHPR, Triggered Updates, and Self Advertisement (from msb to Isb). The configuration here at UPenn has a rip add of the following form:

```
rip add 128.91.0.0 ec0 30 6
```

which causes the rip agent to output route update messages every thirty seconds (as per RFC 1058) to the local network broadcast address and to assume that route update messages received by interface ec0 have been generated by a thirty second interval rip agent also. The flags, decimal 6 = binary 110, designate that this interface Will perform SHPR, Will perform Triggered Updates, and Will Not perform Self Advertisement.

To ignore route update messages from certain other routers, the “rip addrefuse” command is used. If there were a misbehaving local router which I wished to suppress, say for example 128.91.254.34, I would place a rip addrefuse 128.91.254.34

in my autoexec.net to force the routing agent to drop all route update messages from 128.91.254.34 which may arrive.

References

1. Hedrick, C., “Routing Information Protocol - RFC 1058'', SRI Network Information Center, 1988.
2. Mills, D., “Exterior Gateway Protocol Formal Specification - RFC 904'', SRI Network Information Center, 1984.
3. Mills, D., “Distributed Computer Network Protocols - RFC 891'', SRI Network Information Center, 1983.
4. Reynolds, J. ed., “Assigned Numbers - RFC 1010'', SRI Network Information Center, 1988.
5. Goldstein, F., “Radio Shortest Path First (RSPF) Routing Protocol Specification"', as yet unpublished, 1988
6. Postel, J. ed., “Requirements for Internet Gateways - RFC 1021'', SRI Network Information Center, 1987.
7. Tcp-Group electronic mail conference, Discussion on meaning of IP address assignment in network topology, 1987.
8. Each metropolitan area or foreign country with IP activity has been assigned an Amprnet address block. Station IP addresses are to be assigned by a coordinator in each region. The list of coordinators is available from:

   Brian Kantor, wb6cyt
   University of California, San Diego
   Academic Network Operations Group
   Mail Code B-028
   La Jolla, CA 92093

   brian@ucsd.edu

9. Mills, D. & Van Braun, W. “The NSFNET Backbone Network”, Proc. ACM SIGCOMM, Stoweflake, VT, August 1987.
10. Application for an independent Internet network number can be made to:

   Hostmaster
   DDN Network Information Center
   SRI International
   333 Ravenswood Ave
   Menlo Park, CA 94025
   1-(800)-235-3 155

   HOSTMASTER@SRI-NIC.ARPA