

RIP Version 2 Protocol Applicability Statement

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

As required by Routing Protocol Criteria (RFC 1264), this report defines the applicability of the RIP-2 protocol within the Internet. This report is a prerequisite to advancing RIP-2 on the standards track.

1. Protocol Documents

The RIP-2 protocol analysis is documented in RFC 1721 [1].

The RIP-2 protocol description is defined in RFC 1723 [2]. This memo obsoletes RFC 1388, which specifies an update to the "Routing Information Protocol" RFC 1058 (STD 34).

The RIP-2 MIB description is defined in RFC 1724 [3]. This memo will obsolete RFC 1389.

2. Introduction

This report describes how RIP-2 may be useful within the Internet. In essence, the environments in which RIP-2 is the IGP of choice is a superset of the environments in which RIP-1, as defined in RFC 1058 [1], has traditionally been used. It is important to remember that RIP-2 is an extension to RIP-1; RIP-2 is not a new protocol. Thus, the operational aspects of distance-vector routing protocols, and RIP-1 in particular, within an autonomous system are well understood.

It should be noted that RIP-2 is not intended to be a substitute for OSPF in large autonomous systems; the restrictions on AS diameter and complexity which applied to RIP-1 also apply to RIP-2. Rather, RIP-2 allows the smaller, simpler, distance-vector protocol to be used in environments which require authentication or the use of variable

length subnet masks, but are not of a size or complexity which require the use of the larger, more complex, link-state protocol.

The remainder of this report describes how each of the extensions to RIP-1 may be used to increase the overall usefulness of RIP-2.

3. Extension Applicability

3.1 Subnet Masks


The original impetus behind the creation of RIP-2 was the desire to include subnet masks in the routing information exchanged by RIP. This was needed because subnetting was not defined when RIP was first created. As long as the subnet mask was fixed for a network, and well known by all the nodes on that network, a heuristic could be used to determine if a route was a subnet route or a host route. With the advent of variable length subnetting, CIDR, and supernetting, it was no longer possible for a heuristic to reasonably distinguish between network, subnet, and host routes.

By using the 32-bit field immediately following the IP address in a RIP routing entry, it became possible to positively identify a route's type. In fact, one could go so far as to say that the inclusion of the subnet mask effectively creates a 64-bit address which eliminates the network, subnet, host distinction.

Therefore, the inclusion of subnet masks in RIP-2 allows it to be used in an AS which requires precise knowledge of the subnet mask for a given route, but does not otherwise require OSPF.

3.2. Next Hop


The purpose of the Next Hop field is to eliminate packets being routed through extra hops in the system. It is particularly useful when RIP is not being run on all of the routers on a network. Consider the following example topology:



The Internal Routers (IR1 and IR2) are only running RIP-2. The External Routers (XR1 and XR2) are both running BGP, for example; however, only XR1 is running BGP and RIP-2. Since XR2 is not running RIP-2, the IRs will not know of its existence and will never use it

as a next hop, even if it is a better next hop than XR1. Of course, XR1 knows this and can indicate, via the Next Hop field, that XR2 is the better next hop for some routes.

Another use for Next Hop has also been found. Consider the following example topology:



The three links between the Central Office Router (COR) and the Remote Office routers (RO1 and RO2) are all Dial-On-Demand (DOD) links. The link between RO2 and R is a fixed link. Once all of the routers have been initialized, the only routes they know about are the configured static routes for the DOD links. Assume that connections between COR and RO1, and COR and RO2 are established and RIP information is passing between the routers. RO1 will ignore COR's route to RO2 because it already has a better one; however, it will learn to reach R via COR.

If we assume that RO1 and RO2 are only capable of establishing one link at a time, then RO1 will not be able to reach RO2; however, RO1 will be able to reach R. Worse still, if we assume that traffic stops and the DOD links drop due to inactivity, an attempt by RO1 to reach R will trigger the dialing of two links (through COR). Of course, once RO1 establishes a link to RO2, the problem corrects itself because the new route to R is one hop shorter.

To correct this problem, the routers may use the Next Hop field to indicate their next hop. Consider the following route advertisements during the period described above (before the RO1/RO2 link has ever been established):

Sender	Recvr	Route	NextHop	Metric
RO2	COR	R	0	1
COR	RO1	RO2	0	1
		R	RO2	2

When R01 receives the two routes from COR, it will ignore the route for R02, as mentioned above. However, since R is not in R01's routing table, it will add it using a next hop of R02 (because R02 is directly connected, after a fashion). Note that COR does count itself in R's metric; this is less than accurate, but entirely safe and correctable (when the R01/R02 link comes up). Suppose, now, that the R01/R02 link did not exist. R01 would ignore the specification of R02 as the next hop to R and use COR, as it would if no Next Hop had been specified.

Note that this is not a recursive algorithm; it only works to eliminate a single extra hop from the path. There are methods by which this mechanism might be extended to include larger optimizations, but the potential to create routing loops has not been sufficiently analyzed to specify them here.

3.3 Authentication

The need for authentication in a routing protocol is obvious. It is not usually important to conceal the information in the routing messages, but it is essential to prevent the insertion of bogus routing information into the routers. So, while the authentication mechanism specified in RIP-2 is less than ideal, it does prevent anyone who cannot directly access the network (i.e., someone who cannot sniff the routing packets to determine the password) from inserting bogus routing information.

However, the specification does allow for additional types of authentication to be incorporated into the protocol. Unfortunately, because of the original format of RIP packets, the amount of space available for providing authentication information is only 16 octets.

3.4 Multicasting

The RIP-2 protocol provides for the IP multicasting of periodic advertisements. This feature was added to decrease the load on systems which do not support RIP-2. It also provides a mechanism whereby RIP-1 routers will never receive RIP-2 routes. This is a feature when correct use of an advertised route depends on knowing the precise subnet mask, which would be ignored by a RIP-1 router.

4. Conclusion

Because the basic protocol is unchanged, RIP-2 is as correct a routing protocol as RIP-1. The enhancements make RIP-2 useful in environments which RIP-1 could not handle, but which do not necessitate the use of OSPF by virtue of requirements which RIP-2 does not satisfy.

5. References

- [1] Malkin, G., "RIP Version 2 Protocol Analysis", RFC 1721, Xylogics, Inc., November 1994.
- [2] Malkin, G., "RIP Version 2 - Carrying Additional Information", RFC 1723, Xylogics, Inc., November 1994.
- [3] Malkin, G., and F. Baker, "RIP Version 2 MIB Extension", RFC 1724, Xylogics, Inc., Cisco Systems, November 1994.

6. Security Considerations

Security issues are not discussed in this memo.

7. Author's Address

Gary Scott Malkin
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803

Phone: (617) 272-8140
EMail: gmalkin@Xylogics.COM