

RIP Version 2
Carrying Additional Information

Status of this Memo

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

Abstract

This document specifies an extension of the Routing Information Protocol (RIP), as defined in [1], to expand the amount of useful information carried in RIP packets and to add a measure of security. A companion document will define the SNMP MIB objects for RIP-2 [2].

Acknowledgements

I would like to thank the following for their contributions to this document: Fred Baker, Noel Chiappa and Vince Fuller. This memo is a product of the RIP-2 Working Group of the Internet Engineering Task Force (IETF).

Table of Contents

1.	Justification	2
2.	Current RIP	2
3.	Protocol Extensions	2
3.1	Authentication	3
3.2	Routing Domain	4
3.3	Route Tag	4
3.4	Subnet Mask	4
3.5	Next Hop	4
3.6	Multicasting	5
4.	Compatibility	5
4.1	Compatibility Switch	5
4.2	Authentication	6
4.3	Larger Infinity	6
4.4	Addressless Links	6
	Appendix A	6
	References	7

Security Considerations 7
 Author's Address 7

1. Justification

With the advent of OSPF and IS-IS, there are those who believe that RIP is obsolete. While it is true that the newer IGP routing protocols are far superior to RIP, RIP does have some advantages. Primarily, in a small network, RIP has very little overhead in terms of bandwidth used and configuration and management time. RIP is also very easy to implement, especially in relation to the newer IGPs.

Additionally, there are many, many more RIP implementations in the field than OSPF and IS-IS combined. It is likely to remain that way for some years yet.

Given that RIP will be useful in many environments for some period of time, it is reasonable to increase RIP's usefulness. This is especially true since the gain is far greater than the expense of the change.

2. Current RIP

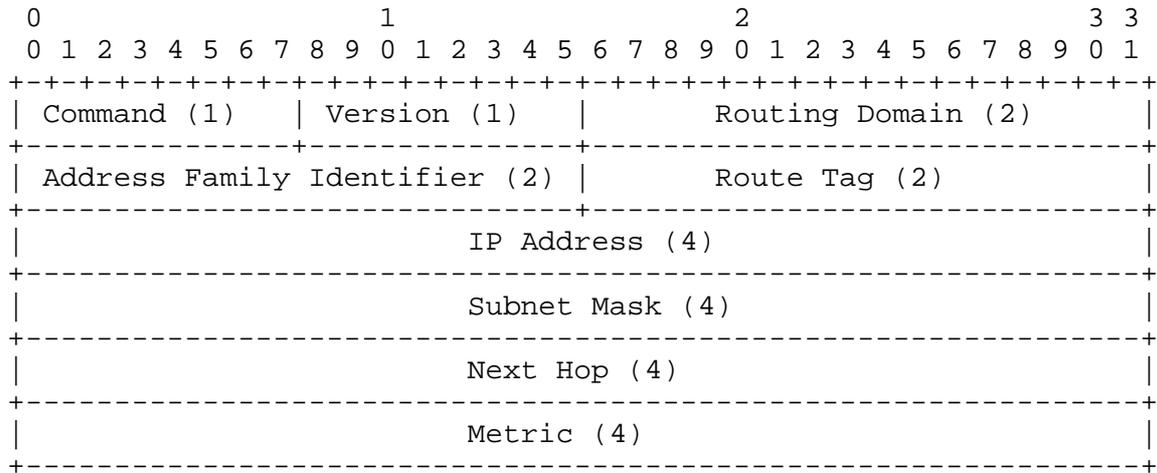
The current RIP packet contains the minimal amount of information necessary for routers to route packets through a network. It also contains a large amount of unused space, owing to its origins.

The current RIP protocol does not consider autonomous systems and IGP/EGP interactions, subnetting, and authentication since implementations of these postdate RIP. The lack of subnet masks is a particularly serious problem for routers since they need a subnet mask to know how to determine a route. If a RIP route is a network route (all non-network bits 0), the subnet mask equals the network mask. However, if some of the non-network bits are set, the router cannot determine the subnet mask. Worse still, the router cannot determine if the RIP route is a subnet route or a host route. Currently, some routers simply choose the subnet mask of the interface over which the route was learned and determine the route type from that.

3. Protocol Extensions

This document does not change the RIP protocol per se. Rather, it provides extensions to the datagram format which allows routers to share important additional information.

The new RIP datagram format is:

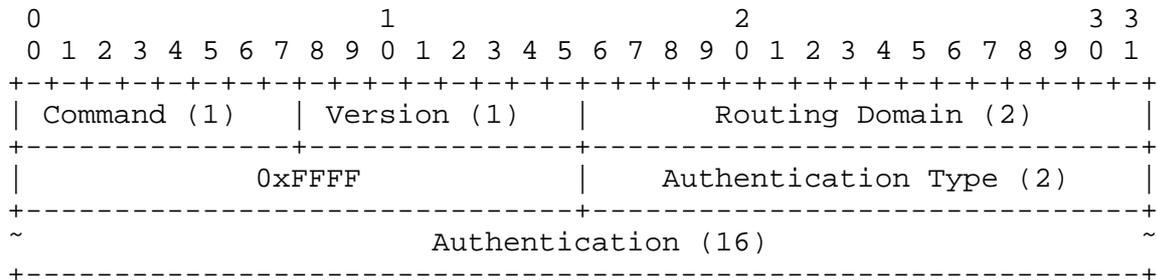


The Command, Address Family Identifier (AFI), IP Address, and Metric all have the meanings defined in RFC 1058. The Version field will specify version number 2 for RIP datagrams which use authentication or carry information in any of the newly defined fields.

All fields are coded in IP network byte order (big-endian).

3.1 Authentication

Since authentication is a per packet function, and since there is only one 2-byte field available in the packet header, and since any reasonable authentication scheme will require more than two bytes, the authentication scheme for RIP version 2 will use the space of an entire RIP entry. If the Address Family Identifier of the first (and only the first) entry in the packet is 0xFFFF, then the remainder of the entry contains the authentication. This means that there can be, at most, 24 RIP entries in the remainder of the packet. If authentication is not in use, then no entries in the packet should have an Address Family Identifier of 0xFFFF. A RIP packet which contains an authentication entry would have the following format:



Currently, the only Authentication Type is simple password and it is type 2. The remaining 16 bytes contain the plain text password. If the password is under 16 bytes, it must be left-justified and padded to the right with nulls (0x00).

3.2 Routing Domain

The Routing Domain (RD) number is the number of the routing process to which this update belongs. This field is used to associate the routing update to a specific routing process on the receiving router. The RD is needed to allow multiple, independent RIP "clouds" to co-exist on the same physical wire. This gives administrators the ability to run multiple, possibly parallel, instances of RIP in order to implement simple policy. This means that a router operating within one routing domain, or a set of routing domains, should ignore RIP packets which belong to another routing domain. RD 0 is the default routing domain.

3.3 Route Tag

The Route Tag (RT) field exists as a support for EGPs. The contents and use of this field are outside the scope of this protocol. However, it is expected that the field will be used to carry Autonomous System numbers for EGP and BGP. Any RIP system which receives a RIP entry which contains a non-zero RT value must re-advertise that value. Those routes which have no RT value must advertise an RT value of zero.

3.4 Subnet mask

The Subnet Mask field contains the subnet mask which is applied to the IP address to yield the non-host portion of the address. If this field is zero, then no subnet mask has been included for this entry.

On an interface where a RIP-1 router may hear and operate on the information in a RIP-2 routing entry the following two rules apply:

- 1) information internal to one network must never be advertised into another network, and
- 2) information about a more specific subnet may not be advertised where RIP-1 routers would consider it a host route.

3.5 Next Hop

The immediate next hop IP address to which packets to the destination specified by this route entry should be forwarded. Specifying a value of 0.0.0.0 in this field indicates that routing should be via

the originator of the RIP advertisement. An address specified as a next hop must, per force, be directly reachable on the logical subnet over which the advertisement is made.

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. A simple example is given in Appendix A. Note that Next Hop is an "advisory" field. That is, if the provided information is ignored, a possibly sub-optimal, but absolutely valid, route may be taken.

3.6 Multicasting

In order to reduce unnecessary load on those hosts which are not listening to RIP-2 packets, an IP multicast address will be used for periodic broadcasts. The IP multicast address is 224.0.0.9. Note that IGMP is not needed since these are inter-router messages which are not forwarded.

In order to maintain backwards compatibility, the use of the multicast address will be configurable, as described in section 4.1. If multicasting is used, it should be used on all interfaces which support it.

4. Compatibility

RFC 1058 showed considerable forethought in its specification of the handling of version numbers. It specifies that RIP packets of version 0 are to be discarded, that RIP packets of version 1 are to be discarded if any Must Be Zero (MBZ) field is non-zero, and that RIP packets of any version greater than 1 should not be discarded simply because an MBZ field contains a value other than zero. This means that the new version of RIP is totally backwards compatible with existing RIP implementations which adhere to this part of the specification.

4.1 Compatibility Switch

A compatibility switch is necessary for two reasons. First, there are implementations of RIP-1 in the field which do not follow RFC 1058 as described above. Second, the use of multicasting would prevent RIP-1 systems from receiving RIP-2 updates (which may be a desired feature in some cases).

The switch has three settings: RIP-1, in which only RIP-1 packets are sent; RIP-1 compatibility, in which RIP-2 packets are broadcast; and RIP-2, in which RIP-2 packets are multicast. The recommended default for this switch is RIP-1 compatibility.

4.2 Authentication

Since an authentication entry is marked with an Address Family Identifier of 0xFFFF, a RIP-1 system would ignore this entry since it would belong to an address family other than IP. It should be noted, therefore, that use of authentication will not prevent RIP-1 systems from seeing RIP-2 packets. If desired, this may be done using multicasting, as described in sections 3.6 and 4.1.

4.3 Larger Infinity

While on the subject of compatibility, there is one item which people have requested: increasing infinity. The primary reason that this cannot be done is that it would violate backwards compatibility. A larger infinity would obviously confuse older versions of rip. At best, they would ignore the route as they would ignore a metric of 16. There was also a proposal to make the Metric a single byte and reuse the high three bytes, but this would break any implementations which treat the metric as a long.

4.4 Addressless Links

As in RIP-1, addressless links will not be supported by RIP-2.

Appendix A

This is a simple example of the use of the next hop field in a rip entry.



Assume that IR1, IR2, and IR3 are all "internal" routers which are under one administration (e.g., a campus) which has elected to use RIP-2 as its IGP. XR1, XR2, and XR3, on the other hand, are under separate administration (e.g., a regional network, of which the campus is a member) and are using some other routing protocol (e.g., OSPF). XR1, XR2, and XR3 exchange routing information among themselves such that they know that the best routes to networks N1 and N2 are via XR1, to N3, N4, and N5 are via XR2, and to N6 and N7 are via XR3. By setting the Next Hop field correctly (to XR2 for N3/N4/N5, to XR3 for N6/N7), only XR1 need exchange RIP-2 routes with IR1/IR2/IR3 for routing to occur without additional hops through XR1. Without the Next Hop (for example, if RIP-1 were used) it would be

necessary for XR2 and XR3 to also participate in the RIP-2 protocol to eliminate extra hops.

References

- [1] Hedrick, C., "Routing Information Protocol", RFC 1058, Rutgers University, June 1988.
- [2] Malkin, G., and F. Baker, "RIP Version 2 MIB Extension", RFC 1389, Xylogics, Inc., Advanced Computer Communications, January 1993.
- [3] Malkin, G., "RIP Version 2 Protocol Analysis", RFC 1387, Xylogics, Inc., January 1993.

Security Considerations

The basic RIP protocol is not a secure protocol. To bring RIP-2 in line with more modern routing protocols, an extensible authentication mechanism has been incorporated into the protocol enhancements. This mechanism is described in sections 3.1 and 4.2.

Author's Address

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

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