

>

Network Working Group
Request for Comments: 2453
Obsoletes: 1723, 1388
STD: 56
Category: Standards Track

G. Malkin
Bay Networks
November 1998

RIP Version 2

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.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

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 messages and to add a measure of security.

A companion document will define the SNMP MIB objects for RIP-2 [2]. An additional document will define cryptographic security improvements for RIP-2 [3].

Acknowledgements

I would like to thank the IETF RIP Working Group for their help in improving the RIP-2 protocol. Much of the text for the background discussions about distance vector protocols and some of the descriptions of the operation of RIP were taken from "Routing Information Protocol" by C. Hedrick [1]. Some of the final editing on the document was done by Scott Bradner.

Table of Contents

1.	Justification	3
2.	Current RIP	3
3.	Basic Protocol	3
3.1	Introduction	3
3.2	Limitations of the Protocol	5
3.3.	Organization of this document	6
3.4	Distance Vector Algorithms	6
3.4.1	Dealing with changes in topology	12
3.4.2	Preventing instability	13
3.4.3	Split horizon	15
3.4.4	Triggered updates	17
3.5	Protocol Specification	18
3.6	Message Format	20
3.7	Addressing Considerations	22
3.8	Timers	24
3.9	Input Processing	25
3.9.1	Request Messages	25
3.9.2	Response Messages	26
3.10	Output Processing	28
3.10.1	Triggered Updates	29
3.10.2	Generating Response Messages.	30
4.	Protocol Extensions	31
4.1	Authentication	31
4.2	Route Tag	32
4.3	Subnet Mask	32
4.4	Next Hop	33
4.5	Multicasting	33
4.6	Queries	33
5.	Compatibility	34
5.1	Compatibility Switch	34
5.2	Authentication	34
5.3	Larger Infinity	35
5.4	Addressless Links	35
6.	Interaction between version 1 and version 2	35
7.	Security Considerations	36
	Appendices	37
	References	37
	Author's Address	38
	Full Copyright Statement	39

With the advent of OSPF and IS-IS, there are those who believe that RIP is obsolete.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-1 message contains the minimal amount of information necessary for routers to route messages through a network. It also contains a large amount of unused space, owing to its origins.

The current RIP-1 protocol does not consider autonomous systems and IGP/EGP interactions, subnetting [11], and authentication since implementations of these postdate RIP-1. 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-1 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-1 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. Basic Protocol

3.1 Introduction

RIP is a routing protocol based on the Bellman-Ford (or distance vector) algorithm. This algorithm has been used for routing computations in computer networks since the early days of the ARPANET. The particular packet formats and protocol described here are based on the program "routed," which is included with the Berkeley distribution of Unix.

Malkin

Standards Track

[Page 3]

RFC 2453

RIP Version 2

November 1998

In an international network, such as the Internet, it is very unlikely that a single routing protocol will be used for the entire network. Rather, the network will be organized as a collection of Autonomous Systems (AS), each of which will, in general, be administered by a single entity. Each AS will have its own routing technology, which may differ among AS's. The routing protocol used

within an AS is referred to as an Interior Gateway Protocol (IGP). A separate protocol, called an Exterior Gateway Protocol (EGP), is used to transfer routing information among the AS's. RIP was designed to work as an IGP in moderate-size AS's. It is not intended for use in more complex environments. For information on the context into which RIP-1 is expected to fit, see Braden and Postel [6].

RIP uses one of a class of routing algorithms known as Distance Vector algorithms. The earliest description of this class of algorithms known to the author is in Ford and Fulkerson [8]. Because of this, they are sometimes known as Ford-Fulkerson algorithms. The term Bellman-Ford is also used, and derives from the fact that the formulation is based on Bellman's equation [4]. The presentation in this document is closely based on [5]. This document contains a protocol specification. For an introduction to the mathematics of routing algorithms, see [1]. The basic algorithms used by this protocol were used in computer routing as early as 1969 in the ARPANET. However, the specific ancestry of this protocol is within the Xerox network protocols. The PUP protocols [7] used the Gateway Information Protocol to exchange routing information. A somewhat updated version of this protocol was adopted for the Xerox Network Systems (XNS) architecture, with the name Routing Information Protocol [9]. Berkeley's routed is largely the same as the Routing Information Protocol, with XNS addresses replaced by a more general address format capable of handling IPv4 and other types of address, and with routing updates limited to one every 30 seconds. Because of this similarity, the term Routing Information Protocol (or just RIP) is used to refer to both the XNS protocol and the protocol used by routed.

RIP is intended for use within the IP-based Internet. The Internet is organized into a number of networks connected by special purpose gateways known as routers. The networks may be either point-to-point links or more complex networks such as Ethernet or token ring. Hosts and routers are presented with IP datagrams addressed to some host. Routing is the method by which the host or router decides where to send the datagram. It may be able to send the datagram directly to the destination, if that destination is on one of the networks that are directly connected to the host or router. However, the interesting case is when the destination is not directly reachable.

Malkin

Standards Track

[Page 4]

RFC 2453

RIP Version 2

November 1998

In this case, the host or router attempts to send the datagram to a router that is nearer the destination. The goal of a routing protocol is very simple: It is to supply the information that is needed to do routing.

3.2 Limitations of the Protocol

This protocol does not solve every possible routing problem. As mentioned above, it is primary intended for use as an IGP in networks of moderate size. In addition, the following specific limitations

are be mentioned:

- The protocol is limited to networks whose longest path (the network's diameter) is 15 hops. The designers believe that the basic protocol design is inappropriate for larger networks. Note that this statement of the limit assumes that a cost of 1 is used for each network. This is the way RIP is normally configured. If the system administrator chooses to use larger costs, the upper bound of 15 can easily become a problem.
- The protocol depends upon "counting to infinity" to resolve certain unusual situations. (This will be explained in the next section.) If the system of networks has several hundred networks, and a routing loop was formed involving all of them, the resolution of the loop would require either much time (if the frequency of routing updates were limited) or bandwidth (if updates were sent whenever changes were detected). Such a loop would consume a large amount of network bandwidth before the loop was corrected. We believe that in realistic cases, this will not be a problem except on slow lines. Even then, the problem will be fairly unusual, since various precautions are taken that should prevent these problems in most cases.
- This protocol uses fixed "metrics" to compare alternative routes. It is not appropriate for situations where routes need to be chosen based on real-time parameters such a measured delay, reliability, or load. The obvious extensions to allow metrics of this type are likely to introduce instabilities of a sort that the protocol is not designed to handle.

Malkin

Standards Track

[Page 5]

RFC 2453

RIP Version 2

November 1998

3.3. Organization of this document

The main body of this document is organized into two parts, which occupy the next two sections:

A conceptual development and justification of distance vector algorithms in general.

The actual protocol description.

Each of these two sections can largely stand on its own. Section 3.4 attempts to give an informal presentation of the mathematical underpinnings of the algorithm. Note that the presentation follows a "spiral" method. An initial, fairly simple algorithm is described.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.