

RIPng Protocol Applicability Statement

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

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

1. Protocol Documents

The RIPng protocol description is defined in RFC 2080.

2. Introduction

This report describes how RIPng may be useful within the new IPv6 Internet. In essence, the environments in which RIPng is the IGP of choice is comparable to the environments in which RIP-2 (RFC 1723) is used in the IPv4 Internet. It is important to remember that RIPng is a simple extrapolation of RIP-2; RIPng has nothing conceptually new. Thus, the operational aspects of distance-vector routing protocols, and RIP-2 in particular, within an autonomous system are well understood.

It should be noted that RIPng is not intended to be a substitute for OSPFng in large autonomous systems; the restrictions on AS diameter and complexity which applied to RIP-2 also apply to RIPng. Rather, RIPng 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 features of RIPng is useful within IPv6.

3. Applicability

A goal in developing RIPng was to make the minimum necessary change to RIP-2 to produce RIPng. In essence, the IPv4 address was expanded into an IPv6 address, the IPv4 subnet mask was replaced with an IPv6 prefix length, the next-hop field was eliminated but the functionality has been preserved, and authentication was removed. The route tag field has been preserved. The maximum diameter of the network (the maximum metric value) is 15; 16 still means infinity (unreachable).

The basic RIP header is unchanged. However, the size of a routing packet is no longer arbitrarily limited. Because routing updates are never forwarded, the routing packet size is now determined by the physical media and the sizes of the headers which precede the routing data (i.e., media MTU minus the combined header lengths). The number routes which may be included in a routing update is the routing data length divided by the size of a routing entry.

3.1 Prefix

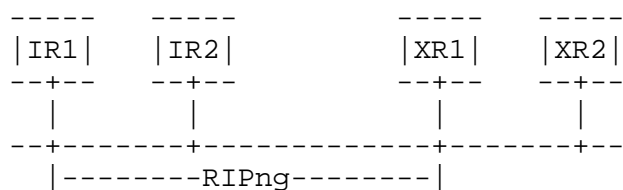
The address field of a routing entry is 128 bits in length, expanded from the 32 bits available in RIP-2. This allows the RIP entry to carry an IPv6 prefix.

3.2 Prefix Length

The 32-bit RIP-2 subnet mask field is replaced by an 8-bit prefix length field. It allows the specification of the number of bits in the prefix which form the actual prefix.

3.3 Next Hop

The ability to specify the next hop, rather than simply allowing the recipient of the update to set the next hop to the sender of the update, allows for the elimination of unnecessary hops through routers which are running multiple routing protocols. Consider following example topology:



The Internal Routers (IR1 and IR2) are only running RIPng. The External Routers (XR1 and XR2) are both running BGP, for example; however, only XR1 is running BGP and RIPng. Since XR2 is not running RIPng, 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 mechanism, that XR2 is the better next hop for some routes.

3.4 Authentication

Authentication, which was added to RIP-2 because RIP-1 did not have it, has been dropped from RIPng. This is safe to do because IPv6, which carries the RIPng packets, has build in security which IPv4 did not have.

3.5 Packet Length

By allowing RIPng routing update packets to be as big as possible, the number of packets which must be sent for a complete update is greatly reduced. This in no way affects the operation of the distance-vector protocol; it is merely a performance enhancement.

3.6 Diameter and Complexity

The limit of 15 cost-1 hops is a function of the distance-vector protocol, which depends on counting to infinity to resolve some routing loops. If infinity is too high, the time it would take to resolve, not to mention the number of routing updates which would be sent, would be prohibitive. If the infinity is too small, the protocol becomes useless in a reasonably sized network. The choice of 16 for infinity was made in the earliest of RIP implementations and experience has shown it to be a good compromise value.

RIPng will efficiently support networks of moderate complexity. That is, topologies without too many multi-hop loops. RIPng also efficiently supports topologies which change frequently because routing table changes are made incrementally and do not require the computation which link-state protocols require to rebuild their maps.

4. Conclusion

Because the basic protocol is unchanged, RIPng is as correct a routing protocol as RIP-2. RIPng serves the same niche for IPv6 as RIP-2 does for IPv4.

5. Security Considerations

RIPng security is discussed in section 3.4.

Author's Address

Gary Scott Malkin
Xylogics/Bay Networks
53 Third Avenue
Burlington, MA 01803

Phone: (617) 238-6237
EMail: gmalkin@xylogics.com

