

## IPv6 Routing Table Size Issues

### STATUS OF THIS MEMO

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its working groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of 6 months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as "work in progress".

This particular Internet Draft is intended for submission to the IETF's IPng working group for possible future publication as a standards-track or informational RFC.

### ABSTRACT

This note describes certain issues relating to the size of the IPv6 routing table in backbone routers. It also suggests several possible approaches to dealing with those issues.

### **1. INTRODUCTION**

**IP version 6 (IPv6) is being developed within the IETF as a long-term replacement for the current IP version 4 (IPv4).** [DH96] As part of IPv6 development, a set of IPv4-compatible IPv6 addresses have been defined. [GN96] These addresses are 128 bits long and consist of a fixed 96 bit prefix and then contain an embedded 32-bit IPv4 address in the low-order portion of the address. [HD96]

In the current IPv4 backbone, the size of the IPv4 routing table is a significant operational issue. Although the deployment of CIDR [FLYV93] has significantly reduced the rate of increase of the total

set of IPv4 routes, the total number of such routes continues to increase. For the purpose of this discussion, a backbone router is defined as a router that contains all IPv4 routes except for the "default" route. As of this writing, a backbone router carries perhaps 47,000 IPv4 routes, including about 30,000 routes that are external to the router's own network and perhaps 17,000 routes that are internal to the router's own network and hence cannot be aggregated. A number of activities performed by a router have costs proportional to the total size of the routing table. Both size and growth rate of the total set of IPv4 routes is a significant operational issue for backbone networks.

Deployment of IPv6 in a backbone will increase the number of total routes that backbone routers have to carry. To the extent that IPv6 addresses can be aggregated more effectively than IPv4 routes have been, the issues posed by deployment of a second network-layer protocol can be mitigated. However, each IPv6 routing prefix is 128 bits long as compared with the 32 bits for an IPv4 routing prefix. The overall size of a routing table entry varies significantly in different implementations of IPv6 routing, but it is very clear that carrying an IPv4 prefix once natively and a second time as if it were an IPv6 prefix causes significant increase in routing table size. Even if an IPv6 route entry were the same size as an IPv4 route entry, the router would need twice the memory for the routing table. This is a significant operational issue in default-less routers.

## **2. ALTERNATIVE APPROACHES**

There are at least three possible approaches to this issue. This section discusses these alternatives and the issues associated with each.

### **2.1 Do Nothing**

**This is the simplest alternative. It involves taking no particular actions to preclude backbone routing tables from increasing in size.** This approach might work if backbone operational networks obtain and deploy routers capable of handling the significantly larger number of backbone routes.

### **2.2 Route Filtering**

**One alternative is for network operators concerned about this issue to filter the routes that they accept via their routing protocols so that IPv4-compatible IPv6 prefixes are not accepted into the routing table and are not propagated within that operator's routing domain.** This can work but potentially consumes significant amounts of a router's resources, potentially adversely impact the forwarding rate of the routers engaged in such filtering. Some believe that this would also break IPv6 routing.

Atkinson

Expires in 6 months

[Page 2]

### **2.3 Avoid carrying IPv4-compatible Routing Prefixes**

Another alternative is to avoid carrying IPv4-compatible IPv6 routing prefixes as if they were native IPv6 routing prefixes. In this alternative, the routing protocol specifications would prohibit IPv4-compatible IPv6 routing prefixes as native IPv6 prefixes. Of course, native IPv4 prefixes that were indicated as being IPv4 prefixes could still be carried by an routing protocol supporting integrated routing (e.g. I/ISIS).

In this approach, a dual-stack (IPv4 and IPv6) router that received a native unencapsulated IPv6 packet destined for an IPv4-compatible IPv6 address that did not have an IPv6 route for that particular IPv4-compatible IPv6 address would immediately encapsulate that IPv6 packet inside IPv4 and forward it in tunneled form to the IPv4 destination address. This would involve increased cost (due to the encapsulation) at the outer edge routers and would eliminate IPv4-compatible IPv6 addresses as a potential source of greatly increased routing tables for the backbone routers.

Because the existence of an IPv4 route in an IPv4 routing table does not provide any information about whether the next hop is also capable of handling IPv6 packets, encapsulation of the IPv6 packet inside the IPv4 packet prior to forwarding is necessary unless the forwarding node has certain knowledge that the next-hop address for the IPv4 route is also capable of handling IPv6 packets.

## **3. PROPOSED ACTION**

It is proposed that the specifications for IPv6-capable routing protocols clearly indicate that router implementations MUST NOT by default inject any part of their logical IPv4 routing table into their logical IPv6 routing table, that such specifications note the operational issue that is described in this note, and that such specifications discourage implementers from permitting such cross-protocol route injection from occurring. Of course, if the operator of the router explicitly chose to inject IPv4-compatible IPv6 routes into the IPv6 routing table and live with the results, this would not be prohibited. The affected routing protocols appear to include OSPFv3 [CFM96], RIPv2 for IPv6 [MM96], I/ISIS, and IDRPv2 [RT96].

Dual-stack routers not having an explicit route for an IPv4-compatible IPv6 address will be required to always encapsulate all IPv6 packets forwarded to an IPv4-compatible IPv6 destination address inside IPv4 and then use the IPv4 routing table to lookup paths to IPv4-compatible destination addresses. Encapsulation is generally required before forwarding for the reason outlined in the previous section.

Atkinson

Expires in 6 months

[Page 3]

#### **4. SECURITY CONSIDERATIONS**

It is possible to take out an operational network through inappropriate injection of incorrect routes. It is also possible to take out an operational network by injecting more routes than that network's routing systems are able to cope with. These can occur through operator error as well as through the malicious actions of third parties. Intentional injection of too many routes into a network's routing system such that the victim network ceases proper operation constitutes a denial of service attack.

#### **REFERENCES**

- [DH96] Steve Deering & Robert Hinden, "IP version 6 Specification", [RFC-1883](#), January 1996.
- [FLYV93] V. Fuller, T. Li, J. Yu, & K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", [RFC-1519](#), September 1993.
- [GN96] Bob Gilligan & Erik Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", [RFC-1933](#), April 1996.
- [HD96] Robert Hinden & Steve Deering, "IP version 6 Addressing Architecture", [RFC-1884](#), January 1996.
- [MM96] Gary Malkin & R. Minnear, "RIPng for IPv6", Internet Draft, August 1996.
- [CFM96] Rob Coltun, Dennis Ferguson, & John Moy, "OSPF for IPv6", Internet Draft, June 1996.
- [RT96] Yakov Rekhter & Paul Traina, "Inter-Domain Routing Protocol Version 2", Internet Draft, June 1996.

#### **AUTHOR'S ADDRESS:**

Randall Atkinson  
cisco Systems  
170 West Tasman Drive  
San Jose, CA 95134-1706 USA

Email: [rja@cisco.com](mailto:rja@cisco.com)  
Telephone: +1 (408) 526-6566

Atkinson

Expires in 6 months

[Page 4]

