

Routing of Scoped Addresses
in the Internet Protocol Version 6 (IPv6)

[<draft-ietf-ipngwg-scoped-routing-00.txt>](#)

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 six months, and 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 a ``working draft'' or ``work in progress.''

To learn the current status of any Internet Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), ftp.isi.edu (US West Coast).

Abstract

This document outlines a mechanism for generating routing tables that include scoped IPv6 addresses. It defines a set of rules for routers to implement in order to forward scoped unicast and multicast addresses regardless of the routing protocol. It should be noted that these rules will apply to all scoped addresses.

Internet Draft

Routing of Scoped Addresses

September 1998

Contents

1. Introduction	1
2. Assumptions and Definitions	1
3. Single Site Routing	1
4. Site Boundary Unicast Routing	2
4.1. Routing Protocols	2
4.2. Packet Forwarding	4
5. Scoped Multicast Routing	5
5.1. Routing Protocols	5
5.2. Packet Forwarding	5
6. Protocol Impact	6

1. Introduction

This document defines a set of rules for the generation of forwarding tables that contain scoped addresses. This document describes the handling of scoped addresses for both single site and site boundary routers. Ideally, these concepts should be included in followup drafts of IPv6 routing protocols.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC 2119](#)].

2. Assumptions and Definitions

This document makes several assumptions concerning sites :

- Site boundaries cut through nodes
- Site boundaries are identical for unicast and multicast traffic
- A single interface can only be in one site
- Each interface participating in a site has a site identifier
- In the absence of explicit configuration, all site identifiers on a node default to the same value

A single site router is defined as a router configured with the same site identifier on all interfaces. A site boundary router is defined as a router that has at least 2 distinct site identifiers configured. This could include a router connected to 2 distinct sites or a router connected to 1 site and a separate global network (Figure 1).

3. Single Site Routing

In a single site router, a routing protocol can advertise and route all addresses and prefixes on all interfaces. This configuration does not require any special handling for site local addresses. The reception and transmission of site local addresses is handled in the same manner as globally scoped addresses. This applies to both unicast and multicast routing protocols.




Figure 1: Multi-Sited Router

4. Site Boundary Unicast Routing

With respect to site boundaries, routers must consider which interfaces a packet can be transmitted on as well as control the propagation of routing information specific to the site. This includes controlling which prefixes can be advertised on an interface.

[4.1.](#) Routing Protocols

When a routing protocol determines that it is a site boundary router, it must perform additional work in order to protect inter site integrity and still maintain intra site connectivity.

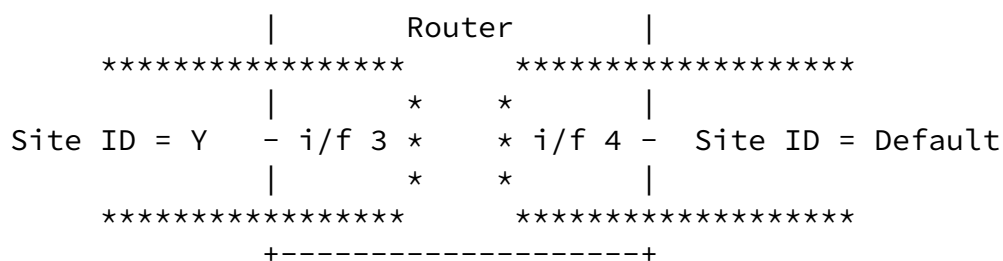
In order to maintain connectivity, the routing protocol must be able to create forwarding information for the global prefixes as well as for all of the site prefixes for each of its attached sites. The most straight forward way of doing this is to create up to $(n + 1)$ routing tables; one for the global prefixes, if any, and one for each of the (n) sites.

To protect inter site integrity, routers must be selective in the forwarding information that is shared with neighboring routers. Routing protocols routinely transmit their routing information to neighboring routers. When a router is transmitting this routing information, it must not include any information about sites other than the site defined on the interface used to reach a neighbor.

```

      *                               *
      *                               *
      *   Site ID = X                 *
      *                               *
      *                               *
      *                               *
+---*---|-----|---*---+
| * i/f 1   i/f 2 * |
| ***** |
|           |
|           |

```



i/f 1 : global prefix = 3FFE:20::/64
 site prefix = FEC0:0:0:N/64

i/f 2 : no global prefix
 site prefix = FEC0:0:0:K/64

i/f 3 : global prefix = 3FFE:40::/64
 site prefix = FEC0:0:0:M/64

i/f 4 : global prefix = 3FFE:80::/64
 no site prefix

Figure 2: Routing Information Exchange

As an example, the router in Figure 2 must advertise routing information on four interfaces. The information advertised is as follows :

- Interface 1
 - * All global prefixes (3FFE:20::/64, 3FFE:40::/64, and 3FFE:80::/64)
 - * Site prefix FEC0:0:0:N/64
 - * Site prefix FEC0:0:0:K/64
- Interface 2
 - * All global prefixes (3FFE:20::/64, 3FFE:40::/64, and

3FFE:80::/64)

- * Site prefix FEC0:0:0:N/64
- * Site prefix FEC0:0:0:K/64
- Interface 3
 - * All global prefixes (3FFE:20::/64, 3FFE:40::/64, and 3FFE:80::/64)
 - * Site prefix FEC0:0:0:M/64
- Interface 4
 - * All global prefixes (3FFE:20::/64, 3FFE:40::/64, and 3FFE:80::/64)

By imposing advertisement rules, site integrity is maintained by keeping all site routing information contained within the site.

[4.2.](#) Packet Forwarding

In addition to the extra cost of generating additional forwarding information for each site, site boundary routers must also do some additional checking when forwarding packets that contain site local addresses.

If a packet being forwarded contains a site local destination address, regardless of the scope of the source address, the router must perform the following :

- Lookup incoming interface's site identifier

- Perform route lookup for destination address in arrival interfaces site scoped routing table

If a packet being forwarded contains a site local source address and a globally scoped destination address, the following must

be performed :

- Lookup outgoing interface's site identifier
- Compare inbound and outbound interfaces' site identifiers

If the site identifiers match, the packet can be forwarded. If they do not match, an ICMPv6 destination unreachable message must be sent to the sender with a new code value, code = 5 (Scope Mismatch).

This ICMPv6 message will indicate that the destination address is not reachable with the specified source address.

[5. Scoped Multicast Routing](#)

With IPv6 multicast, there are multiple scopes supported. Multicast routers must be able to control the propagation of scoped packets based on administratively configured boundaries.

[5.1. Routing Protocols](#)

Multicast routing protocols will have to follow the same rules as the unicast protocols. They will be required to maintain information about global prefixes as well as information about all scope boundaries that pass through the router. Multicast protocols that rely on underlying unicast protocols (i.e. PIM) will not suffer as much of a performance impact since the unicast protocol will handle the forwarding table generation. They must be able to handle the additional scope boundaries used in multicast addresses.

Multicast protocols that generate and maintain their own routing tables will have to perform the additional route calculations for scope. All multicast protocols will be forced to handle 14 additional scoping identifiers above the site identifiers supported in IPv6 unicast addresses.

[5.2. Packet Forwarding](#)

The forwarding rules for multicast can be described by the following combinations :

- Global multicast destination address / Global unicast source address
- Global multicast destination address / Site local unicast source address
- Scoped multicast destination address / Global unicast source address
- Scoped multicast destination address / Site local unicast source address

The first combination requires no special processing over what is currently in place for global IPv6 multicast. Combinations 2,3, and 4 should result in the router performing the same identifiers check as outlined for site local unicast addresses. Since IPv6 supports fifteen unique multicast scopes, it is assumed that scopes 0x1 through 0x4 are strictly less than the unicast site scope, scope 0x5 (site) is equal to the unicast site scope, scopes 0x6 through 0xd are strictly greater than the unicast site scope and strictly less than the unicast global scope, and scope 0xe is equal to the unicast global scope.

6. Protocol Impact

The performance impact on routing protocols is obvious. Routers implementing scoped address support will be forced to perform an additional check in the main forwarding path to determine if the source address is scoped. This will add overhead to the processing of every packet flowing through the router. In addition, there will be some storage overhead for the scope identifiers. If scoped addresses are going to be realized, this performance impact may be acceptable.

References

[RFC 2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP14](#), March 1997.

Security Considerations

This document specifies a set of guidelines that allow routers to prevent site

Internet Draft

Routing of Scoped Addresses

September 1998

specific information from leaking out of each site. If site boundary routers allow site routing information to be forwarded outside of the site, the integrity of the site could be compromised.

Acknowledgements

I would like to thank Thomas Narten, Steve Deering, and Erik Nordmark for their comments and reviews of this document.

Author's Address

Brian Haberman
IBM Corporation
800 Park Office Drive
Research Triangle Park, NC 27709
USA
+1-919-254-2673
haberman@raleigh.ibm.com

Haberman

Expires March 1999

[Page 7]