

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

[<draft-haberman-ipv6-site-route-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 view the entire list of current Internet-Drafts, please check the "lid-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), or ftp.isi.edu (US West Coast).

## Abstract

This document outlines a mechanism for generating routing tables that include site-scoped IPv6 addresses. It defines a set of rules for routers to implement in order to forward site-scoped addresses regardless of the routing protocol.

Internet Draft

Routing of Site-Scoped Addresses

April 1998


## Contents

Status of This Memo	i
Abstract	i
1. Introduction	1
2. Assumptions and Definitions	2
3. Single Site Routing	3
4. Site-Boundary Unicast Routing	3
<a href="#">4.1.</a> Routing Protocols . . . . .	<a href="#">3</a>
<a href="#">4.2.</a> Packet Forwarding . . . . .	<a href="#">5</a>
5. Site-Boundary Multicast Routing	6
<a href="#">5.1.</a> Routing Protocols . . . . .	<a href="#">6</a>
<a href="#">5.2.</a> Packet Forwarding . . . . .	<a href="#">6</a>
6. Protocol Impact	6

## 1. Introduction

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

A preferred model for site-boundaries is depicted in Figure 1. In this model, each router is responsible for generating forwarding information for the global prefixes and exactly 1 site. The link between the routers is in neither site. In addition, routing information about Site Y is never propagated into Site X and routing information about Site X is never propagated into Site Y. This model is similar to that of net 10 routing in IPv4.



\* \* \* \* \*

\* \* \* \* \*

Figure 1: Site Boundary Model

The rest of this document describes the modifications needed in order to combine the functions of Router 1 and Router 2 into a single router.

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].

Haberman

Expires October 1998

[Page 1]

Internet Draft

## Routing of Site-Scoped Addresses

April 1998

## 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 router default to the same value
- Routers only advertise site prefixes on interfaces configured with site identifiers and site prefixes

[illegible]




Figure 2: Multi-Sited Router

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 2).

### 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.

#### 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 site-specific routing information. This includes controlling which prefixes can be advertised on an interface and what forwarding information can be sent to other routers out that 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 defined in the router's site identifiers. 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. This will increase the protocol processing time, but is necessary for connectivity within 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.

As an example, the router in Figure 3 must advertise routing information on four interfaces. The information advertised is as follows :

- Interface 1

```

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

```

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

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




Figure 3: Routing Information Exchange

- \* 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
- Lookup outgoing interface's site identifier
- Verify that the two site identifiers match

The last two steps are required due to the fact that the two interfaces could be in different sites but have identical site-local prefixes.

If a packet being forwarded contains a site-local source address and a globally-scoped destination address, there are two possible ways of handling it. The first way is to forward the packet based solely on the destination address. This leads to the possibility that the destination cannot send a response. The second option is to perform the same checks as outlined above for a site-local destination address. If this method is chosen, a new ICMPv6 message could be returned to the sender indicating there is a scoping mismatch. This would allow the sender the chance to use a higher scoped address.



## [5. Site-Boundary Multicast Routing](#)

### [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 sites defined on the box. Protocols that rely on underlying unicast protocols will not suffer as much of a performance impact since the unicast protocol will handle the forwarding table generation. However, multicast protocols that generate and maintain their own routing tables will have to perform the additional route calculations for each site.

### [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
- Site-scoped multicast destination address / Global unicast source address
- Site-scoped multicast destination address / Site-local unicast source address

The first combination should follow the same forwarding mechanisms that are in place today for IPv4 multicast. Combinations 3 and 4 should result in the router performing the same site identifiers check as outlined for site-local unicast destination addresses. Combination 2 could be handled in either manner. By performing the site identifier check, a source could be notified that there is a scoping mismatch and give it the opportunity to choose a higher scoped source address. This would require the definition of a new scoping mismatch ICMPv6 message.

## [6. Protocol Impact](#)

The performance impact on routing protocols is obvious. However, this impact should only be felt by those routers that exist on site boundaries. Realistically, a router would probably only be a

boundary between either two sites or a site

Haberman

Expires October 1998

[Page 6]

---

Internet Draft

Routing of Site-Scoped Addresses

April 1998

and the Internet. If site-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 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 for his 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](mailto:haberman@raleigh.ibm.com)

