

OSPF
Internet-Draft
Intended status: Standards Track
Expires: March 01, 2014

F.J. Baker
Cisco Systems
August 28, 2013

IPv6 Source/Destination Routing using OSPFv3
draft-baker-ipv6-ospf-dst-src-routing-03

Abstract

This note describes the changes necessary for OSPFv3 to route IPv6 traffic from a specified prefix to a specified prefix.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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 inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 01, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	2
2.	Theory of Routing	2
2.1.	Notation	3
2.2.	Dealing with ambiguity	3
2.3.	Interactions with other constraints	4
3.	Extensions necessary for IPv6 Source/Destination Routing in OSPFv3	4
3.1.	IPv6 Source Prefix TLV	4
4.	IANA Considerations	5
5.	Security Considerations	5
6.	Acknowledgements	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	6
Appendix A.	Change Log	6
	Author's Address	6

[1.](#) Introduction

This specification builds on OSPF for IPv6 [[RFC5340](#)] and the extensible LSAs defined in [[I-D.acee-ospfv3-lsa-extend](#)]. It defines the TLV for an IPv6 [[RFC2460](#)] Source Prefix, to define routes from a source prefix to a destination prefix.

This implies not simply routing "to a destination", but routing "to that destination AND from a specified source". It may be combined with other qualifying attributes, such as "traffic going to that destination AND using a specified flow label AND from a specified source prefix". The obvious application is egress routing, as required for a multihomed entity with a provider-allocated prefix from each of several upstream networks. Traffic within the network could be source/destination routed as well, or could be implicitly or explicitly routed from "any prefix", `::/0`. Other use cases are described in [[I-D.baker-rtgwg-src-dst-routing-use-cases](#)]. If a FIB contains a route to a given destination from one or more prefixes not including `::/0`, and a given packet destined there that has a source address that is in none of them, the packet in effect has no route, just as if the destination itself were not in the route table.

[1.1.](#) Requirements Language

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 [[RFC2119](#)].

[2.](#) Theory of Routing

Baker

Expires March 01, 2014

[Page 2]

Both IS-IS and OSPF perform their calculations by building a lattice of routers and links from the router performing the calculation to each router, and then use routes (sequences in the lattice) to get to destinations that those routes advertise connectivity to. Following the SPF algorithm, calculation starts by selecting a starting point (typically the router doing the calculation), and successively adding {link, router} pairs until one has calculated a route to every router in the network. As each router is added, including the original router, destinations that it is directly connected to are turned into routes in the route table: "to get to 2001:db8::/32, route traffic to {interface, list of next hop routers}". For immediate neighbors to the originating router, of course, there is no next hop router; traffic is handled locally.

In this context, the route is qualified by a source prefix; It is installed into the FIB with the destination prefix, and the FIB applies the route if and only if the IPv6 source address also matches the advertised prefix. Of course, there may be multiple LSAs in the LSDB with the same destination and differing source prefixes; these may also have the same or differing next hop lists. The intended forwarding action is to forward matching traffic to one of the next hop routers associated with this destination and source prefix, or to discard non-matching traffic as "destination unreachable".

LSAs that lack a source prefix TLV match any source address (i.e., the source prefix TLV defaults to ::/0), by definition.

2.1. Notation

For the purposes of this document, a route from the prefix A to the prefix B (in other words, whose source prefix is A and whose destination prefix is B) is expressed as A->B. A packet with the source address A and the destination address B is similarly described as A->B.

2.2. Dealing with ambiguity

In any routing protocol, there is the possibility of ambiguity. For example, one router might advertise a fairly general prefix - a default route, a discard prefix (which consumes all traffic that is not directed to an instantiated subnet), or simply an aggregated prefix while another router advertises a more specific one. In source/destination routing, potentially ambiguous cases include cases in which the link state database contains two routes A->B' and A'->B, in which A' is a more specific prefix within the prefix A and B' is a more specific prefix within the prefix B. Traditionally, we have dealt with ambiguous destination routes using a "longest match first" rule. If the same datagram matches more than one destination prefix

Baker

Expires March 01, 2014

[Page 3]

advertised within an area, we follow the route with the longest matching prefix.

With source/destination routes, as noted in [[I-D.baker-rtgwg-src-dst-routing-use-cases](#)], we follow a similar but slightly different rule; the FIB lookup MUST yield the route with the longest matching destination prefix that also matches the source prefix constraint. In the event of a tie on the destination prefix, it MUST also match the longest matching source prefix among those options.

An example of the issue is this. Suppose we have two routes:

1. 2001:db8:1::/48 -> 2001:db8:3:3::/64
2. 2001:db8:2::/48 -> 2001:db8:3::/48

and a packet

2001:db8:2::1 -> 2001:db8:3:3::1

If we require the algorithm to follow the longest destination match without regard to the source, the destination address matches 2001:db8:3:3::/64 (the first route), and the source address doesn't match the constraint of the first route; we therefore have no route. The FIB algorithm, in this example, must therefore match the second route, even though it is not the longest destination match, because it also matches the source address.

2.3. Interactions with other constraints

In the event that there are other constraints on routing, such as proposed in [[I-D.baker-ipv6-ospf-dst-flowlabel-routing](#)], the effect is a logical AND. The FIB lookup must yield the route with the longest matching destination prefix that also matches each of the constraints.

3. Extensions necessary for IPv6 Source/Destination Routing in OSPFv3

The extensible LSA format defined in [[I-D.acee-ospfv3-lsa-extend](#)] requires one additional option to accomplish source/destination routing: the source prefix. This is defined here. As noted in [Section 2](#), any prefix LSA that does not specify a source prefix is understood to as specifying ::/0 as the source prefix.

3.1. IPv6 Source Prefix TLV

[I-D.acee-ospfv3-lsa-extend]
Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", [draft-acee-ospfv3-lsa-extend-00](#) (work in progress), May 2013.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), July 2008.

[7.2. Informative References](#)

- [I-D.baker-ipv6-ospf-dst-flowlabel-routing]
Baker, F., "Using OSPFv3 with Role-Based Access Control", [draft-baker-ipv6-ospf-dst-flowlabel-routing-02](#) (work in progress), May 2013.
- [I-D.baker-rtgwg-src-dst-routing-use-cases]
Baker, F., "Requirements and Use Cases for Source/Destination Routing", [draft-baker-rtgwg-src-dst-routing-use-cases-00](#) (work in progress), August 2013.

[Appendix A. Change Log](#)

Initial Version: February 2013

First revision: April 2013

Correction: Corrected the reference to [[I-D.acee-ospfv3-lsa-extend](#)]

Use Case: Remove appendices, clarify the discussion of ambiguity and refer to [[I-D.baker-rtgwg-src-dst-routing-use-cases](#)]

Author's Address

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com

