Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2015

F.J. Baker
Cisco Systems
D. Lamparter
NetDEF
October 20, 2014

IPv6 Source/Destination Routing using IS-IS
draft-baker-ipv6-isis-dst-src-routing-02

Abstract

This note describes the changes necessary for IS-IS 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 April 23, 2015.

Copyright Notice

Table of Contents

1. Introduction

This specification builds on IS-IS for IPv6 [RFC5308] and the critical extension TLV in [critical-subtlvs]. This note defines the sub-TLV for an IPv6 [RFC2460] Source Prefix, to define routes from a source prefix to a destination prefix.

This implements the Destination/Source Routing mechanism described in [dst-src-routing]. 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

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 LSPs in the RIB 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".

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

When resolving Destination/Source Reachabilities, the SPF calculation results used MUST reflect a calculation performed including only routers that advertise support for the critical Source Prefix TLV

defined in section 3. The mechanism for signaling this is described in [critical-subtlvs]. Routers that support this extension MUST advertise support as described there.

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 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-isis-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. The general mechanics for this are described in [extra-qualifiers].

2.4. Multi-topology Routing

While not mandatory, IS-IS is often implemented as Multi Topology Routing [RFC5120] with IPv4 or other protocols in the same or different topologies. The TLV structure in [critical-subtlvs] is topology-agnostic in that it always includes the topology ID, which may be zero to indicate the default topology.

The mechanism in this document and its Sub-TLV are applicable to any topology that carries routing information used for IPv6 Unicast routing. Destination/Source reachability information SHOULD NOT be placed differently from "plain" destination reachabilities.

A system MUST NOT originate Destination/Source Reachabilities in a topology that is exclusively configured for multicast RPF operation. If a topology is shared between unicast lookups and multicast reverse path lookups, reachabilities with a source prefix other than ::/0 MUST be ignored for multicast reverse path lookups.

The statements in the previous two paragraphs currently result in applicability of Destination/Source routes as:

MT-ID	designated usage	applicability
0	default topology	yes
1	IPv4 management	no
2	IPv6 default	yes
3	IPv4 multicast	no
4	IPv6 multicast	no
5	IPv6 management	yes

Applicability of Destination/Source IPv6 Reachabilities

3. Extensions necessary for IPv6 Source/Destination Routing in IS-IS

Section 2 of [RFC5308] defines the "IPv6 Reachability TLV", and carries in it destination prefix advertisements. It has the capability of extension, using sub-TLVs.

We define the Source Prefix Sub-TLV as in Section 3.1. As noted in Section 2, any IPv6 Reachability TLV that does not specify a source prefix is understood to as specifying ::/0 (any IPv6 address) as the source prefix.

3.1. Source Prefix sub-TLV

The following Sub-TLV is defined for the critical part of TLV TBD2 defined in [critical-subtlvs]:

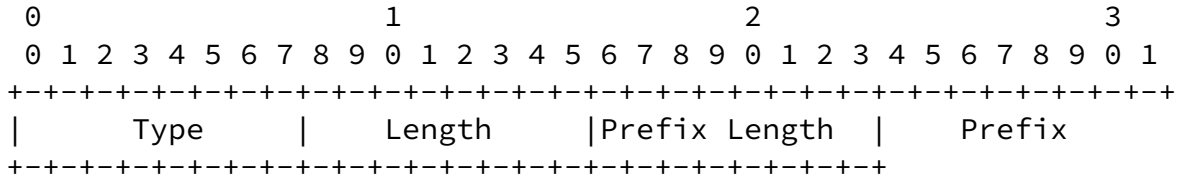

Source Prefix Sub-TLV

Source Prefix Type: assigned by IANA

TLV Length: Length of the sub-TLV in octets

Prefix Length: Length of the prefix in bits

Prefix: (source prefix length+7)/8 octets of prefix

4. IANA Considerations

The "Sub-TLVs for TLVs TBD1 (critical) and TBD2 (critical)" registry defined in [critical-subtlvs] is extended by the following element:

Source Prefix Type: assigned by IANA

Description: IPv6 Source Prefix

Applicable to TLV TBD1 (IPv4): No

Applicable to TLV TBD2 (IPv6): Yes

5. Security Considerations

There are no security considerations specific to this document. However, the considerations from [dst-src-routing] and [critical-subtlvs] are particularly relevant to this document.

6. Acknowledgements

7. References

7.1. Normative References

[IS-IS] ISO/IEC, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 2002.

[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.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.

[critical-subtlvs]
Lamparter, D., "IS-IS Reachability with critical Sub-TLVs", draft-lamparter-isis-reachability-critical-subtlvs-00 (work in progress), October 2014.

[dst-src-routing]
Lamparter, D., "Destination/Source Routing", draft-lamparter-rtgwg-dst-src-routing-00 (work in progress), October 2014.

[extra-qualifiers]
Lamparter, D., "Considerations and Registry for extending IP route lookup", draft-lamparter-rtgwg-routing-discriminators-00 (work in progress), October 2014.

7.2. Informative References

[I-D.baker-ipv6-isis-dst-flowlabel-routing]
Baker, F., "Using IS-IS with Token-based Access Control", draft-baker-ipv6-isis-dst-flowlabel-routing-01 (work in progress), August 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

updated Version: August 2013

Added MTR: August 2014

Split into 4 drafts: October 2014

Authors' Addresses

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com

David Lamparter
NetDEF
Leipzig 04103
Germany

Email: david@opensourcerouting.org