

Routing IPv6 with IS-IS

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies a method for exchanging IPv6 routing information using the IS-IS routing protocol. The described method utilizes two new TLVs: a reachability TLV and an interface address TLV to distribute the necessary IPv6 information throughout a routing domain. Using this method, one can route IPv6 along with IPv4 and OSI using a single intra-domain routing protocol.

1. Overview

IS-IS is an extendible intra-domain routing protocol. Each router in the routing domain issues an Link State Protocol Data Unit (LSP) that contains information pertaining to that router. The LSP contains typed variable-length data, often referred to as TLVs (type-length-values). We extend the protocol with two new TLVs to carry information required to perform IPv6 routing.

In [RFC1195], a method is described to route both OSI and IPv4. We utilize this same method with some minor changes to allow for IPv6. To do so, we must define two new TLVs, namely "IPv6 Reachability" and "IPv6 Interface Address", and a new IPv6 protocol identifier. In our new TLVs, we utilize the extended metrics and up/down semantics of [RFC5305].

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 RFC 2119 [RFC2119].

2. IPv6 Reachability TLV

The "IPv6 Reachability" TLV is TLV type 236 (0xEC).

[RFC1195] defines two Reachability TLVs, "IP Internal Reachability Information" and "IP External Reachability Information". We provide the equivalent IPv6 data with the "IPv6 Reachability" TLV and an "external" bit.

The "IPv6 Reachability" TLV describes network reachability through the specification of a routing prefix, metric information, a bit to indicate if the prefix is being advertised down from a higher level, a bit to indicate if the prefix is being distributed from another routing protocol, and OPTIONALLY the existence of Sub-TLVs to allow for later extension. This data is represented by the following structure:

```

0                                     1                                     2                                     3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type = 236  |      Length      |      Metric ..      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      .. Metric      |U|X|S| Reserve |  Prefix Len  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Prefix ...      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Sub-TLV Len(*) | Sub-TLVs(*) ...
* - if present

```

U - up/down bit
X - external original bit
S - subtlv present bit

The above IPv6 Reachability TLV MAY appear any number of times (including none) within an LSP. Link-local prefixes MUST NOT be advertised using this TLV.

As is described in [RFC5305]: "The up/down bit SHALL be set to 0 when a prefix is first injected into IS-IS. If a prefix is advertised from a higher level to a lower level (e.g. level 2 to level 1), the bit SHALL be set to 1, indicating that the prefix has traveled down the hierarchy. Prefixes that have the up/down bit set to 1 may only be advertised down the hierarchy, i.e., to lower levels".

If the prefix was distributed into IS-IS from another routing protocol, the external bit SHALL be set to 1. This information is useful when distributing prefixes from IS-IS to other protocols.

If the Sub-TLV bit is set to 0, then the octets of Sub-TLVs are not present. Otherwise, the bit is 1 and the octet following the prefix will contain the length of the Sub-TLV portion of the structure.

The prefix is "packed" in the data structure. That is, only the required number of octets of prefix are present. This number can be computed from the prefix length octet as follows:

prefix octets = integer of ((prefix length + 7) / 8)

Just as in [RFC5305], if a prefix is advertised with a metric larger than MAX_V6_PATH_METRIC (0xFE000000), this prefix MUST not be considered during the normal Shortest Path First (SPF) computation. This will allow advertisement of a prefix for purposes other than building the normal IPv6 routing table.

If Sub-TLVs are present, they have the same form as normal TLVs, as shown below.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      |      Value(*) ..
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
* - if present

```

Length indicates how many octets of value are present and can be 0.

3. IPv6 Interface Address TLV

The "IPv6 Interface Address" TLV is TLV type 232 (0xE8).

TLV 232 maps directly to "IP Interface Address" TLV in [RFC1195]. We necessarily modify the contents to be 0-15 16-octet IPv6 interface addresses instead of 0-63 4-octet IPv4 interface addresses.

```

0                                     1                                     2                                     3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type = 232  |      Length      |  Interface Address 1(*) ..  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     .. Interface Address 1(*) ..  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     .. Interface Address 1(*) ..  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     .. Interface Address 1(*) ..  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Interface Address 1(*) ..  |  Interface Address 2(*) ..  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
* - if present

```

We further restrict the semantics of this TLV depending on where it is advertised. For Hello PDUs, the "Interface Address" TLV MUST contain only the link-local IPv6 addresses assigned to the interface that is sending the Hello. For LSPs, the "Interface Address" TLVs MUST contain only the non-link-local IPv6 addresses assigned to the IS.

4. IPv6 NLPID

The value of the IPv6 Network Layer Protocol ID (NLPID) is 142 (0x8E).

As with [RFC1195] and IPv4, if the IS supports IPv6 routing using IS-IS, it MUST advertise this in the "NLPID" TLV by adding the IPv6 NLPID.

5. Operation

We utilize the same changes to [RFC1195] as made in [RFC5305] for the processing of prefix information. These changes are both related to the SPF calculation.

Since the metric space has been extended, we need to redefine the MAX_PATH_METRIC (1023) from the original specification in [RFC1195]. This new value MAX_V6_PATH_METRIC is the same as in [RFC5305] (0xFE000000). If, during the SPF, a path metric would exceed MAX_V6_PATH_METRIC, it SHALL be considered to be MAX_V6_PATH_METRIC.

The order of preference between paths for a given prefix MUST be modified to consider the up/down bit. The new order of preference is as follows (from best to worst).

1. Level 1 up prefix
2. Level 2 up prefix
3. Level 2 down prefix
4. Level 1 down prefix

If multiple paths have the same best preference, then selection occurs based on metric. Any remaining multiple paths SHOULD be considered for equal-cost multi-path routing if the router supports this; otherwise, the router can select any one of the multiple paths.

6. IANA Considerations

IANA has updated the IS-IS codepoint registry so that TLV codes 232 and 236 refer to this RFC.

IANA has also created the following new codepoint registry for Sub-TLVs of TLV 236. The range of values for Type is 0-255. Allocations within the registry require documentation of the use and requires approval by the Designated Expert assigned by the IESG [RFC5226]. All codepoints are currently unassigned.

7. Security Considerations

This document raises no new security considerations. Security considerations for the IS-IS protocol are covered in [ISO10589] and in [RFC5304].

8. References

8.1. Normative References

- [ISO10589] ISO, "Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", International Standard 10589:2002, Second Edition, 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, October 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

Author's Address

Christian E. Hopps
Cisco Systems
170 W. Tasman Dr.
San Jose, California 95134
USA

EMail: chopps@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

