Internet Engineering Task Force
INTERNET-DRAFT
Expires June 2003

Christian E. Hopps
Allegro Networks
8 January 2003

Routing IPv6 with IS-IS

<draft-ietf-isis-ipv6-05.txt>

Status of this Memo

Copyright Notice

Abstract

This draft specifies a method for exchanging IPv6 routing information using the IS-IS routing protocol. The described method utilizes 2 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. Terms

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. Overview

IS-IS [0] is an extendible intra-domain routing protocol. Each router in the routing domain issues an 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 2 new TLVs to carry information required to perform IPv6 routing.

In [1] 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 2 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 [2].

3. IPv6 Reachability TLV

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

[1] defines 2 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
```

This structure MAY appear any number of times (including none) within the TLV.

As is described in [2], "the up/down bit is set to 0 when a prefix is first injected into IS-IS. If a prefix is redistributed from a higher level to a lower level (e.g., level two to level one), the bit SHALL be set to 1 to indicate that the prefix has travelled down the hierarchy. If a prefix is redistributed from an area to another area at the same level then the up/down bit SHALL be set to 1."

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 [2], if a prefix is advertised with a metric larger than MAX_V6_PATH_METRIC (0xFE000000), this prefix MUST not be considered during the normal 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.

4. IPv6 Interface Address TLV

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

This TLV maps directly to [1]'s "IP Interface Address" TLV. We necessarily modify the contents to be 0-15 16 octet IPv6 interface addresses instead of 0-63 4 octet IPv4 interface address.

```
 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 "Interfaces Address" TLV MUST contain only the link-local IPv6 addresses assigned to the interface which is sending the Hello. For LSPs the "Interfaces Address" TLVs MUST contain only the non-link-local IPv6 addresses assigned to the IS.

5. IPv6 NLPID

The value of the IPv6 NLPID is 142 (0x8E).

As with [1] and IPv4, if the IS supports IPv6 routing using IS-IS, it MUST advertise this in the "NLPID" TLV by adding the IPv6 NLPID.

6. Operation

We utilize the same changes to [1] as made in [2] 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 [1]. This new value MAX_V6_PATH_METRIC is the same as in [2] (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.

7. Security Considerations

This document raises no new security considerations.

8. References

[0] "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO 10589, 1992.

[1] Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", RFC 1195, December 1990.

[2] Smit, H., and T. Li, "IS-IS extensions for Traffic Engineering", Work in Progress, August 2001.

9. Author's Address

Christian E. Hopps
Allagro Networks
6399 San Ignacio Ave.
San Jose, CA 95119
U.S.A.
Phone: +1 408 281 5500
Email: chopps@allegronetworks.com

10. Full Copyright Statement