

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: December 31, 2020

T. Li
S. Chen
V. Ilangoan
Arista Networks
G. Mishra
Verizon Inc.
June 29, 2020

Area Proxy for IS-IS
draft-ietf-lsr-isis-area-proxy-00

Abstract

Link state routing protocols have hierarchical abstraction already built into them. However, when lower levels are used for transit, they must expose their internal topologies to each other, leading to scale issues.

To avoid this, this document discusses extensions to the IS-IS routing protocol that would allow level 1 areas to provide transit, yet only inject an abstraction of the level 1 topology into level 2. Each level 1 area is represented as a single level 2 node, thereby enabling greater scale.

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 <https://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 December 31, 2020.

Copyright Notice

Copyright (c) 2020 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 (<https://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	3
1.1.	Requirements Language	4
2.	Area Proxy	4
2.1.	Segment Routing	5
3.	Inside Router Functions	6
3.1.	The Area Proxy Router Capability	6
3.2.	Level 2 SPF Computation	6
3.3.	The Inside Node TLV	7
4.	Area Leader Functions	7
4.1.	Area Leader Election	7
4.2.	Redundancy	8
4.3.	The Area Proxy TLV	8
4.3.1.	The Area Proxy System Id Sub-TLV	8
4.3.2.	The Area Segment SID Sub-TLV	9
4.4.	Area Proxy LSP Generation	10
4.4.1.	The Protocols Supported TLV	10
4.4.2.	The Area Address TLV	10
4.4.3.	The Dynamic Hostname TLV	10
4.4.4.	The IS Neighbors TLV	10
4.4.5.	The Extended IS Neighbors TLV	11
4.4.6.	The MT Intermediate Systems TLV	11
4.4.7.	Reachability TLVs	11
4.4.8.	The Router Capability TLV	12
4.4.9.	The Multi-Topology TLV	12
4.4.10.	The SID/Label Binding and The Multi-Topology SID/Label Binding SID TLV	13
4.4.11.	The Area Segment SID TLV	13
4.4.11.1.	Flags	14
4.4.12.	The SRv6 Locator TLV	14
4.4.13.	Traffic Engineering Information	14
5.	Inside Edge Router Functions	14
5.1.	Generating L2 IIHs to Outside Routers	15
5.2.	Filtering LSP information	15
6.	Acknowledgments	15
7.	IANA Considerations	16
8.	Security Considerations	16

9.	References	17
9.1.	Normative References	17
9.2.	Informative References	18
9.3.	URIs	19
	Authors' Addresses	19

[1.](#) Introduction

The IS-IS routing protocol IS-IS [[IS010589](#)] currently supports a two-level hierarchy of abstraction. The fundamental unit of abstraction is the 'area', which is a (hopefully) connected set of systems running IS-IS at the same level. Level 1, the lowest level, is abstracted by routers that participate in both Level 1 and Level 2, and they inject area information into Level 2. Level 2 systems seeking to access Level 1, use this abstraction to compute the shortest path to the Level 1 area. The full topology database of Level 1 is not injected into Level 2, only a summary of the address space contained within the area, so the scalability of the Level 2 Link State Database (LSDB) is protected.

This works well if the Level 1 area is tangential to the Level 2 area. This also works well if there are several routers in both Level 1 and Level 2 and they are adjacent, so Level 2 traffic will never need to transit Level 1 only routers. Level 1 will not contain any Level 2 topology, and Level 2 will only contain area abstractions for Level 1.

Unfortunately, this scheme does not work so well if the Level 1 only area needs to provide transit for Level 2 traffic. For Level 2 shortest path first (SPF) computations to work correctly, the transit topology must also appear in the Level 2 LSDB. This implies that all routers that could provide transit, plus any links that might also provide Level 2 transit must also become part of the Level 2 topology. If this is a relatively tiny portion of the Level 1 area, this is not overly painful.

However, with today's data center topologies, this is problematic. A common application is to use a Layer 3 Leaf-Spine (L3LS) topology, which is a folded 3-stage Clos [[Clos](#)] fabric. It can also be thought of as a complete bipartite graph. In such a topology, the desire is to use Level 1 to contain the routing dynamics of the entire L3LS topology and then to use Level 2 for the remainder of the network. Leaves in the L3LS topology are appropriate for connection outside of the data center itself, so they would provide connectivity for Level 2. If there are multiple connections to Level 2 for redundancy, or other areas, these too would also be made to the leaves in the topology. This creates a difficulty because there are now multiple

Level 2 leaves in the topology, with connectivity between the leaves provided by the spines.

Following the current rules of IS-IS, all spine routers would necessarily be part of the Level 2 topology, plus all links between a Level 2 leaf and the spines. In the limit, where all leaves need to support Level 2, it implies that the entire L3LS topology becomes part of Level 2. This is seriously problematic as it more than doubles the LSDB held in the L3LS topology and eliminates any benefits of the hierarchy.

This document discusses the handling of IP traffic. Supporting MPLS based traffic is a subject for future work.

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 [BCP 14](#) [1] [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Area Proxy

To address this, we propose to completely abstract away the details of the Level 1 area topology within Level 2, making the entire area look like a single proxy system directly connected to all of the area's Level 2 neighbors. By only providing an abstraction of the topology, Level 2's requirement for connectivity can be satisfied without the full overhead of the area's internal topology. It then becomes the responsibility of the Level 1 area to ensure the forwarding connectivity that's advertised.

For this discussion, we'll consider a single Level 1 IS-IS area to be the Inside Area, and the remainder of the Level 2 area is the Outside Area. All routers within the Inside Area speak Level 1 and Level 2 IS-IS on all of the links within the topology. We propose to implement Area Proxy by having a Level 2 Proxy Link State Protocol Data Unit (PDU, LSP) that represents the entire Inside Area. This is the only LSP from the area that will be flooded into the overall Level 2 LSDB.

There are four classes of routers that we need to be concerned with in this discussion:

Inside Router A router within the Inside Area that runs Level 1 and Level 2 IS-IS. A router is recognized as an Inside Router by the existence of its LSP in the Level 1 LSDB.

Area Leader The Area Leader is an Inside Router that is elected to represent the Level 1 area by injecting the Proxy LSP into the Level 2 LSDB. There may be multiple candidates for Area Leader, but only one is elected at a given time.

Inside Edge Router An Inside Edge Router is an Inside Area Router that has at least one Level 2 interface outside of the Inside Area. An interface on an Inside Edge Router that is connected to an Outside Edge Router is an Area Proxy Boundary.

Outside Edge Router An Outside Edge Router is a Level 2 router that is outside of the Inside Area that has an adjacency with an Inside Edge Router.

All Inside Edge Routers learn the Area Proxy System Identifier from the Level 1 LSDB and use that as the system identifier in their Level 2 IS-IS Hello PDUs (IIHs) on all Outside interfaces. Outside Edge Routers should then advertise an adjacency to the Area Proxy System Identifier. This allows all Outside Routers to use the Proxy LSP in their SPF computations without seeing the full topology of the Inside Area.

Area Proxy functionality assumes that all circuits on Inside Routers are either Level 1-2 circuits within the Inside Area, or Level 2 circuits between Outside Edge Routers and Inside Edge Routers.

Area Proxy Boundary multi-access circuits (i.e. Ethernet in LAN mode) with multiple Inside Edge Routers on them are not supported. The Inside Edge Router on any boundary LAN MUST NOT flood Inside Router LSPs on this link. Boundary LANs SHOULD NOT be enabled for Level 1. An Inside Edge Router may be elected the DIS for a Boundary LAN. In this case using the Area Proxy System Id as the basis for the LAN pseudonode identifier could create a collision, so the Inside Edge Router SHOULD compose the pseudonode identifier using its native system identifier.

2.1. Segment Routing

If the Inside Area supports Segment Routing [[RFC8402](#)], then all Inside Nodes MUST advertise an SR Global Block (SRGB). The first value of the SRGB advertised by all Inside Nodes MUST start at the same value. The range advertised for the area will be the minimum of all Inside Nodes.

To support Segment Routing, the Area Leader will take the global SID information found in the L1 LSDB and convey that to L2 through the Proxy LSP. Prefixes with SID assignments will be copied to the Proxy

LSP. Adjacency SIDs for Outside Edge Nodes will be copied to the Proxy LSP.

To further extend Segment Routing, it would be helpful to have a SID that refers to the entire Inside Area. This allows a path to refer to an area and have any node within that area accept and forward the packet. In effect, this becomes an anycast SID that is accepted by all Inside Edge Nodes. The information about this SID is distributed in the Area Segment SID Sub-TLV, as part of the Area Leader's Area Proxy TLV ([Section 4.3.2](#)). The Inside Edge Nodes MUST establish forwarding based on this SID. The Area Leader SHALL also include the Area Segment SID TLV in the Area Proxy LSP so that the remainder of L2 can use it for path construction ([Section 4.4.11](#)). These two TLVs are similar in structure, so care must be taken not to confuse them.

3. Inside Router Functions

All Inside Routers run Level 1-2 IS-IS and must be explicitly instructed to enable the Area Proxy functionality. To signal their readiness to participate in Area Proxy functionality, they will advertise the Area Proxy Router Capability as part of its Level 1 Router Capability TLV.

3.1. The Area Proxy Router Capability

The Area Proxy Router Capability is a sub-TLV of the Router Capability TLV [[RFC7981](#)] and has the following format:

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
| TLV Type      | TLV Length  |
+---+---+---+---+---+---+---+---+

```

TLV Type: LLL

TLV Length: 0

A router advertising this TLV indicates that it is running Level 1-2 and is prepared to perform Area Proxy functions.

3.2. Level 2 SPF Computation

When Outside Routers perform a Level 2 SPF computation, they will use the Area Proxy LSP for computing a path transiting the Inside Area. Because the topology has been abstracted away, the cost for transiting the Inside Area will be zero.

When Inside Routers perform a Level 2 SPF computation, they MUST ignore the Area Proxy LSP. Further, because these systems do see the Inside Area topology, the link metrics internal to the area are visible. This could lead to different and possibly inconsistent SPF results, potentially leading to forwarding loops.

To prevent this, the Inside Routers MUST consider the metrics of links outside of the Inside Area (inter-area metrics) separately from the metrics of the Inside Area links (intra-area metrics). Intra-area metrics MUST be treated as less than any inter-area metric. Thus, if two paths have different total inter-area metrics, the path with the lower inter-area metric would be preferred, regardless of any intra-area metrics involved. However, if two paths have equal inter-area metrics, then the intra-area metrics would be used to compare the paths.

Point-to-Point links between two Inside Routers are considered to be Inside Area links. LAN links which have a pseudonode LSP in the Level 1 LSDB are considered to be Inside Area links.

3.3. The Inside Node TLV

To simplify determining which nodes belong to the Inside Area, all Inside Nodes MUST insert the Inside Node TLV into their LSP and into any Inside Area pseudonode LSPs. The format of the Inside Node TLV is:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|           Type           | Length |
+---+---+---+---+---+---+---+---+

```

Type: ZZZ

Length: Zero (0)

4. Area Leader Functions

The Area Leader has several responsibilities. First, it MUST inject the Area Proxy System Identifier into the Level 1 LSDB. Second, the Area Leader MUST generate the Proxy LSP for the Inside Area.

4.1. Area Leader Election

The Area Leader is selected using the election mechanisms and TLVs described in Dynamic Flooding for IS-IS [[I-D.ietf-lsr-dynamic-flooding](#)].

4.2. Redundancy

If the Area Leader fails, another candidate may become Area Leader and MUST regenerate the Area Proxy LSP. The failure of the Area Leader is not visible outside of the area and appears to simply be an update of the Area Proxy LSP.

For consistency, all Area Leader candidates SHOULD be configured with the same Proxy System Id, Proxy Hostname, and any other information that may be inserted into the Proxy LSP.

4.3. The Area Proxy TLV

The Area Proxy TLV is a container for sub-TLVs with Area Proxy Information. This TLV is injected into the Area Leader's Level 1 LSP.

The format of the Area Proxy TLV is:

```

      0                   1                   2
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+
| TLV Type      | TLV Length    | Sub-TLVs ...
+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type: YYY

TLV Length: length of the sub-TLVs

4.3.1. The Area Proxy System Id Sub-TLV

The Area Proxy System Id Sub-TLV MUST be used by the Area Leader to distribute the Area Proxy System Id. This is an additional system identifier that is used by Inside Nodes. The format of this sub-TLV is:

```

      0                   1                   2
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length    |      Proxy System ID      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|      Proxy System Identifier continued
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: AAA

Length: length of a system ID (6)

Proxy System Identifier: the Area Proxy System Identifier.

The Area Leader SHOULD advertise the Area Proxy System Identifier Sub-TLV when it observes that all Inside Routers are advertising the Area Proxy Router Capability. Their advertisements indicate that they are individually ready to perform Area Proxy functionality. The Area Leader then advertises the Area Proxy System Identifier TLV to indicate that the Inside Area SHOULD enable Area Proxy functionality.

Other candidates for Area Leader MAY also advertise the Area Proxy System Identifier when they observe that all Inside Routers are advertising the Area Proxy Router Capability. All candidates advertising the Area Proxy System Identifier TLV MUST be advertising the same system identifier. Multiple proxy system identifiers in a single area is a misconfiguration and each unique occurrence SHOULD be logged.

The Area Leader and other candidates for Area Leader MAY withdraw the Area Proxy System Identifier when one or more Inside Routers are not advertising the Area Proxy Router Capability. This will disable Area Proxy functionality. However, before withdrawing the Area Proxy System Identifier, an implementation SHOULD protect against unnecessary churn from transients by delaying the withdrawal. The amount of delay is implementation-dependent.

4.3.2. The Area Segment SID Sub-TLV

The Area Segment SID Sub-TLV allows the Area Leader to advertise a SID that represents the entirety of the Inside Area to the Outside Area. This sub-TLV is learned by all of the Inside Edge Nodes who should consume this SID at forwarding time. The Area Segment SID Sub-TLV has the format:

```

      0               1               2
      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   |   Flags   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SID/Index/Label (variable)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: BBB

Length: variable (1 + SID length)

Flags: 1 octet, see [Section 4.4.11.1](#)

SID/Index/Label: as defined in [\[RFC8667\] Section 2.1.1.1](#)

[4.4.](#) Area Proxy LSP Generation

Each Inside Router generates a Level 2 LSP, and the Level 2 LSPs for the Inside Edge Routers will include adjacencies to Outside Edge Routers. Unlike normal Level 2 operations, these LSPs are not advertised outside of the Inside Area and MUST be filtered by all Inside Edge Routers to not be flooded to Outside Routers. Only the Area Proxy LSP is injected into the overall Level 2 LSDB.

The Area Leader uses the Level 2 LSPs generated by the Inside Edge Routers to generate the Area Proxy LSP. This LSP is originated using the Area Proxy System Identifier. The Area Leader MAY also insert the following additional TLVs into the Area Proxy LSP for additional information for the Outside Area. LSPs generated by unreachable nodes MUST NOT be considered.

[4.4.1.](#) The Protocols Supported TLV

The Area Leader SHOULD insert a Protocols Supported TLV (129) [\[RFC1195\]](#) into the Area Proxy LSP. The values included in the TLV SHOULD be the protocols supported by the Inside Area.

[4.4.2.](#) The Area Address TLV

The Area Leader SHOULD insert an Area Addresses TLV (1) [\[ISO10589\]](#) into the Area Proxy LSP.

[4.4.3.](#) The Dynamic Hostname TLV

It is RECOMMENDED that the Area Leader insert the Dynamic Hostname TLV (137) [\[RFC5301\]](#) into the Area Proxy LSP. The contents of the hostname may be specified by configuration. The presence of the hostname helps to simplify debugging the network.

[4.4.4.](#) The IS Neighbors TLV

The Area Leader MAY insert the IS Neighbors TLV (2) [\[ISO10589\]](#) into the Area Proxy LSP for Outside Edge Routers. The Area Leader learns of the Outside Edge Routers by examining the LSPs generated by the Inside Edge Routers copying any IS Neighbors TLVs referring to Outside Edge Routers into the Proxy LSP. Since the Outside Edge Routers advertise an adjacency to the Area Proxy System Identifier, this will result in a bi-directional adjacency.

An entry for a neighbor in both the IS Neighbors TLV and the Extended IS Neighbors would be functionally redundant, so the Area Leader SHOULD NOT do this.

4.4.5. The Extended IS Neighbors TLV

The Area Leader MAY insert the Extended IS Reachability TLV (22) [[RFC5305](#)] into the Area Proxy LSP. The Area Leader SHOULD copy each Extended IS Reachability TLV advertised by an Inside Edge Router about an Outside Edge Router into the Proxy LSP.

If the Inside Area supports Segment Routing and Segment Routing selects a SID where the L-Flag is unset, then the Area Lead SHOULD include an Adjacency Segment Identifier sub-TLV (31) [[RFC8667](#)] using the selected SID.

If the inside area supports SRv6, the Area Leader SHOULD copy the "SRv6 End.X SID" and "SRv6 LAN End.X SID" sub-TLVs of the extended IS reachability TLVs advertised by Inside Edge Routers about Outside Edge Routers.

If the inside area supports Traffic Engineering (TE), the Area Leader SHOULD copy TE related sub-TLVs [[RFC5305](#) Section 3] to each Extended IS Reachability TLV in the Proxy LSP.

4.4.6. The MT Intermediate Systems TLV

If the Inside Area supports Multi-Topology, then the Area Leader SHOULD copy each Outside Edge Router advertisement that is advertised by an Inside Edge Router in a MT Intermediate Systems TLV into the Proxy LSP.

4.4.7. Reachability TLVs

The Area Leader SHOULD insert additional TLVs describing any routing prefixes that should be advertised on behalf of the area. These prefixes may be learned from the Level 1 LSDB, Level 2 LSDB, or redistributed from another routing protocol. This applies to all of various types of TLVs used for prefix advertisement:

IP Internal Reachability Information TLV (128) [[RFC1195](#)]

IP External Reachability Information TLV (130) [[RFC1195](#)]

Extended IP Reachability TLV (135) [[RFC5305](#)]

IPv6 Reachability TLV (236) [[RFC5308](#)]

Multi-Topology Reachable IPv4 Prefixes TLV (235) [[RFC5120](#)]

Multi-Topology Reachable IPv6 Prefixes TLV (237) [[RFC5120](#)]

For TLVs in the Level 1 LSDB, for a given TLV type and prefix, the Area Leader SHOULD select the TLV with the lowest metric and copy that TLV into the Area Proxy LSP.

When examining the Level 2 LSDB for this function, the Area Leader SHOULD only consider TLVs advertised by Inside Routers. Further, for prefixes that represent Boundary links, the Area Leader SHOULD copy all TLVs that have unique sub-TLV contents.

If the Inside Area supports Segment Routing and the selected TLV includes a Prefix Segment Identifier sub-TLV (3) [[RFC8667](#)], then the sub-TLV SHOULD be copied as well. The P-Flag SHOULD be set in the copy of the sub-TLV to indicate that penultimate hop popping SHOULD NOT be performed for this prefix. The E-Flag SHOULD be reset in the copy of the sub-TLV to indicate that an explicit NULL is not required. The R-Flag SHOULD simply be copied.

4.4.8. The Router Capability TLV

The Area Leader MAY insert the Router Capability TLV (242) [[RFC7981](#)] into the Area Proxy LSP. If Segment Routing is supported by the inside area, as indicated by the presence of an SRGB being advertised by all Inside Nodes, then the Area Leader SHOULD advertise an SR-Capabilities sub-TLV (2) [[RFC8667](#)] with an SRGB. The first value of the SRGB is the same value as the first value advertised by all Inside Nodes. The range advertised for the area will be the minimum of all ranges advertised by Inside Nodes. The Area Leader SHOULD use its own Router Id in the Router Capability TLV.

If SRv6 Capability sub-TLV [[RFC7981](#)] is advertised by all Inside Routers, the Area Leader should insert an SRv6 Capability sub-TLV in the Router Capability TLV. Each flag in the SRv6 Capability sub-TLV should be set if the flag is set by all Inside Routers.

If the Node Maximum SID Depth (MSD) sub-TLV [[RFC8491](#)] is advertised by all Inside Routers, the Area Leader should advertise common MSD types and the smallest supported MSD values for each type.

4.4.9. The Multi-Topology TLV

If the Inside Area supports multi-topology, then the Area Leader SHOULD insert the Multi-Topology TLV (229) [[RFC5120](#)], including the topologies supported by the Inside Nodes.

If any Inside Node is advertising the 'O' (Overload) bit for a given topology, then the Area Leader MUST advertise the 'O' bit for that topology. If any Inside Node is advertising the 'A' (Attach) bit for a given topology, then the Area Leader MUST advertise the 'A' bit for that topology.


4.4.10. The SID/Label Binding and The Multi-Topology SID/Label Binding SID TLV

If an Inside Node advertises the SID/Label Binding or Multi-Topology SID/Label Binding SID TLV [[RFC8667](#)], then the Area Leader MAY copy the TLV to the Area Proxy LSP.

4.4.11. The Area Segment SID TLV

If the Area Leader is advertising an Area Segment SID in the Area Segment SID sub-TLV of the Area Proxy TLV, then the Area Leader SHOULD advertise the Area Segment SID TLV in the Proxy LSP. The advertisement in the Proxy LSP informs the remainder of the network that packets directed to the SID will be forwarded by one of the Inside Edge Nodes and the Area Segment SID will be consumed.

This TLV is not specific to Area Proxy and MAY be used by Edge Routers in conventional areas. The Area Segment SID TLV has the format:



where:

Type: XXX

Length: variable (1 + SID length)

Flags: 1 octet, see below

SID/Index/Label: as defined in [\[RFC8667\] Section 2.1.1.1](#)

4.4.11.1. Flags

The Flags octet is defined as follows:

```

0 1 2 3 4 5 6 7
+--+--+--+--+--+
|F|V|L|      |
+--+--+--+--+--+

```

where:

F: Address-Family Flag. If unset, then this proxy SID is used when forwarding IPv4-encapsulated traffic. If set, then this proxy SID is used when forwarding IPv6-encapsulated traffic.

V: Value Flag. If set, then the proxy SID carries a value.

L: Local Flag. If set, then the value/index carried by the proxy SID has local significance.

Other bits: MUST be zero when originated and ignored when received.

4.4.12. The SRv6 Locator TLV

If the inside area supports SRv6, the Area Leader SHOULD copy all SRv6 locator TLVs [[I-D.ietf-lsr-isis-srv6-extensions](#)] advertised by Inside Routers to the Proxy LSP.

4.4.13. Traffic Engineering Information

If the inside area supports TE, the Area Leader SHOULD advertise a TE Router ID TLV (134) [[RFC5305](#)] in the Proxy LSP. It SHOULD copy the Shared Risk Link Group (SRLS) TLVs (138) [[RFC5307](#)] advertised by Inside Edge Routers about links to Outside Edge Routers.

If the inside area supports IPv6 TE, the Area Leader SHOULD advertise an IPv6 TE Router ID TLV (140) [[RFC6119](#)] in the Proxy LSP. It SHOULD also copy the IPv6 SRLG TLVs (139) [[RFC6119](#)] advertised by Inside Edge Routers about links to Outside Edge Routers.

5. Inside Edge Router Functions

The Inside Edge Router has two additional and important functions. First, it MUST generate IIHs that appear to have come from the Area Proxy System Identifier. Second, it MUST filter the L2 LSPs, Partial Sequence Number PDUs (PSNPs), and Complete Sequence Number PDUs (CSNPs) that are being advertised to Outside Routers.

5.1. Generating L2 IIHs to Outside Routers

The Inside Edge Router has one or more Level 2 interfaces to Outside Routers. These may be identified by explicit configuration or by the fact that they are not also Level 1 circuits. On these Level 2 interfaces, the Inside Edge Router MUST NOT send an IIH until it has learned the Area Proxy System Id from the Area Leader. Then, once it has learned the Area Proxy System Id, it MUST generate its IIHs on the circuit using the Proxy System Id as the source of the IIH.

Using the Proxy System Id causes the Outside Router to advertise an adjacency to the Proxy System Id, not to the Inside Edge Router, which supports the proxy function. The normal system id of the Inside Edge Router MUST NOT be used as it will cause unnecessary adjacencies to form and subsequently flap.

5.2. Filtering LSP information

For the proxy abstraction to be effective the L2 LSPs generated by the Inside Routers MUST be restricted to the Inside Area. The Inside Routers know which system ids are members of the Inside Area based on the Level 1 LSDB. To prevent unwanted LSP information from escaping the Inside Area, the Inside Edge Router MUST perform filtering of LSP flooding, CSNPs, and PSNPs. Specifically:

A Level 2 LSP with a source system identifier that is found in the Level 1 LSDB MUST never be flooded to an Outside Router.

A Level 2 CSNP sent to an Outside Router MUST NOT contain any information about an LSP with a system identifier found in the Level 1 LSDB. If an Inside Edge Router filters a CSNP and there is no remaining content, then the CSNP MUST NOT be sent. The source address of the CSNP MUST be the Area Proxy System Id.

A Level 2 PSNP sent to an Outside Router MUST NOT contain any information about an LSP with a system identifier found in the Level 1 LSDB. If an Inside Edge Router filters a PSNP and there is no remaining content, then the PSNP MUST NOT be sent. The source address of the PSNP MUST be the Area Proxy System Id.

6. Acknowledgments

The authors would like to thank Bruno Decraene, Vivek Ilangovan, and Gunter Van De Velde for their many helpful comments. The authors would also like to thank a small group that wishes to remain anonymous for their valuable contributions.

7. IANA Considerations

This memo requests that IANA allocate and assign one code point from the IS-IS TLV Codepoints registry for the Area Segment SID TLV (XXX), one code point for the Area Proxy TLV (YYY), and one code point for the Inside Node TLV (ZZZ). The registry fields for all three should be: IIH:n, LSP:y, SNP:n, Purge:n.

In association with this, this memo requests that IANA create a registry for code points for the sub-TLVs of the Area Proxy TLV.

Name of the registry: Sub-TLVs for TLV YYY (Area Proxy TLV)

Required information for registrations: Temporary registrations may be made under the Early IANA Allocation of Standards Track Code Points policy. [[RFC7120](#)] Permanent registrations require the publication of an RFC describing the usage of the code point.

Applicable registration policy: RFC Required and Expert Review. We propose the initial experts be Chris Hopps, Tony Li, and Sarah Chen.

Size, format, and syntax of registry entries: Value (0-255), Name, and Reference

Initial assignments and reservations: IANA is requested to assign the following code points:

Value	Name	Reference
AAA	Area Proxy System Identifier	This document
BBB	Area Segment SID	This document

IANA is also requested to allocate and assign one code point from the IS-IS Router Capability TLV sub-TLV registry for the Area Proxy Capability (LLL).

8. Security Considerations

This document introduces no new security issues. Security of routing within a domain is already addressed as part of the routing protocols themselves. This document proposes no changes to those security architectures.

9. References

9.1. Normative References

- [I-D.ietf-lsr-dynamic-flooding]
Li, T., Psenak, P., Ginsberg, L., Chen, H., Przygienda, T., Cooper, D., Jalil, L., Dontula, S., and G. Mishra, "Dynamic Flooding on Dense Graphs", [draft-ietf-lsr-dynamic-flooding-07](#) (work in progress), June 2020.
- [I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", [draft-ietf-lsr-isis-srv6-extensions-08](#) (work in progress), April 2020.
- [ISO10589]
International Organization for Standardization, "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, Nov. 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5301] McPherson, D. and N. Shen, "Dynamic Hostname Exchange Mechanism for IS-IS", [RFC 5301](#), DOI 10.17487/RFC5301, October 2008, <<https://www.rfc-editor.org/info/rfc5301>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 5307](#), DOI 10.17487/RFC5307, October 2008, <<https://www.rfc-editor.org/info/rfc5307>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", [RFC 6119](#), DOI 10.17487/RFC6119, February 2011, <<https://www.rfc-editor.org/info/rfc6119>>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", [RFC 7981](#), DOI 10.17487/RFC7981, October 2016, <<https://www.rfc-editor.org/info/rfc7981>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", [RFC 8491](#), DOI 10.17487/RFC8491, November 2018, <<https://www.rfc-editor.org/info/rfc8491>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [RFC 8667](#), DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

9.2. Informative References

- [Clos] Clos, C., "A Study of Non-Blocking Switching Networks", The Bell System Technical Journal Vol. 32(2), DOI 10.1002/j.1538-7305.1953.tb01433.x, March 1953, <<http://dx.doi.org/10.1002/j.1538-7305.1953.tb01433.x>>.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 7120](#), DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.

9.3. URIs

[1] <https://tools.ietf.org/html/bcp14>

Authors' Addresses

Tony Li
Arista Networks
5453 Great America Parkway
Santa Clara, California 95054
USA

Email: tony.li@tony.li

Sarah Chen
Arista Networks
5453 Great America Parkway
Santa Clara, California 95054
USA

Email: sarahchen@arista.com

Vivek Ilangovan
Arista Networks
5453 Great America Parkway
Santa Clara, California 95054
USA

Email: ilangovan@arista.com

Gyan S. Mishra
Verizon Inc.

13101 Columbia Pike

Silver Spring

,

MD 20904

United States of America

Phone:

301 502-1347

Email:

gyan.s.mishra@verizon.com

