

Internet Engineering Task Force (IETF)
Request for Comments: 9666
Category: Experimental
ISSN: 2070-1721

T. Li
Juniper Networks
S. Chen
V. Ilango
Arista Networks
G. Mishra
Verizon Inc.
October 2024

Area Proxy for IS-IS

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, thereby leading to scaling issues.

To avoid such issues, this document discusses extensions to the IS-IS routing protocol that allow Level 1 (L1) areas to provide transit but only inject an abstraction of the Level 1 topology into Level 2 (L2). Each Level 1 area is represented as a single Level 2 node, thereby enabling a greater scale.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9666>.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
 - 1.1. Requirements Language

- 2. Area Proxy
 - 2.1. Segment Routing
- 3. Inside Router Functions
 - 3.1. The Area Proxy TLV
 - 3.2. Level 2 SPF Computation
 - 3.3. Responsibilities Concerning the Proxy LSP
- 4. Area Leader Functions
 - 4.1. Area Leader Election
 - 4.2. Redundancy
 - 4.3. Distributing Area Proxy Information
 - 4.3.1. The Area Proxy System Identifier Sub-TLV
 - 4.3.2. The Area SID Sub-TLV
 - 4.4. Proxy LSP Generation
 - 4.4.1. The Protocols Supported TLV
 - 4.4.2. The Area Address TLV
 - 4.4.3. The Dynamic Hostname TLV
 - 4.4.4. The IS Neighbors TLV
 - 4.4.5. The Extended IS Neighbors TLV
 - 4.4.6. The MT Intermediate Systems TLV
 - 4.4.7. Reachability TLVs
 - 4.4.8. The Router Capability TLV
 - 4.4.9. The Multi-Topology TLV
 - 4.4.10. The SID/Label Binding and the Multi-Topology SID/Label Binding TLV
 - 4.4.11. The SRv6 Locator TLV
 - 4.4.12. Traffic Engineering Information
 - 4.4.13. The Area SID
- 5. Inside Edge Router Functions
 - 5.1. Generating L2 IIHs to Outside Routers
 - 5.2. Filtering LSP Information
- 6. IANA Considerations
- 7. Security Considerations
- 8. References
 - 8.1. Normative References
 - 8.2. Informative References
- Acknowledgements
- Authors' Addresses

1. Introduction

The IS-IS routing protocol [ISO10589] 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, rather, only a summary of the address space contained within the area is injected. Therefore, 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 Levels 1 and 2 and they are adjacent to one another, 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 fabric [Clos]. 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 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 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", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Area Proxy

In this specification, we 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 provide 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 to be 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 PDU (LSP) that represents the entire Inside Area. We will refer to this as the Proxy LSP. 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. Any Inside Router can be

the Area Leader.

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.




Figure 1: An Example of Router Classes

All Inside Edge Routers learn the Area Proxy System Identifier from the Area Proxy TLV advertised by the Area Leader and use that as the system identifier in their Level 2 IS-IS Hello (IIH) PDUs on all Outside interfaces. Outside Edge Routers will 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., Ethernets 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 as the Designated Intermediate System (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 Insider Edge Router SHOULD compose the pseudonode identifier using its originally configured system identifier. This choice of pseudonode identifier may confuse neighbors with an extremely strict implementation. In this case, the Inside Edge Router may be configured with priority 0, causing an Outside Router to be elected as the DIS.

2.1. Segment Routing

If the Inside Area supports Segment Routing (SR) [RFC8402], then all Inside Nodes MUST advertise a Segment Routing Global Block (SRGB). The first value of the SRGB advertised by all Inside Nodes MUST start at the same value. If the Area Leader detects SRGBs that do not start with the same value, it MUST log an error and not advertise an SRGB in the Proxy LSP. The range advertised for the area will be the minimum of that advertised by all Inside Nodes.

To support SR, the Area Leader will take the SRGB information found in the L1 LSDB and convey that to L2 through the Proxy LSP. Prefixes with Segment Identifier (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 SR, it is helpful to have a segment 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 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 SID in the Proxy LSP so that the remainder of L2 can use it for path construction. (Section 4.4.13).

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 TLV in their L2 LSP.

3.1. The Area Proxy TLV

The Area Proxy TLV serves multiple functions:

- * The presence of the Area Proxy TLV in a node's LSP indicates that the node is enabled for Area Proxy.
- * An LSP containing the Area Proxy TLV is also an Inside Node. All Inside Nodes, including pseudonodes, MUST advertise the Area Proxy TLV.
- * It is a container for sub-TLVs with Area Proxy information.

A node advertises the Area Proxy TLV in fragment 0 of its L2 LSP. Nodes MUST NOT advertise the Area Proxy TLV in an L1 LSP. Nodes MUST ignore the Area Proxy TLV if it is found in an L1 LSP. The Area Proxy TLV is not used in the Proxy 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: 20

TLV Length: Length of the sub-TLVs.

3.2. Level 2 SPF Computation

When Outside Routers perform a Level 2 SPF computation, they will use the 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 Proxy LSP. Because these systems 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 that have a pseudonode LSP in the Level 1 LSDB are considered to be Inside Area links.

3.3. Responsibilities Concerning the Proxy LSP

The Area Leader will generate a Proxy LSP that will be flooded across the Inside Area. Inside Routers MUST flood the Proxy LSP and MUST ignore its contents. The Proxy LSP uses the Area Proxy System Identifier as its Source ID.

4. Area Leader Functions

The Area Leader has several responsibilities. First, it MUST inject the Area Proxy System Identifier into the Level 2 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 on Dense Graphs" [RFC9667].

4.2. Redundancy

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

The Area Leader is responsible for distributing information about the area to all Inside Nodes. In particular, the Area Leader distributes the Proxy System ID and the Area SID. This is done using two sub-TLVs of the Area Proxy TLV.

4.3.1. The Area Proxy System Identifier Sub-TLV

The Area Proxy System Identifier 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 as an indication that Area Proxy is active. 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 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      | Proxy System Identifier |
+-----+-----+-----+-----+-----+-----+-----+-----+
|
```

Type: 1

Length: Length of a system ID (6).

Proxy System Identifier: The Area Proxy System Identifier.

The Area Leader MUST advertise the Area Proxy System Identifier sub-TLV when it observes that all Inside Routers are advertising the Area Proxy TLV. 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 MUST 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 TLV. All candidates advertising the Area Proxy System Identifier TLV SHOULD be advertising the same system identifier. Multiple proxy system identifiers in a single area is a misconfiguration and each unique occurrence SHOULD be logged. Systems should use the Proxy System ID advertised by the Area Leader.

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 TLV. 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 SID Sub-TLV

The Area SID sub-TLV allows the Area Leader to advertise a prefix and SID that represent 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 SID sub-TLV has the following 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)
+-----+-----+-----+-----+-----+-----+-----+
| Prefix Length | Prefix (variable)
+-----+-----+-----+-----+-----+-----+-----+
```

where:

Type: 2

Length: Variable (1 + SID length)

Flags: 1 octet, defined as follows.

```

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

F: Address-Family Flag. If this flag is not set, 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, as

defined in [RFC8667], Section 2.1.1.1.

L: Local Flag. If set, then the value/index carried by the proxy SID has local significance, as defined in [RFC8667], Section 2.1.1.1.

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

SID/Index/Label: As defined in [RFC8667], Section 2.1.1.1.

Prefix Length: 1 octet

Prefix: 0-16 octets

4.4. 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 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 Proxy LSP. This LSP is originated using the Area Proxy System Identifier. The Area Leader can also insert the following additional TLVs into the 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 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 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 Proxy LSP. The contents of the hostname may be specified by configuration. The presence of the hostname helps to simplify network debugging.

4.4.4. The IS Neighbors TLV

The Area Leader can insert the IS Neighbors TLV (2) [ISO10589] into the 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 bidirectional adjacency.

An entry for a neighbor in both the IS Neighbors TLV and the Extended IS Neighbors TLV would be functionally redundant, so the Area Leader SHOULD NOT do this. The Area Leader MAY omit either the IS Neighbors TLV or the Extended IS Neighbors TLV, but it MUST include at least one of them.

4.4.5. The Extended IS Neighbors TLV

The Area Leader can insert the Extended IS Reachability TLV (22) [RFC5305] into the 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 not set, 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 (MT), then the Area Leader SHOULD copy each Outside Edge Router advertisement that is advertised by an Inside Edge Router in an 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 the 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 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 SR 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 Proxy LSP. If SR 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 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 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 the intersection of the advertised 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 TLV

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

4.4.11. The SRv6 Locator TLV

If the inside area supports SRv6, the Area Leader SHOULD copy all SRv6 locator TLVs [RFC9352] advertised by Inside Routers to the Proxy LSP.

4.4.12. 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.

4.4.13. The Area SID

When SR is enabled, it may be useful to advertise an Area SID that will direct traffic to any of the Inside Edge Routers. The information for the Area SID is distributed to all Inside Edge Routers using the Area SID sub-TLV (Section 4.3.2) by the Area Leader.

The Area Leader SHOULD advertise the Area SID information in the

Proxy LSP as a Node SID as defined in [RFC8667], Section 2.1. The advertisement in the Proxy LSP informs the Outside Area that packets directed to the SID will be forwarded to one of the Inside Edge Nodes and the Area SID will be consumed.

Other uses of the Area SID and Area SID prefix are outside the scope of this document. Documents that define other use cases for the Area SID MUST specify whether the SID value should be the same or different from that used in support of Area Proxy.

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 the 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.

5.2. Filtering LSP Information

For the area 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 advertisement of the Area Proxy TLV. 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 NOT be flooded to an Outside Router.
- * A Level 2 LSP that contains the Area Proxy TLV MUST NOT 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. IANA Considerations

IANA has assigned code point 20 from the "IS-IS TLV Codepoints"

registry for the Area Proxy TLV. The registry fields are I1H:n, LSP:y, SNP:n, and Purge:n.

In association with this, IANA has created a "IS-IS Sub-TLVs for the Area Proxy TLV" registry. Temporary registrations may be made via early allocation [RFC7120].

The registration procedure is Expert Review [RFC8126]. The values are from 0-255, and the fields are Value, Name, and Reference. The initial assignments are as follows.

Value	Name	Reference
1	Area Proxy System Identifier	RFC 9666
2	Area SID	RFC 9666

Table 1

7. 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. Security for IS-IS is provided by "IS-IS Cryptographic Authentication" [RFC5304] and "IS-IS Generic Cryptographic Authentication" [RFC5310].

8. References

8.1. Normative References

- [ISO10589] ISO, "Information technology - Telecommunications and information exchange between systems - 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)", Second Edition, ISO/IEC 10589:2002, November 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>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [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>>.

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>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [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>>.
- [RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.
- [RFC9667] Li, T., Ed., Psenak, P., Ed., Chen, H., Jalil, L., and S. Dontula, "Dynamic Flooding on Dense Graphs", RFC 9667, DOI 10.17487/RFC9667, October 2024, <<https://www.rfc-editor.org/info/rfc9667>>.

8.2. Informative References

- [Clos] Clos, C., "A study of non-blocking switching networks", The Bell System Technical Journal, Volume 32, Issue 2, pp. 406-424, DOI 10.1002/j.1538-7305.1953.tb01433.x, March 1953, <<https://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>>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

Acknowledgements

The authors would like to thank Bruno Decraene 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.

Authors' Addresses

Tony Li
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States of America
Email: tony.li@tony.li

Sarah Chen
Arista Networks
5453 Great America Parkway
Santa Clara, CA 95054
United States of America
Email: sarahchen@arista.com

Vivek Ilangovan
Arista Networks
5453 Great America Parkway
Santa Clara, CA 95054
United States of America
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