

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: May 2, 2020

T. Li
Arista Networks
L. Ginsberg
P. Wells
Cisco Systems
October 30, 2019

IS-IS Extended Hierarchy
draft-ietf-lsr-isis-extended-hierarchy-00

Abstract

The IS-IS routing protocol was originally defined with a two level hierarchical structure. This was adequate for the networks at the time. As we continue to expand the scale of our networks, it is apparent that additional hierarchy would be a welcome degree of flexibility in network design.

This document defines IS-IS Levels 3 through 8.

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 May 2, 2020.

Copyright Notice

Copyright (c) 2019 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	2
1.1.	Requirements Language	3
2.	PDU changes	3
2.1.	Circuit Type	3
2.2.	PDU Type	4
3.	Additional PDUs	4
3.1.	Level n LAN IS to IS hello PDU (Ln-LAN-HELLO-PDU)	4
3.2.	Level n Point-to-point IS to IS hello PDU (Ln-P2P-HELLO-PDU)	5
4.	Level Specific Area Identifiers	5
4.1.	IS-IS Area Identifier TLV	5
4.2.	Adjacency Formation Rules	6
4.2.1.	Level 3-8 Adjacency Formation Rules	7
4.2.2.	Special Level-2 Adjacency Formation Rules	7
5.	New Flooding Scopes	7
6.	MAC Addresses	8
7.	Inheritance of TLVs	9
8.	Behavior of Level n	9
9.	Relationship between levels	9
10.	Acknowledgements	9
11.	IANA Considerations	9
11.1.	PDU Type	9
11.2.	New PDUs	10
11.3.	New TLVs	10
11.4.	New Flooding Scopes	10
11.5.	New MAC Addresses	11
12.	Security Considerations	12
13.	Normative References	12
	Authors' Addresses	12

[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.

Practical considerations, such as the size of an area's link state database, cause network designers to restrict the number of routers in any given area. Concurrently, the dominance of scale-out

architectures based around small routers has created a situation where the scalability limits of the protocol are going to become critical in the foreseeable future.

The goal of this document is to enable additional hierarchy within IS-IS. Each additional level of hierarchy has a multiplicative effect on scale, so the addition of six levels should be a significant improvement. While all six levels may not be needed in the short term, it is apparent that the original designers of IS-IS reserved enough space for these levels, and defining six additional levels is only slightly harder than adding a single level, so it makes sense to expand the design for the future.

The modifications described herein are designed to be fully backward compatible and have no effect on existing networks. The modifications are also designed to have no effect whatsoever on networks that only use Level 1 and/or Level 2.

Section references in this document are references to sections of IS-IS [[IS010589](#)].

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. PDU changes

In this section, we enumerate all of the redefinitions of protocol header fields necessary to add additional levels.

2.1. Circuit Type

In the fixed header of some IS-IS PDUs, a field is named 'Reserved/Circuit Type' ([Section 9.5](#)). The high order six bits are reserved, with the low order two bits indicating Level 1 (bit 1) and Level 2 (bit 2).

This field is renamed to be 'Circuit Type'. The bits are redefined as follows:

1. Level 1
2. Level 2
3. Level 3

4. Level 4
5. Level 5
6. Level 6
7. Level 7
8. Level 8

The value of zero (no bits set) is reserved. PDUs with a Circuit Type of zero SHALL be ignored.

The set bits of the Circuit Type MUST be contiguous. If bit n and bit m are set in the Circuit Type, then all bits in the interval [n:m] must be set.

2.2. PDU Type

The fixed header of IS-IS PDUs contains an octet with three reserved bits and the 'PDU Type' field. The three reserved bits are transmitted as zero and ignored on receipt. ([Section 9.5](#))

To allow for additional PDU space, this entire octet is renamed the 'PDU Type' field.

3. Additional PDUs

3.1. Level n LAN IS to IS hello PDU (Ln-LAN-HELLO-PDU)

The 'Level n LAN IS to IS hello PDU' (Ln-LAN-HELLO-PDU) is identical in format to the 'Level 2 LAN IS to IS hello PDU' ([Section 9.6](#)), except that the PDU Types are defined as follows:

Level 3 (L3-LAN-HELLO-PDU): 33 (Suggested - to be assigned by IANA)

Level 4 (L4-LAN-HELLO-PDU): 34 (Suggested - to be assigned by IANA)

Level 5 (L5-LAN-HELLO-PDU): 35 (Suggested - to be assigned by IANA)

Level 6 (L6-LAN-HELLO-PDU): 36 (Suggested - to be assigned by IANA)

Level 7 (L7-LAN-HELLO-PDU): 37 (Suggested - to be assigned by IANA)

Level 8 (L8-LAN-HELLO-PDU): 38 (Suggested - to be assigned by IANA)

The Circuit Type field MUST be set to indicate all levels supported on that circuit.

3.2. Level n Point-to-point IS to IS hello PDU (Ln-P2P-HELLO-PDU)

The 'Point-to-point IS to IS hello PDU' ([Section 9.7](#)) is used on Level 1 and Level 2 circuits. Legacy systems will not expect the circuit type field to indicate other levels, so a new PDU is used if the circuit supports other levels. The additional PDU is the 'Level n Point-to-point IS to IS hello PDU' (Ln-P2P-HELLO-PDU) and has PDU Type 39 (Suggested - to be assigned by IANA). The format of this PDU is identical to the existing Point-to-Point IS to IS hello PDU. Both PDUs may be used on the same circuit.

4. Level Specific Area Identifiers

[ISO10589] defines an Area Address to uniquely identify a Level-1 area. A given area may have multiple synonymous area addresses - which is useful in support of hitless merging or splitting of areas. Area address matching is part of the adjacency formation rules defined in [Section 8](#) which determine whether a given adjacency supports Level-1, Level-2, or both. Area addresses are advertised in IIHs and LSPs using the Area Address TLV.

With the extensions defined in this document, there is a need to define an equivalent identifier for Levels 2-8. This identifier is a 32 bit value and is advertised using the new Area Identifier TLV defined in the following section. There is no relationship between the Level-1 Area Addresses and the new Level Specific Area Identifier.

Just as with Area Addresses, multiple synonymous Area Identifiers may be assigned to a given level. This supports hitless merging or splitting of the level specific area. Although it is legal to do so, it is generally not useful to define more than two Area Identifiers for a given level.

4.1. IS-IS Area Identifier TLV

The Area Identifier TLV is added to IS-IS to allow nodes to indicate which areas they participate in for Levels 2-8. Area Identifiers are locally administered 32 bit numbers. Each level may have multiple Area Identifiers. The format of the TLV is:


```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  TLV Type   | TLV Length   |    Level    |    Count    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Area Identifier                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

TLV Type: ZZZ

TLV Length: (2 * number of Levels) + (4 * each Count field)

Level: The level number of the area. (1 octet) Legal values are 2-8

Count: The number of Area Identifier fields (1 octet)

Area Identifier: One or more identifiers associated with the area. (4 * count octets)

The Level/Count/Area Identifier tuple MAY be repeated as necessary.

The Area Identifier TLV MAY appear in all types of IIHs except for Level 1 LAN IS to IS hellos.

The Area Identifier TLV MAY appear in LSP #0 of non-pseudo-node Level 3-8 Flooding Scoped LSPs defined in [Section 5](#). It MUST NOT be present in any LSP with non-zero LSP number. If present in an LSP with non-zero LSP number it MUST be ignored on receipt.

A system may participate in more than one level. At a given level, an area may have a number of synonymous identifiers. A system MUST advertise all of the levels it supports and the associated Area Identifiers.

[4.2](#). Adjacency Formation Rules

Adjacency formation rules for Levels 1 and 2 are defined in [\[ISO10589\]](#) and are not altered by these extensions except where noted below.

Adjacency Formation rules for Levels 3 and above are defined to insure that adjacency support for a given level is only enabled when there is a matching Area Identifier. Adjacency formation rules also are defined so as to prevent interconnection of neighbors which will connect to different areas at levels above any supported level.

4.2.1. Level 3-8 Adjacency Formation Rules

When the Area Identifier TLV appears in a Level N Point-to-point IS to IS hello PDU or a Level N LAN IS to IS Hello PDU, the Circuit Type field is inspected. For all levels with their corresponding bit set in the Circuit Type in the range 3-8 the following checks are performed:

- o Check for a matching Area Identifier at the same level
- o Check for a matching Area Identifier at all supported levels greater than the level being checked

If both checks pass, then an adjacency can be formed supporting the level. If any of the checks fail, then that level **MUST NOT** be supported by an adjacency formed on that circuit.

On a Point-to-Point circuit, a single adjacency is formed which supports all of the levels which pass the above checks.

On a LAN circuit, an adjacency is formed only for the level specified by the PDU type. Nevertheless, the checks for all levels with the corresponding bit set in the Circuit Type **MUST** be performed.

Note that (as previously specified) the set of levels supported **MUST** be contiguous.

4.2.2. Special Level-2 Adjacency Formation Rules

The Area Identifier TLV **MAY** appear in a Point-to-point IS to IS hello PDU or Level 2 LAN IS to IS Hello PDU (both specified in [[IS010589](#)]). In such a case, the neighbor may or may not support the Area Identifier TLV. If the Area Identifier TLV is present and Level 2 is indicated as being supported in the Circuit Type field, then in addition to the checks specified in [[IS010589](#)] the checks specified in the previous section **SHOULD** be performed for Level 2.

5. New Flooding Scopes

For levels 3-8, all link state information, PSNPs, and CSNPs are relayed in conformance with [RFC 7356](#) [[RFC7356](#)]. Additional flooding scopes are defined for each new level, for both circuit flooding scope and level flooding scope. Level flooding scopes are defined for both Standard and Extended TLV formats. The list of additional flooding scopes is:

Value	Description	FS LSP ID Format/ TLV Format
6	Level 3 Circuit Flooding Scope	Extended/Standard
7	Level 4 Circuit Flooding Scope	Extended/Standard
8	Level 5 Circuit Flooding Scope	Extended/Standard
9	Level 6 Circuit Flooding Scope	Extended/Standard
10	Level 7 Circuit Flooding Scope	Extended/Standard
11	Level 8 Circuit Flooding Scope	Extended/Standard
12	Level 3 Flooding Scope	Extended/Standard
13	Level 4 Flooding Scope	Extended/Standard
14	Level 5 Flooding Scope	Extended/Standard
15	Level 6 Flooding Scope	Extended/Standard
16	Level 7 Flooding Scope	Extended/Standard
17	Level 8 Flooding Scope	Extended/Standard
18	Level 3 Flooding Scope	Standard/Standard
19	Level 4 Flooding Scope	Standard/Standard
20	Level 5 Flooding Scope	Standard/Standard
21	Level 6 Flooding Scope	Standard/Standard
22	Level 7 Flooding Scope	Standard/Standard
23	Level 8 Flooding Scope	Standard/Standard
70	Level 3 Circuit Flooding Scope	Extended/Extended
71	Level 4 Circuit Flooding Scope	Extended/Extended
72	Level 5 Circuit Flooding Scope	Extended/Extended
73	Level 6 Circuit Flooding Scope	Extended/Extended
74	Level 7 Circuit Flooding Scope	Extended/Extended
75	Level 8 Circuit Flooding Scope	Extended/Extended
76	Level 3 Flooding Scope	Extended/Extended
77	Level 4 Flooding Scope	Extended/Extended
78	Level 5 Flooding Scope	Extended/Extended
79	Level 6 Flooding Scope	Extended/Extended
80	Level 7 Flooding Scope	Extended/Extended
81	Level 8 Flooding Scope	Extended/Extended

6. MAC Addresses

On a broadcast network, PDUs are currently sent to the AllL1Iss or AllL2Iss MAC addresses. We will need additional MAC addresses for Levels 3-8.

AllL3ISs: MAC3

AllL4ISs: MAC4

AllL5ISs: MAC5

AllL6ISs: MAC6

AllL7ISs: MAC7

AllL8ISs: MAC8

When operating in Point-to-Point mode on a broadcast network [[RFC5309](#)], a Level N Point-to-Point Hello PDU will be sent. Any of the above MAC addresses could be used in this case, but it is recommended to use the AllL3ISs MAC address.

7. Inheritance of TLVs

All existing Level 2 TLVs may be used in the corresponding Level 3 through Level 8 PDUs. When used in a Level 3 through Level 8 PDU, the semantics of these TLVs will be applied to the Level of the containing PDU. If the original semantics of the PDU was carrying a reference to Level 1 in a Level 2 TLV, then the semantics of the TLV at level N will be a reference to level N-1. The intent is to retain the original semantics of the TLV at the higher level.

8. Behavior of Level n

The behavior of Level n is analogous to the behavior of Level 2.

9. Relationship between levels

The relationship between Level n and Level n-1 is analogous to the relationship between Level 2 and Level 1.

An area at Level n has at most one parent at Level n+1.

10. Acknowledgements

The authors would like to thank Dinesh Dutt for inspiring this document and Huaimo Chen for his comments.

11. IANA Considerations

This document makes many requests to IANA, as follows:

11.1. PDU Type

The existing IS-IS PDU registry currently supports values 0-31. This should be expanded to support the values 0-255. The existing value assignments should be retained. Value 255 should be reserved.

11.2. New PDUs

IANA is requested to allocate values from the IS-IS PDU registry for the following:

L3-LAN-HELLO-PDU: 33 (Suggested - to be assigned by IANA)

L4-LAN-HELLO-PDU: 34 (Suggested - to be assigned by IANA)

L5-LAN-HELLO-PDU: 35 (Suggested - to be assigned by IANA)

L6-LAN-HELLO-PDU: 36 (Suggested - to be assigned by IANA)

L7-LAN-HELLO-PDU: 37 (Suggested - to be assigned by IANA)

L8-LAN-HELLO-PDU: 38 (Suggested - to be assigned by IANA)

Ln-P2P-HELLO-PDU: 39 (Suggested - to be assigned by IANA)

11.3. New TLVs

IANA is requested to allocate values from the IS-IS TLV registry for the following:

Area Identifier: ZZZ

11.4. New Flooding Scopes

IANA is requested to allocate the following values from the IS-IS Flooding Scope Identifier Registry.

Value	Description	FS LSP ID Format/ TLV Format	IIH Announce Lx-P2P Lx-LAN
6	Level 3 Circuit Flooding Scope	Extended/Standard	Y Y
7	Level 4 Circuit Flooding Scope	Extended/Standard	Y Y
8	Level 5 Circuit Flooding Scope	Extended/Standard	Y Y
9	Level 6 Circuit Flooding Scope	Extended/Standard	Y Y
10	Level 7 Circuit Flooding Scope	Extended/Standard	Y Y
11	Level 8 Circuit Flooding Scope	Extended/Standard	Y Y
12	Level 3 Flooding Scope	Extended/Standard	Y Y
13	Level 4 Flooding Scope	Extended/Standard	Y Y
14	Level 5 Flooding Scope	Extended/Standard	Y Y
15	Level 6 Flooding Scope	Extended/Standard	Y Y
16	Level 7 Flooding Scope	Extended/Standard	Y Y
17	Level 8 Flooding Scope	Extended/Standard	Y Y
18	Level 3 Flooding Scope	Standard/Standard	Y Y
19	Level 4 Flooding Scope	Standard/Standard	Y Y
20	Level 5 Flooding Scope	Standard/Standard	Y Y
21	Level 6 Flooding Scope	Standard/Standard	Y Y
22	Level 7 Flooding Scope	Standard/Standard	Y Y
23	Level 8 Flooding Scope	Standard/Standard	Y Y
70	Level 3 Circuit Flooding Scope	Extended/Extended	Y Y
71	Level 4 Circuit Flooding Scope	Extended/Extended	Y Y
72	Level 5 Circuit Flooding Scope	Extended/Extended	Y Y
73	Level 6 Circuit Flooding Scope	Extended/Extended	Y Y
74	Level 7 Circuit Flooding Scope	Extended/Extended	Y Y
75	Level 8 Circuit Flooding Scope	Extended/Extended	Y Y
76	Level 3 Flooding Scope	Extended/Extended	Y Y
77	Level 4 Flooding Scope	Extended/Extended	Y Y
78	Level 5 Flooding Scope	Extended/Extended	Y Y
79	Level 6 Flooding Scope	Extended/Extended	Y Y
80	Level 7 Flooding Scope	Extended/Extended	Y Y
81	Level 8 Flooding Scope	Extended/Extended	Y Y

11.5. New MAC Addresses

IANA is requested to allocate values from the IANA Multicast 48-bit MAC Addresses block for the following:

AllL3Iss: MAC3

AllL4Iss: MAC4

AllL5Iss: MAC5

AllL6Iss: MAC6

AllL7Iss: MAC7

AllL8Iss: MAC8

12. 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.

13. Normative References

[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.

[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>>.

[RFC5309] Shen, N., Ed. and A. Zinin, Ed., "Point-to-Point Operation
over LAN in Link State Routing Protocols", [RFC 5309](#),
DOI 10.17487/RFC5309, October 2008,
<<https://www.rfc-editor.org/info/rfc5309>>.

[RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding
Scope Link State PDUs (LSPs)", [RFC 7356](#),
DOI 10.17487/RFC7356, September 2014,
<<https://www.rfc-editor.org/info/rfc7356>>.

Authors' Addresses

Tony Li
Arista Networks
5453 Great America Parkway
Santa Clara, California 95054
United States of America

Email: tony.li@tony.li

Les Ginsberg
Cisco Systems
United States of America

Email: ginsberg@cisco.com

Paul Wells
Cisco Systems
United States of America

Email: pauwells@cisco.com