

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 7, 2019

T. Li
Arista Networks
L. Ginsberg
P. Wells
Cisco Systems
June 5, 2019

Hierarchical IS-IS
draft-li-lsr-isis-hierarchical-isis-01

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 December 7, 2019.

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)	4
4. IS-IS Area Identifier TLV	5
5. New Flooding Scopes	5
6. Inheritance of TLVs	6
7. Relationship between levels	7
8. Acknowledgements	7
9. IANA Considerations	7
9.1. PDU Type	7
9.2. New PDUs	7
9.3. New TLVs	7
9.4. New Flooding Scopes	8
10. Security Considerations	8
11. Normative References	9
Authors' Addresses	9

1. Introduction

The IS-IS routing protocol IS-IS [ISO10589] 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 [ISO10589].

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): AA3

Level 4 (L4-LAN-HELLO-PDU): AA4

Level 5 (L5-LAN-HELLO-PDU): AA5

Level 6 (L6-LAN-HELLO-PDU): AA6

Level 7 (L7-LAN-HELLO-PDU): AA7


Level 8 (L8-LAN-HELLO-PDU): AA8

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 TTT with the same format. Both PDUs may be used on the same circuit.

4. IS-IS Area Identifier TLV

The Area Identifier TLV is added to IS-IS to allow nodes to indicate which areas they participate in. Area Identifiers are locally administered 32 bit numbers. The format of the TLV is:



TLV Type: ZZZ

TLV Length: 7

Level: The level number of the area.

Area Identifier: The identifier associated with the area.

The Area Identifier TLV may appear in IIHs or in LSPs. When the Area Identifier TLV appears in a PDU, it indicates that the system is participating in the specified area at the indicated level. When the Area Identifier TLV appears in a IIH, the receiving system MUST NOT form an adjacency unless an Area Identifier TLV corresponds to the receiver's own Area Identifier for the given level.

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. 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.

7. Relationship between levels

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

8. Acknowledgements

The author would like to thank Dinesh Dutt for inspiring this document. The author would also like to thank Les Ginsberg and Paul Wells for their helpful comments.

9. IANA Considerations

This document makes many requests to IANA, as follows:

9.1. 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.

9.2. New PDUs

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

L3-LAN-HELLO-PDU: AA3

L4-LAN-HELLO-PDU: AA4

L5-LAN-HELLO-PDU: AA5

L6-LAN-HELLO-PDU: AA6

L7-LAN-HELLO-PDU: AA7

L8-LAN-HELLO-PDU: AA8

Ln-P2P-HELLO-PDU: TTT

To allow for PDU types to be defined independent of this document, the above values should be allocated from the range 32-254.

9.3. New TLVs

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

Area Identifier: ZZZ

9.4. 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

10. 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.

11. 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>>.
- [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