

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: December 30, 2018

T. Li  
Arista Networks  
June 28, 2018

Hierarchical IS-IS  
draft-li-hierarchical-isis-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 December 30, 2018.

## Copyright Notice

Copyright (c) 2018 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

Internet-Draft

Hierarchical IS-IS

June 2018

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">PDU changes</a>	<a href="#">3</a>
<a href="#">2.1.</a>	<a href="#">Circuit Type</a>	<a href="#">3</a>
<a href="#">2.2.</a>	<a href="#">PDU Type</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Additional PDUs</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">LAN IS to IS hello PDU (LAN-HELLO-PDU)</a>	<a href="#">4</a>
<a href="#">3.2.</a>	<a href="#">Point-to-point IS to IS hello PDU (P2P-HELLO-PDU)</a>	<a href="#">4</a>
<a href="#">3.3.</a>	<a href="#">Level n Link State PDU (Ln-LSP-PDU)</a>	<a href="#">4</a>
<a href="#">3.4.</a>	<a href="#">Level n complete sequence numbers PDU (Ln-CSNP-PDU)</a>	<a href="#">5</a>
<a href="#">3.5.</a>	<a href="#">Level n partial sequence numbers PDU (Ln-PSNP-PDU)</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Inheritance of TLVs</a>	<a href="#">5</a>
<a href="#">5.</a>	<a href="#">Acknowledgements</a>	<a href="#">6</a>
<a href="#">6.</a>	<a href="#">IANA Considerations</a>	<a href="#">6</a>
<a href="#">6.1.</a>	<a href="#">PDU Type</a>	<a href="#">6</a>
<a href="#">6.2.</a>	<a href="#">New PDUs</a>	<a href="#">6</a>
<a href="#">7.</a>	<a href="#">Security Considerations</a>	<a href="#">7</a>
<a href="#">8.</a>	<a href="#">Normative References</a>	<a href="#">7</a>
	<a href="#">Author's Address</a>	<a href="#">7</a>

## [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 by creating additional hierarchy. 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 some sense to expand the design for the future.

The modifications described herein are designed to be fully backward compatible.

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.](#) LAN IS to IS hello PDU (LAN-HELLO-PDU)

The 'LAN IS to IS hello PDU' (LAN-HELLO-PDU) is identical in format to the 'Level 2 LAN IS to IS hello PDU' ([Section 9.6](#)), except that the PDU Type has value AAA. The LAN-HELLO-PDU MUST be used instead of the 'Level 1 LAN IS to IS hello PDU' ([Section 9.5](#)) or the 'Level 2 LAN IS to IS hello PDU' on any circuit that has one or more of Level 3 through Level 8 enabled.

### [3.2.](#) Point-to-point IS to IS hello PDU (P2P-HELLO-PDU)

The 'Point-to-point IS to IS hello PDU' can be used on circuits of any Level without modification.

### [3.3.](#) Level n Link State PDU (Ln-LSP-PDU)

The 'Level n Link State PDU' (Ln-LSP-PDU) has the same format as the 'Level 2 Link State PDU' ([Section 9.9](#)), except for the PDU Type. The PDU Types for Levels 3 through 8 are defined as follows:

Level 3 (L3-LSP-PDU): BBB

Level 4 (L4-LSP-PDU): CCC

Level 5 (L5-LSP-PDU): DDD

Level 6 (L6-LSP-PDU): EEE

Level 7 (L7-LSP-PDU): FFF

Level 8 (L8-LSP-PDU): GGG

### [3.4.](#) Level n complete sequence numbers PDU (Ln-CSNP-PDU)

The 'Level n complete sequence numbers PDU' (Ln-CSNP-PDU) has the same format as the 'Level 2 complete sequence numbers PDU' ([Section 9.11](#)), except for the PDU Type. The PDU Types for Levels 3 through 8 are defined as follows:

Level 3 (L3-CSNP-PDU): HHH

Level 4 (L4-CSNP-PDU): III

Level 5 (L5-CSNP-PDU): JJJ

Level 6 (L6-CSNP-PDU): KKK

Level 7 (L7-CSNP-PDU): LLL

Level 8 (L8-CSNP-PDU): MMM

### [3.5.](#) Level n partial sequence numbers PDU (Ln-PSNP-PDU)

The 'Level 2 partial sequence numbers PDU' (Ln-PSNP-PDU) has the same

format as the 'Level 2 partial sequence numbers PDU' ([Section 9.13](#)), except for the PDU Type. The PDU Types for Levels 3 through 8 are defined as follows:

Level 3 (L3-PSNP-PDU): NNN

Level 4 (L4-PSNP-PDU): 000

Level 5 (L5-PSNP-PDU): PPP

Level 6 (L6-PSNP-PDU): QQQ

Level 7 (L7-PSNP-PDU): RRR

Level 8 (L8-PSNP-PDU): SSS

#### [4.](#) 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.

#### [5.](#) Acknowledgements

The author would like to thank Dinesh Dutt for inspiring this document.

#### [6.](#) IANA Considerations

This document makes many requests to IANA, as follows:

##### [6.1.](#) 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.

## [6.2.](#) New PDUs

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

LAN-HELLO-PDU: AAA

L3-LSP-PDU: BBB

L4-LSP-PDU: CCC

L5-LSP-PDU: DDD

L6-LSP-PDU: EEE

L7-LSP-PDU: FFF

L8-LSP-PDU: GGG

L3-CSNP-PDU: HHH

L4-CSNP-PDU: III

L5-CSNP-PDU: JJJ

L6-CSNP-PDU: KKK

L7-CSNP-PDU: LLL

L8-CSNP-PDU: MMM

L3-PSNP-PDU: NNN

L4-PSNP-PDU: 000

L5-PSNP-PDU: PPP

L6-PSNP-PDU: QQQ

L7-PSNP-PDU: RRR

L8-PSNP-PDU: SSS

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

## 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.

## 8. 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>>.

### Author's Address

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

Email: [tony.li@tony.li](mailto:tony.li@tony.li)