

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 30 December 2022

T. Li
Juniper Networks
L. Ginsberg
P. Wells
Cisco Systems
28 June 2022

IS-IS Extended Hierarchy
draft-ietf-lsr-isis-extended-hierarchy-06

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 30 December 2022.

Copyright Notice

Copyright (c) 2022 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	3
1.1. Requirements Language	3
2. PDU changes	4
2.1. Circuit Type	4
2.2. PDU Type	4
3. Additional PDUs	5
3.1. Level n LAN IS to IS hello PDU (Ln-LAN-HELLO-PDU)	5
3.2. Level n Point-to-point IS to IS hello PDU (Ln-P2P-HELLO-PDU)	5
4. Level Specific Area Identifiers	6
4.1. IS-IS Area Hierarchy TLV	6
4.2. Adjacency Formation Rules	8
4.2.1. Level 3-8 Adjacency Formation Rules	8
4.2.2. Special Level-1 and Level-2 Adjacency Formation Rules	9
4.2.2.1. Actions on a Point-to-Point Circuit	9
4.2.2.2. Actions on a LAN Circuit	9
4.2.2.3. Reporting of Mismatched Area Hierarchies	9
5. New Flooding Scopes	10
6. MAC Addresses	12
7. Inheritance of TLVs	13
8. Behavior of Level n	13
9. Relationship between levels	13
10. Acknowledgements	13
11. IANA Considerations	13
11.1. PDU Type	13
11.2. New PDUs	13
11.3. New TLVs	14
11.4. New Flooding Scopes	14
11.5. New MAC Addresses	15
12. Security Considerations	16
13. Normative References	16
Appendix A. Preventing Cross Branching in the Hierarchy	16
Appendix B. Guidelines for Introducing a new level	17
Authors' Addresses	18

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].

Note that [ISO10589] uses a bit encoding convention where bit numbers are 1 based and Bit 1 is the Least Significant Bit (LSB) of the datatype. Traditionally IETF documents have used a bit encoding convention where bit numbers are 0 based and Bit 0 is the Most Significant Bit (MSB) of the datatype. This document uses [ISO10589] conventions throughout.

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. 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 - not just the level associated with the containing PDU type.

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. The Level Specific Area Identifier (LSAI) is a 16 bit value and is advertised using the new Area Hierarchy TLV defined in Section 4.1. There is no relationship between a Level-1 Area Address and an LSAI.

Just as with Area Addresses, multiple synonymous LSAIs 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.

A node MAY support any set of contiguous levels. Support for non-contiguous levels is undefined.

4.1. IS-IS Area Hierarchy TLV

The Area Hierarchy TLV specifies the set of LSAIs which comprise the branch of the network hierarchy to which the advertising node is connected. The TLV MUST include at least one LSAI for Levels 2-N, where N is ≥ 2 and N represents the highest level supported in the IS-IS domain. It is RECOMMENDED that $N == 8$ even when not all 8 levels are currently in use, but in cases where a network does not support higher levels a number less than 8 MAY be used.

Note that the levels advertised MAY include levels which are not supported by the advertising node.

The Area Hierarchy TLV has the following format:

```

      8 7 6 5 4 3 2 1
+---+---+---+---+---+---+
|   TLV Type   |
+---+---+---+---+---+---+
| TLV Length   |
+---+---+---+---+---+---+
| Supp-Levels  |
+---+---+---+---+---+---+

```

Followed by one or more Level Specific Area ID Sets:

```

      1                               0
      6 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Level           | # of LSAIs |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Level Specific Area Id(s) |
...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type: ZZZ (1 octet)

TLV Length: Variable (1 octet)

Supp-Levels: A contiguous bitmask representing the set of levels supported by the advertising node (1 octet)

Bit #8 of this field is set if Level 8 is supported.

Bit #7 of this field is set if Level 7 is supported.

Bit #6 of this field is set if Level 6 is supported.

Bit #5 of this field is set if Level 5 is supported.

Bit #4 of this field is set if Level 4 is supported.

Bit #3 of this field is set if Level 3 is supported.

Bit #2 of this field is set if Level 2 is supported.

Bit #1 of this field is set if Level 1 is supported

If the Supp-level bit mask is non-contiguous all advertised LSAIs are ignored.

Each Level Specific Area ID Set consists of:

Level: 2-8 (1 octet)

of LSAI: >=1 (1 octet)

LSAIs: The set of synonomous LSAIs associated with this level
(2 * # of LSAIs octets)

The Area Hierarchy TLV MUST appear in all new IIH PDUs defined in Section 3. It MAY appear in P2P-HELLO-PDUs, L1-LAN-HELLO-PDUs, or L2-LAN-HELLO-PDUs.

The Area Hierarchy TLV MUST appear in LSP #0 of non-pseudo-node Level 3-8 Flooding Scoped LSPs defined in Section 5. It MAY appear in L1 or L2 LSP #0. 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.

Multiple Area Hierarchy TLVs MUST NOT be sent. In the event multiple Area Hierarchy TLVs are received, the first such TLV in the PDU is used. Subsequent TLVs in the same PDU MUST be ignored.

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.

The checks discussed below need to be performed on receipt of an IIH.

4.2.1. Level 3-8 Adjacency Formation Rules

The Area Hierarchy TLV MUST be present in a Level N Point-to-point IS to IS hello PDU or a Level N LAN IS to IS Hello PDU and the TLV content MUST adhere to the definition in Section 4.1. Beginning with the lowest level supported by the receiving node on this circuit and including all higher levels for which the receiver has an assigned LSAI regardless as to whether the higher levels are supported on this circuit, the set of LSAIs defined on the receiving node is compared against the set of LSAIs advertised in the received TLV. A matching LSAI MUST be found for each level.

If all of the checks pass then a new adjacency is formed or an existing adjacency is maintained.

NOTE: The absence of the advertisement of an LSAI for a given level is considered as a failure to find a matching LSAI.

On a Point-to-Point circuit, a single adjacency is formed which supports all of the levels supported by both nodes on this circuit.

On a LAN circuit, an adjacency is formed supporting only the level specified by the PDU type.

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

4.2.2. Special Level-1 and Level-2 Adjacency Formation Rules

The Area Hierarchy TLV MAY appear in a Point-to-point IS to IS hello PDU, Level 1 LAN IS to IS Hello PDU, or Level 2 LAN IS to IS Hello PDU (PDUs specified in [ISO10589]). In such a case, the neighbor may or may not support the Area Hierarchy TLV. The following sub-sections define modified adjacency formation rules for point-to-point and LAN circuits.

4.2.2.1. Actions on a Point-to-Point Circuit

If the Area Hierarchy TLV is present, then in addition to the checks specified in [ISO10589] the checks specified in Section 4.2.1 MUST be performed for all levels for which the receiver has an assigned LSAI beginning with Level 2. If those checks fail an adjacency MUST NOT be formed and any existing matching adjacency MUST transition to DOWN state.

4.2.2.2. Actions on a LAN Circuit

Adjacency formation MUST follow the rules defined in [ISO10589]. If the Area Hierarchy TLV is present in the Level 1 or Level 2 LAN IS to IS Hello PDU then the checks specified in Section 4.2.1 SHOULD be performed for all levels for which the receiver has an assigned LSAI beginning with Level 2. If those checks fail an error SHOULD be reported, but the level specific adjacency is still allowed. This prevents violation of the assumption of transitivity on the LAN in the presence of systems which do not support the extensions defined in this document.

4.2.2.3. Reporting of Mismatched Area Hierarchies

When forming adjacencies at Level-1 and/or Level-2, it is possible to have a mixture of legacy nodes (which do NOT support the extensions defined in this document) and new nodes which do support the extensions.

In Point-to-Point mode, legacy nodes will not advertise the new Area Hierarchy TLV and will not have an assigned LSAI for Level-2. It then becomes possible for new nodes with mismatched Area Hierarchies to form adjacencies with legacy nodes and form an L1 or L2 area where not all new nodes have a matching Area Hierarchy. This cannot be detected when forming adjacencies if the new nodes are not directly connected - but it can be detected after the adjacencies have been formed by inspecting the set of Area Hierarchy TLVs in the level specific LSPs of all routers in the area.

Similarly in LAN mode, the transitivity requirement means that new nodes MUST form adjacencies with all nodes connected to the LAN even when the Area Hierarchy TLV mismatch check fails (see Section 4.2.2.2). This can occur both at Level-1 and Level-2.

New nodes MUST report these inconsistencies.

5. New Flooding Scopes

For levels 3-8, all link state information, PSNPs, and CSNPs are relayed in conformance with [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

The final octet of the header of a Flooding Scoped LSP as defined in [RFC7356] contains Reserved/LSPDBOL/IS Type information. This field is redefined for the new flooding scopes defined in this document as follows:

Reserved/ATT/LSPDBOL

Bits 8-5 Reserved

Transmitted as 0 and ignored on receipt

Bit 4 ATT

If set to 1 indicates that the sending IS is attached to
routers in other Level N areas via Level N+1

Bit 3 LSDBOL

As defined in RFC7356

Bits 2-1

Transmitted as 0 and ignored on receipt.

Note that the levels supported (analogous to the IS-type information in L1 and L2 LSPs) can be obtained from the Area Hierarchy TLV advertised in the associated LSP #0.

Note that the definition of the ATT bit specified above also applies to L2 LSPs. Previously this bit would have no meaning as [ISO10589] does not define support for Level 3.

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. The authors would also like to thank Tony Pryzienda for his careful review and excellent suggestions.

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 Hierarchy: 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] ISO, "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, November 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>>.
- [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>>.

Appendix A. Preventing Cross Branching in the Hierarchy

The use of additional levels requires careful interconnection of routers which support multiple levels. Consistent association of LSAs is required not only for validating the connections between routers in a level specific area but also for all levels above a given level to which any of the routers may be connected (directly or indirectly). Failure to do so can result in interconnecting different branches of a tree leading to interarea loops. This leads to the requirement that all routers advertise an LSA for all levels regardless of whether a given router is configured to participate in a given level or not.

At first glance it may seem that it would be sufficient for each router to advertise LSAs only for the levels that the router is configured to support. However, the following simple example illustrates why this is problematic.



Since Router B does not support Level 4, it chose not to advertise any Area for Level 4. This means that neither Router A nor Router C can tell by inspecting hellos that not all routers in Level 3 area 30 have been configured to support the same Level 4 area. It is possible for Rtr A and Rtr C to discover the LSAs advertised by all routers by inspecting the Level 3 LSPs - however this requires that Level 3 adjacencies be formed and maintained even when routing cannot be safely performed via all adjacencies in a given area. It then needs to be decided how routing over existing adjacencies should be limited. A number of possibilities exist:

Treat the area as if it were two partitions. In the example Router A would be in one partition and Router C would be in another partition. But Router B could belong to either partition.

Select a winning Level 4 Area among the set of Level 4 areas advertised in L3 LSPs and only allow leaking of routes to/from that level

But either of these options introduce the possibility that a previously fully connected hierarchy becomes partially disconnected as a result of a single configuration change on a single router and/or the bringup of a new router.

The choice made was then to require all routers supporting the extensions in this document to advertise an LSAI for all levels regardless of what specific levels an individual router is configured to support. This guarantees that any inconsistency between the intended connectivity of a router at all levels - direct and indirect - can be detected during exchange of hellos and therefore adjacency bringup can always be blocked when necessary.

Appendix B. Guidelines for Introducing a new level

It is desirable to be able to introduce support for a new level without disruption. This section discusses ways to do this.

Initial deployment may require only the support of one additional level (Level 3). However, in the future increased network scale may make introduction of an additional level (Level 4) desirable. It is suggested that all routers be configured to advertise a single candidate LSAI for Level 4 - for the purposes of the example let's use LSAI 44. When ready to deploy Level 4, it is then only necessary to enable Level 4 on those routers who will be participating in the additional level.

However, perhaps at the time of deploying Level 3 the administrator has no idea what LSAI will be used for Level 4 in the future. In such a case a "dummy" LSAI should be configured for Level 4 on all routers - let's use "0" in this example. In this case, what needs to be done when ready to enable Level 4 is to go to every router (regardless of whether it will actively participate in the new level) and configure the intended LSAI for Level 4. If LSAI 45 is the intended Level 4 area, then LSAI 45 is configured on each router. Each router is then advertising two LSAIs for Level 4: (0, 45). Once this is completed, go to every router and remove the "dummy" Level 4 LSAI (0) and the network is now ready to have this Level 4 area enabled.

In the event that support for a new level needs to be introduced and no LSAI was ever advertised for that level, the introduction of LSAI for the new level will cause temporary adjacency flaps as the advertisement of the LSAI for the new level is introduced. To avoid this, implementations would need to introduce support for temporary disablement of the LSAI check for the new level until the configuration of the new LSAI is complete on all nodes. Support for this transition mode is outside the scope of this document. The need for a transition mode can be avoided if an LSAI is configured for levels 2-8 from day one.

Authors' Addresses

Tony Li
Juniper Networks
1133 Innovation Way
Sunnyvale, California 95089
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