Network Working Group
Request for Comments: 5311
Obsoletes: 3786

D. McPherson, Ed.
Arbor Networks
L. Ginsberg
S. Previdi
M. Shand
Cisco Systems
February 2009

Simplified Extension of Link State PDU (LSP) Space for IS-IS

Status of This Memo

Copyright Notice

Abstract

This document describes a simplified method for extending the Link State PDU (LSP) space beyond the 256 LSP limit. This method is intended as a preferred replacement for the method defined in RFC 3786.

Table of Contents

1. Overview

[IS-IS] defines the set of LSPs that may be originated by a system at each level. This set is limited to 256 LSPs. [IS-IS] also defines a maximum value for an LSP (originatingLxLSPBufferSize) as 1492 bytes. The carrying capacity of an LSP set, while bounded, has thus far been sufficient for advertisements associated with an area/domain in existing deployment scenarios. However, the definition of additional information to be included in LSPs (e.g., multi-topology support, traffic engineering information, router capabilities, etc.) has the potential to exceed the carrying capacity of an LSP set.

This issue first drew interest when traffic engineering extensions were introduced. This interest resulted in the solution defined in [RFC3786]. However, that solution suffers from restrictions required to maintain interoperability with systems that do not support the extensions.

This document defines extensions that allow a system to exceed the 256 LSP limit and do so in a way that has no interoperability issues with systems that do not support the extension. It is seen as a simpler, and therefore preferred, solution to the problem.

2. Specification of Requirements

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 [BCP14].

3. Definition of Commonly Used Terms

This section provides definitions for terms that are used throughout the text. The terminology is consistent with that used in RFC 3786.

Originating System: A physical IS running the IS-IS protocol. As this document describes a method that allows a single physical IS to originate LSPs on behalf of multiple virtual ISs, the Originating System represents the single physical IS.

Normal system-id: The system-id of an Originating System as defined by [IS-IS].

Additional system-id: A system-id other than the "Normal system-id" that is assigned by the network administrator to an Originating System in order to allow the generation of Extended LSPs. The Additional system-id, like the Normal system-id, must be unique throughout the routing area (Level-1) or domain (Level-2).

Original LSP: An LSP using the Normal system-id in its LSP ID.

Extended LSP: An LSP using an Additional system-id in its LSP ID.

LSP set: All LSPs of a given level having the same system ID and Pseudonode ID. (The LSPID field then only varies in the LSP number octet.) This constitutes the complete set of link state information at a given level originated using that system ID/Pseudonode ID. This term is defined to resolve the ambiguity between a logical LSP and a single Link State PDU -- which is sometimes called an LSP fragment. The latter is the unit of information handled by the update process.

Extended LSP set: An LSP set consisting of LSPs using an Additional system-id.

Extension-capable IS: An IS implementing the mechanisms described in this document.

Virtual IS: The system, identified by an Additional system-id, advertised as originating the Extended LSPs. These LSPs specify the Additional system-id in their LSP IDs.

4. Utilizing Additional System IDs

This extension allows an Originating System to be assigned additional system-ids that may be used to generate additional LSP sets. The additional system-ids are subject to the same restrictions as normal system-ids, i.e., when used at Level-1, the additional system-id MUST be unique within the Level-1 area. When used at Level-2, the additional system-id MUST be unique within the domain.

Extended LSPs are treated by the IS-IS Update Process in the same manner as normal LSPs, i.e., the same rules as to generation, flooding, purging, etc. apply. In particular, if the Extended LSP with LSP number zero and remaining lifetime > 0 is not present for a particular additional system-id, then none of the Extended LSPs in that Extended LSP set shall be processed.

4.1. Additional Information in Extended LSPs

The LSP number zero of an Extended LSP set MUST include the new IS alias ID TLV defined in Section 4.4. This allows the Extended LSP set to be associated with the Originating System that generated the LSP(s).

4.2. Extended LSP Restrictions

The following restrictions on the information that may appear in an Extended LSP are defined in order to avoid interoperability issues with systems that do not support the extensions defined in this document. All TLV references are based on the current definitions in the IANA IS-IS TLV Codepoints Registry.

4.2.1. TLVs That MUST NOT Appear

The following TLVs MUST NOT appear in an Extended LSP:

TLV Name (#)
ES Neighbors (3)
Part. DIS (4)
Prefix Neighbors (5)

If any of the TLVs listed above appear in an Extended LSP, an Extension Capable IS MUST ignore those TLVs on receipt and SHOULD report an error. Other TLVs in that Extended LSP set MUST be processed normally.

4.2.2. Leaf Advertisements in Extended LSPs

Advertisement of leaf information in Extended LSPs is allowed. Inclusion of such information requires the advertisement of a neighbor between the Originating System and the Virtual IS associated with the Extended LSP set in which the leaf advertisements appear. See Section 4.2.3.

When leaf advertisements for multiple topologies (see [RFC5120]) are included in an Extended LSP set, the multi-topology TLV (229) MUST include all topologies for which a leaf advertisement is included.

The following TLVs fall into this category:

TLV Name (#)
IP Int. Reach (128)
IP Ext. Address (130)
The extended IP reachability TLV (135)
MT IP Reach (235)
IPv6 IP Reach (236)
MT IPv6 IP Reach (237)

4.2.3. IS Neighbor Advertisement Restrictions

Advertisement of IS Neighbor Reachability in an Extended LSP is restricted to advertisement of neighbor reachability to the Originating System. A neighbor to the Originating System MUST be advertised in Extended LSPs. If multi-topology capability [RFC5120] is supported, an MT IS Neighbor advertisement to the Originating System IS MUST be included for every topology advertised in the Extended LSP set. Neighbor advertisement(s) to the Originating System in an Extended LSP MUST use a non-zero metric and SHOULD use a metric of MaxLinkMetric-1.

The restrictions defined here apply to all TLVs used to advertise neighbor reachability. These include the following TLVs:

TLV Name (#)
IIS Neighbors (2)
The extended IS reachability TLV (22)
MT-ISN (222)

4.2.4. Area Addresses

LSP number zero of an Extended LSP set MUST include an Area Address TLV. The set of area addresses advertised MUST be a subset of the set of Area Addresses advertised in the normal LSP number zero at the corresponding level. Preferably, the advertisement SHOULD be syntactically identical to that included in the normal LSP number zero at the corresponding level.

4.2.5. Overload, Attached, Partition Repair Bits

The Overload (OL), Attached (ATT), and Partition Repair (P) bits MUST be set to 0 in all Extended LSPs.

Note that ISs NOT supporting these extensions will interpret these bits normally in Extended LSPs they receive. If the ATT bit were set in an Extended LSP, this could indicate that the Virtual IS is attached to other areas when the Originating System is not. This might cause legacy systems to use the Virtual IS as a default exit point from the area.

4.3. Originating LSP Requirements

The Original LSP set MUST include a neighbor to the Virtual IS associated with each Extended LSP set generated. If multi-topology capability [RFC5120] is supported, an MT IS Neighbor advertisement to the Virtual IS MUST be included for every topology advertised in the Extended LSP set. The neighbor advertisement(s) in the Original LSP MUST specify a metric of zero. This guarantees that the two-way connectivity check between Originating System and Virtual IS will succeed and that the cost of reaching the Virtual IS is the same as the cost to reach the Originating System.

4.4. IS Alias ID TLV (IS Alias ID)

The IS-Alias TLV allows extension-capable ISs to recognize the Originating System of an Extended LSP set. It identifies the Normal system-id of the Originating System.

Type 24
Length # of octets in the value field (7 to 255)
Value

	No. of octets
Normal System-id	6
Sub-TLV length	1
Sub-TLVs (optional)	0 to 248

Normal system-id

The Normal system-id of the Originating System.

Sub-TLVs length

Total length of all sub-TLVs.

Sub-TLVs

No sub-TLVs are defined in this document. Should future extensions define sub-TLVs, the sub-TLVs MUST be formatted as described in [RFC5305].

4.5. New TLVs in Support of IS Neighbor Attributes

One of the major sources of additional information in LSPs is the sub-TLV information associated with the extended IS reachability TLV (22) and MT-ISN TLV (222). This includes (but is not limited to) information required in support of Traffic Engineering (TE) as defined in [RFC5305] and [RFC5307]. The restrictions defined in this document prohibit the presence of TLV 22 and/or TLV 222 in Extended LSPs except to advertise the neighbor relationship to the Originating System. In the event that there is a need to advertise in Extended LSPs such information associated with neighbors of the Originating System, it is necessary to define new TLVs to carry the sub-TLV information.

Two new TLVs are therefore defined.

1) IS Neighbor Attribute TLV (23). It is identical in format to the extended IS reachability TLV (22).

2) MT IS Neighbor Attribute TLV (223). It is identical in format to the MT-ISN TLV (222).

These new TLVs MAY be included in Original LSPs or Extended LSPs. Regardless of the type of LSP in which the TLVs appear, the information pertains to the neighbor relationship between the Originating System and the IS identified in the TLV.

These TLVs MUST NOT be used to infer that a neighbor relationship exists in the absence of TLV 22 or TLV 222 (whichever applies) in the Originating LSP set for the specified neighbor. This restriction is necessary in order to maintain compatibility with systems that do not support these extensions.

5. Comparison with the RFC 3786 Solution

This document utilizes the same basic mechanism (additional system-ids) as RFC 3786 to allow an originating system to generate more than 256 LSPs. It differs from RFC 3786 in that it restricts the content of Extended LSPs to information that does NOT impact the building of a Shortest Path Tree (SPT).

Legacy IS-IS implementations which do not support the extensions defined in this document see the Extended LSPs as information associated with a system that is reachable only via the Originating System. As no other systems are reachable via the Virtual ISs, the Shortest Path First (SPF) calculation in legacy ISs is therefore consistent with that performed by extension-capable ISs. There is therefore no need for the two different operating modes defined in RFC 3786.

There is also no need for the special handling of the original LSP set and the Extended LSP set(s) as a single Logical LSP during the SPF as specified in Section 5 of RFC 3786.

6. Deployment Considerations

There are a number of deployment considerations that limit the usefulness of Extended LSPs unless all systems are extension-capable ISs.

6.1. Advertising New TLVs in Extended LSPs

As Extended LSPs MAY be utilized to advertise TLVs associated with other protocol extensions (definition of which is outside the scope of this document) and/or the extensions defined in Section 4.4 of this document, it is obvious that the utilization of the information in Extended LSPs by legacy IS-IS implementations will be limited. The implication of this is that as implementations are revised to support the protocol extensions that define new TLVs/sub-TLVs that MAY be advertised in Extended LSPs; the implementation SHOULD also be revised to support the extensions defined in this document so that it is capable of processing the new information whether it appears in normal or Extended LSPs.

6.2. Reachability and Non-SPF TLV Staleness

In cases where non-SPF information is advertised in LSPs, it is necessary to determine whether the system that originated the advertisement is reachable in order to guarantee that a receiving IS does not use or leak stale information. As long as the OL bit is NOT set by the Originating System in normal LSPs, reachability to the Virtual IS will be consistent with reachability to the Originating System. Therefore, no special rules are required in this case.

6.3. Normal LSP OL State and Use of Extended LSPs

If the Originating System sets the OL bit in a normal LSP, legacy systems will see the Virtual ISs associated with that Originating System as unreachable and therefore will not use the information in the corresponding Extended LSPs. Under these circumstances, Extension-capable ISs MUST also see the Virtual ISs as unreachable. This avoids potential routing loops in cases where leaf information is advertised in Extended LSPs.

6.4. Moving Neighbor Attribute INFO LSPs

Section 4.4 defines new TLVs that MAY be used to advertise neighbor attribute information in Extended LSPs. In cases where neighbor attribute information associated with the same context (e.g., the same link) appears in both an Original LSP and in one or more Extended LSP sets, the following rules apply for each attribute:

- If the attribute information does not conflict, it MUST be considered additive.
- If the attribute information conflicts, then the information in the Original LSP, if present, MUST be used. If no information is

in the Original LSP, then the information from the Extended LSP with the lowest system-id SHALL be preferred.

- In cases where information about the same neighbor/link/attribute appears in both TLV 22 and TLV 23 (or TLV 222 and TLV 223 for the same MTID) then the information in TLV 22 (or TLV 222) MUST be used and the information in TLV 23 (or TLV 223) MUST be ignored.

Utilization of the new TLVs for neighbor attribute information would provide additional benefits that include:

- Elimination of the need for redundant IS neighbor TLVs to be processed as part of the SPF.

- Easier support for a set of TE information associated with a single link that exceeds the 255-byte TLV limit by allowing the interpretation of multiple TLVs to be considered additive rather than mutually exclusive.

6.5. Advertising Leaf INFO Extended LSPs

The need to advertise leaf information in Extended LSPs may arise because of extensive leaking of inter-level information or because of the support of multiple topologies as described in [RFC5120]. When leaf information is advertised in Extended LSPs, these LSPs now contain information that MUST be processed in order to correctly update the forwarding plane of an IS. This may increase the frequency of events that trigger forwarding plane updates by ISs in the network. It is therefore recommended that, when possible, leaf information be restricted to the normal LSP set.

7. Security Considerations

This document raises no new security issues for IS-IS. For general security considerations for IS-IS, see [RFC5304].

8. IANA Considerations

This document defines the following new ISIS TLVs that are reflected in the ISIS TLV codepoint registry:

Type	Description	IIH	LSP	SNP
23	IS Neighbor Attribute	n	y	n
24	IS Alias ID	n	y	n
223	MT IS Neighbor Attribute	n	y	n

9. References

9.1. Normative References

[IS-IS] ISO, "Intermediate system to Intermediate system routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)," ISO/IEC 10589:2002, Second Edition.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5307] Kompella, K., Ed., and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, October 2008.

[BCP14] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

9.2. Informative References

[RFC3786] Hermelin, A., Previdi, S., and M. Shand, "Extending the Number of Intermediate System to Intermediate System (IS-IS) Link State PDU (LSP) Fragments Beyond the 256 Limit", RFC 3786, May 2004.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, October 2008.

Authors' Addresses

Danny McPherson (editor)
Arbor Networks, Inc.
EMail: danny@arbor.net

Les Ginsberg
Cisco Systems
EMail: ginsberg@cisco.com

Stefano Previdi
Cisco Systems
EMail: sprevidi@cisco.com

Mike Shand
Cisco Systems
EMail: mshand@cisco.com