

Networking Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: August 19, 2016

L. Ginsberg  
A. Bashandy  
C. Filsfils  
S. Previdi  
Cisco Systems  
M. Nanduri  
Microsoft  
E. Aries  
Private Contributor  
February 16, 2016

Advertising L2 Bundle Member Link Attributes in IS-IS  
draft-ginsberg-isis-l2bundles-02.txt

Abstract

There are deployments where the Layer 3 interface on which IS-IS operates is a Layer 2 interface bundle. Existing IS-IS advertisements only support advertising link attributes of the Layer 3 interface. If entities external to IS-IS wish to control traffic flows on the individual physical links which comprise the Layer 2 interface bundle link attribute information about the bundle members is required.

This document introduces the ability for IS-IS to advertise the link attributes of layer 2 (L2) bundle members.

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].

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 <http://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 August 19, 2016.

#### Copyright Notice

Copyright (c) 2016 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 (<http://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  |
| 2. L2 Bundle Member Attributes TLV . . . . .                   | 3  |
| 2.1. Parallel L3 Adjacencies . . . . .                         | 5  |
| 2.2. Shared Attribute sub-TLVs . . . . .                       | 5  |
| 3. Advertising L2 Bundle Member Adj-SIDs . . . . .             | 5  |
| 3.1. L2 Bundle Member Adjacency Segment Identifier sub-TLV . . | 6  |
| 3.2. L2 Bundle Member LAN Adjacency Segment Identifier sub-TLV | 7  |
| 4. IANA Considerations . . . . .                               | 9  |
| 5. Security Considerations . . . . .                           | 11 |
| 6. Acknowledgements . . . . .                                  | 11 |
| 7. References . . . . .  | 11 |
| 7.1. Normative References . . . . .                            | 11 |
| 7.2. Informational References . . . . .                        | 12 |
| Authors' Addresses . . . . .                                   | 12 |

#### 1. Introduction

There are deployments where the Layer 3 interface on which an IS-IS adjacency is established is a Layer 2 interface bundle, for instance a Link Aggregation Group (LAG) [IEEE802.1AX]. This reduces the number of adjacencies which need to be maintained by the routing protocol in cases where there are parallel links between the neighbors. Entities external to IS-IS such as Path Computation Elements (PCE) [RFC4655] may wish to control traffic flows on individual members of the underlying Layer 2 bundle. In order to do so link attribute information about individual bundle members is required - but currently IS-IS only supports advertising link attributes for the Layer 3 interfaces on which it operates.

This document introduces a new TLV to advertise link attribute information for each of the L2 bundle members which comprise the Layer 3 interface on which IS-IS operates.

[SR] introduces a new link attribute - adjacency segment identifier (Adj-SID) - which can be used as an instruction to forwarding to send traffic over a specific link. This document introduces additional sub-TLVs to advertise Adj-SIDs for L2 Bundle members.

Note that the new advertisements defined in this document are intended to be provided to external entities.

## 2. L2 Bundle Member Attributes TLV

A new TLV is introduced to advertise L2 Bundle member attributes. Although much of the information is identical to and uses the same sub-TLVs included in Extended IS-Neighbor advertisements (TLVs 22 and 222), a new TLV is used so that changes to the advertisement of the L2 Bundle member link attributes does not trigger unnecessary action by the [ISO10589] Decision process.

This new TLV utilizes the sub-TLV space defined for TLVs 22, 23, 141, 222, and 223.

The following new TLV is introduced:

## L2 Bundle Member Attributes

Type: 25 (suggested - to be assigned by IANA)

Length: Number of octets to follow

## Parent L3 Neighbor Descriptor

L3 Neighbor System ID + pseudonode ID (7 octets)

Flags: 1 octet field of following flags:

```

  0 1 2 3 4 5 6 7
  +-----+
  |P|                                     |
  +-----+

```

where:

P-flag: When set to 1 one of the sub-TLVs described in Section 2.1 immediately follows the flags field. If the P-flag is set to 0, then none of the sub-TLVs described in Section 2.1 are present.

Other bits: MUST be zero when originated and ignored when received.

One or more of the following:

## L2 Bundle Attribute Descriptors

Length of L2 Bundle Attribute Descriptor (1 octet)

NOTE: This includes all fields described below.

Number of L2 Bundle Member Descriptors (1 octet)

L2 Bundle Member Link Local Identifiers

(4 \* Number of L2 Bundle Member Descriptors octets)

NOTE: An L2 Bundle Member Descriptor is a Link Local Identifier as defined in [RFC5307].

## sub-TLV(s)

A sub-TLV may define an attribute common to all of the bundle members listed or a sub-TLV may define an attribute unique to each bundle member. Use of these two classes of sub-TLVs is described in the following sections.

NOTE: Only one Parent L3 Neighbor Descriptor is present in a given TLV. Multiple L2 Bundle Attribute Descriptors may be present in a single TLV.

### 2.1. Parallel L3 Adjacencies

When there exist multiple L3 adjacencies to the same neighbor additional information is required to uniquely identify the L3 Neighbor. One and only one of the following three sub-TLVs is used to uniquely identify the L3 adjacency:

- o IPv4 Interface Address (sub-TLV 6 defined in [RFC5305])
- o IPv6 Interface Address (sub-TLV 12 defined in [RFC6119])
- o Link Local/Remote Identifiers (sub-TLV 4 defined in [RFC5307])

When the P-bit is set in the flags field in the Parent L3 Neighbor Descriptor one and only one of the above sub-TLVs MUST be present. The chosen sub-TLV MUST immediately follow the flags field described in Section 2.

These sub-TLVs MAY be omitted if no parallel adjacencies to the neighbor exist.

### 2.2. Shared Attribute sub-TLVs

These sub-TLVs advertise a single copy of an attribute (e.g. link bandwidth). The attribute applies to all of the L2 Bundle Members in the set advertised under the preceding L2 Bundle Member Attribute Descriptor. No more than one copy of a given sub-TLV in this category may appear in the set of sub-TLVs under the preceding L2 Bundle Member Attribute Descriptor. If multiple copies of a given sub-TLV are present both MUST be ignored.

The set of L2 Bundle Member Descriptors which may be advertised under a single L2 Bundle Member Attribute Descriptor is therefore limited to bundle members which share the set of attributes advertised in the shared attribute sub-TLVs.

All existing sub-TLVs defined in the IANA Sub-TLVs for TLVs 22, 23, 141, 222, and 223 registry are in the category of shared attribute sub-TLVs unless otherwise specified in this document.

### 3. Advertising L2 Bundle Member Adj-SIDs

[SR] defines sub-TLVs to advertise Adj-SIDs for L3 adjacencies. However these sub-TLVs only support a advertisement of a single Adj-SID. As it is expected that each L2 Bundle member will have unique Adj-SIDs in many deployments it is desirable to define a new sub-TLV which allows more efficient encoding of a set of Adj-SIDs in a single sub-TLV. Two new sub-TLVs are therefore introduced to support

advertising Adj-SIDs for L2 Bundle members. The format of the new sub-TLVs is similar to that used for L3 adjacencies, but is optimized to allow advertisement of a set of Adj-SIDs (one per L2 Bundle Member) in a single sub-TLV.

The two new sub-TLVs defined in the following sections do not fall into the category of shared attribute sub-TLVs.

### 3.1. L2 Bundle Member Adjacency Segment Identifier sub-TLV

This sub-TLV is used to advertise Adj-SIDs for L2 Bundle Members associated with a parent L3 adjacency which is Point-to-Point. The following format is defined for this sub-TLV:

Type: 41 (suggested value to be assigned by IANA) (1 octet)  
Length: variable (1 octet)

Flags: 1 octet field of following flags:

```

  0 1 2 3 4 5 6 7
  +-----+
  |F|*|V|L|S|   |
  +-----+

```

where:

\* - Is a flag used in the L3 Adj-SID sub-TLV but which is NOT used in this sub-TLV. These bits SHOULD be sent as 0 and MUST be ignored on receipt

F-Flag: Address-Family flag. If unset, then the Adj-SID refers to an L2 Bundle Member with outgoing IPv4 encapsulation. If set then the Adj-SID refers to an L2 Bundle Member with outgoing IPv6 encapsulation.

V-Flag: Value flag. If set, then the Adj-SID carries a value. By default the flag is SET.

L-Flag: Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.

S-Flag. Set Flag. When set, the S-Flag indicates that the Adj-SID refers to a set of L2 Bundle Members (and therefore MAY be assigned to other L2 Bundle Members as well).

Other bits: MUST be zero when originated and ignored when

received.

Weight: 1 octet. The value represents the weight of the Adj-SID for the purpose of load balancing. The use of the weight is defined in [SR-ARCH].

NOTE: Flags and weight are shared by all L2 Bundle Members listed in the L2 Bundle Attribute Descriptor.

L2 Bundle Member Adj-SID Descriptors. There MUST be one descriptor for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor. Each descriptor consists of one of the following fields:

SID/Index/Label: according to the V and L flags, it contains either:

- \* A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- \* A 4 octet index defining the offset in the SID/Label space advertised by this router. See [SR]. In this case V and L flags MUST be unset.
- \* A 16 octet IPv6 address. In this case the V flag MUST be set. The L flag MUST be unset if the IPv6 address is globally unique.

### 3.2. L2 Bundle Member LAN Adjacency Segment Identifier sub-TLV

This sub-TLV is used to advertise Adj-SIDs for L2 Bundle Members associated with a parent L3 adjacency which is a LAN adjacency. In LAN subnetworks, the Designated Intermediate System (DIS) is elected and originates the Pseudonode-LSP (PN-LSP) including all neighbors of the DIS. When Segment Routing is used, each router in the LAN MAY advertise the Adj-SID of each of its neighbors on the LAN. Similarly, for each L2 Bundle Member a router MAY advertise an Adj-SID to each neighbor on the LAN.

The following format is defined for this sub-TLV:

Type: 42 (suggested value to be assigned by IANA) (1 octet)  
Length: variable (1 octet)  
Neighbor System ID: 6 octets

Flags: 1 octet field of following flags:

```

  0 1 2 3 4 5 6 7
+-----+-----+
|F|*|V|L|S|   |
+-----+-----+

```

where:

\* - Is a flag used in the L3 Adj-SID sub-TLV but which is NOT used in this sub-TLV. These bits SHOULD be sent as 0 and MUST be ignored on receipt

F-Flag: Address-Family flag. If unset, then the Adj-SID refers to an L2 Bundle Member with outgoing IPv4 encapsulation. If set then the Adj-SID refers to an L2 Bundle Member with outgoing IPv6 encapsulation.

V-Flag: Value flag. If set, then the Adj-SID carries a value. By default the flag is SET.

L-Flag: Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.

S-Flag. Set Flag. When set, the S-Flag indicates that the Adj-SID refers to a set of L2 Bundle Members (and therefore MAY be assigned to other L2 Bundle Members as well).

Other bits: MUST be zero when originated and ignored when received.

Weight: 1 octet. The value represents the weight of the Adj-SID for the purpose of load balancing. The use of the weight is defined in [SR-ARCH].

NOTE: Flags and weight are shared by all L2 Bundle Members listed in the L2 Bundle Attribute Descriptor.

L2 Bundle Member LAN Adj-SID Descriptors. There MUST be one descriptor for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor. Each descriptor consists of one of the following fields:

SID/Index/Label: according to the V and L flags, it contains either:

\* A 3 octet local label where the 20 rightmost bits are used

for encoding the label value. In this case the V and L flags MUST be set.

- \* A 4 octet index defining the offset in the SID/Label space advertised by this router. See [SR]. In this case V and L flags MUST be unset.
- \* A 16 octet IPv6 address. In this case the V flag MUST be set. The L flag MUST be unset if the IPv6 address is globally unique.

#### 4. IANA Considerations

This document adds the following new TLV to the IS-IS TLV Codepoints registry.

Value: 25 (suggested - to be assigned by IANA)

Name: L2 Bundle Member Attributes

The name of the Sub-TLVs for TLVs 22, 23, 141, 222, and 223 registry needs to be changed to Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 registry. An additional column needs to be added to the registry to indicate which sub-TLVs may appear in the new L2 Bundle Member Attributes TLV. The following table indicates the appropriate settings for all currently defined sub-TLVs as regards their use in the new L2 Bundle Member Attributes TLV.

3 Administrative group (color) y  
 4 Link Local/Remote Identifiers y  
 6 IPv4 interface address y  
 8 IPv4 neighbor address y  
 9 Maximum link bandwidth y  
 10 Maximum reservable link bandwidth y  
 11 Unreserved bandwidth y  
 12 IPv6 Interface Address y  
 13 IPv6 Neighbor Address y  
 14 Extended Administrative Group y  
 18 TE Default metric y  
 19 Link-attributes y  
 20 Link Protection Type y  
 21 Interface Switching Capability Descriptor y  
 22 Bandwidth Constraints y  
 23 Unconstrained TE LSP Count y  
 24 Remote AS number n  
 25 IPv4 remote ASBR Identifier n  
 26 IPv6 remote ASBR Identifier n  
 27 Interface Adjustment Capability Descriptor (IACD) y  
 28 MTU n  
 29 SPB-Metric y  
 30 SPB-A-OALG y

This document adds the following new sub-TLVs to the sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 registry.

Value: 41 (suggested - to be assigned by IANA)

Name: L2 Bundle Member Adj-SID

This sub-TLV is allowed in the following TLVs:

|    |    |    |     |     |     |
|----|----|----|-----|-----|-----|
| 22 | 23 | 25 | 141 | 222 | 223 |
| n  | n  | y  | n   | n   | n   |

Value: 42 (suggested to be assigned by IANA)

Name: L2 Bundle Member LAN Adj-SID

This sub-TLV is allowed in the following TLVs:

|    |    |    |     |     |     |
|----|----|----|-----|-----|-----|
| 22 | 23 | 25 | 141 | 222 | 223 |
| n  | n  | y  | n   | n   | n   |

## 5. Security Considerations

None.

## 6. Acknowledgements

The authors would like to thank Jon MITchell for his careful review.

## 7. References

### 7.1. Normative References

[IEEE802.1AX]

Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation.", ISO/IEC 10589:2002, Second Edition, Nov 2008.

[ISO10589]

International Organization for Standardization, "Intermediate system to Intermediate system intra-domain 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, 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, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.

[RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <<http://www.rfc-editor.org/info/rfc5307>>.

[RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", RFC 6119, DOI 10.17487/RFC6119, February 2011, <<http://www.rfc-editor.org/info/rfc6119>>.

## 7.2. Informational References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [SR] "IS-IS Extensions for Segment Routing, draft-ietf-isis-segment-routing-extensions-06(work in progress)", December 2015.
- [SR-ARCH] "Segment Routing Architecture, draft-ietf-spring-segment-routing-07(work in progress)", December 2015.

## Authors' Addresses

Les Ginsberg  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas, CA 95035  
USA

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

Ahmed Bashandy  
Cisco Systems  
170 West Tasman Drive  
San Jose, Ca 95134  
US

Clarence Filsfils  
Cisco Systems

Email: [cf@cisco.com](mailto:cf@cisco.com)

Stefano Previdi  
Cisco Systems  
Via Del Serafico 200  
Rome 0144  
Italy

Email: [sprevidi@cisco.com](mailto:sprevidi@cisco.com)

Mohan Nanduri  
Microsoft

Email: [mnanduri@microsoft.com](mailto:mnanduri@microsoft.com)

Ebben Aries  
Private Contributor

Email: [exa@dscp.org](mailto:exa@dscp.org)

Networking Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: February 1, 2016

L. Ginsberg  
P. Wells  
S. Previdi  
Cisco Systems  
B. Decraene  
Orange  
T. Przygienda  
Ericsson  
H. Gredler  
Juniper Networks, Inc  
July 31, 2015

IS-IS Minimum Remaining Lifetime  
draft-ginsberg-isis-remaining-lifetime-00.txt

Abstract

Corruption of the Remaining Lifetime Field in a Link State PDU can go undetected. In certain scenarios this may cause or exacerbate flooding storms. It is also a possible denial of service attack vector. This document defines a backwards compatible solution to this problem.

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].

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 <http://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 February 1, 2016.

## Copyright Notice

Copyright (c) 2015 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 (<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Table of Contents

|   |   |
|---|---|
| 1. Problem Statement . . . . .                | 2 |
| 2. Solution . . . . .                         | 4 |
| 3. Deployment Considerations . . . . .        | 5 |
| 3.1. Inconsistent Values for MaxAge . . . . . | 5 |
| 3.2. Reporting Corrupted Lifetime . . . . .   | 6 |
| 3.3. Impact of Delayed LSP Purging . . . . .  | 6 |
| 4. IANA Considerations . . . . .              | 7 |
| 5. Security Considerations . . . . .          | 7 |
| 6. Acknowledgements . . . . .                 | 7 |
| 7. References . . . . .                       | 7 |
| 7.1. Normative References . . . . .           | 7 |
| 7.2. Informational References . . . . .       | 8 |
| Authors' Addresses . . . . .                  | 8 |

## 1. Problem Statement

Each Link State PDU (LSP) includes a Remaining Lifetime field. This field is set by the originator based on local configuration and then decremented by all systems once the entry is stored in their Link

State PDU Database (LSPDB) consistent with the passing of time. This allows all Intermediate Systems (ISs) to age out the LSP at approximately the same time.

Each LSP also has a checksum field to allow receiving systems to detect errors which may have occurred during transmission. As the Remaining Lifetime field changes as it is flooded and as the checksum field MUST NOT be altered by receiving ISs the Remaining Lifetime is deliberately excluded from the checksum calculation. In cases where cryptographic authentication is included in an LSP ([RFC5304] or [RFC5310]) the Remaining Lifetime field is also excluded from the hash calculation. If the Remaining Lifetime field gets corrupted during flooding this corruption is therefore undetectable. The consequences of such corruption depend upon how the Remaining Lifetime is altered.

In cases where the Remaining Lifetime becomes larger than the originator intended the impact is benign. As the originator is responsible for refreshing the LSP before it ages out a new version of the LSP will be generated before the LSP ages out - so no harm is done.

In cases where the Remaining Lifetime field becomes smaller than the originator intended the LSP may age out prematurely (i.e. before the originator refreshes the LSP). This has two negative consequences:

1. The LSP will be purged by an IS when the Remaining Lifetime expires. This will cause a temporary loss of reachability to destinations impacted by the content of that LSP.
2. Unnecessary LSP flooding will occur as a result of the premature purge and subsequent regeneration/flooding of a new version of the LSP by the originator

If the corrupted Remaining Lifetime is only modestly shorter than the lifetime assigned by the originator the negative impacts are also modest. If, however, the corrupted Remaining Lifetime becomes very small, then the negative impacts can become significant - especially in cases where the cause of the corruption is persistent so that the cycle repeats itself frequently.

A backwards compatible solution to this problem is defined in the following sections.

## 2. Solution

As discussed in the previous section, the problematic case is when Remaining Lifetime is corrupted and becomes much smaller than it should be. The goal of the solution is then to prevent premature purging.

Under normal circumstances updates to an LSP - including purging if appropriate - are the responsibility of the originator of the LSP. There is a maximum time between generations of a given LSP. Once this time has expired it is the responsibility of the originator to refresh the LSP (i.e. issue a new version with higher sequence number) even if the contents of the LSP have not changed. [ISO10589] specifies that maximumLSPGenerationInterval MUST be sufficiently less than the maximum lifetime of an LSP so that the new version can be flooded network wide before the old version ages out on any IS.

There are two cases where a system other than the originator of an LSP is allowed to purge an LSP:

1. The LSP ages out. This should only occur if the originating IS is no longer reachable and therefore is unable to update the LSP
2. There is a Designated Intermediate System (DIS) change on a LAN. The pseudo-node LSPs generated by the previous DIS are no longer required and MAY be purged by the new DIS.

In both of these cases purging is not necessary for correct operation of the protocol. It is provided as an optimization to remove stale entries from the LSPDB.

In cases where the Remaining Lifetime in a received LSP has been corrupted and is smaller than the remaining lifetime at the originating node when the RemainingLifetime expires on the receiving node it can appear as if the originating IS has failed to regenerate the LSP (case #1 above) when in fact the LSP still has significant lifetime remaining. To prevent this from having a negative impact a modest change to the storage of "new" LSPs in the LSPDB is specified.

[ISO10589] Section 7.3.16 defines the rules to determine whether a received LSP is older, the same, or newer than the copy of the same LSP in the receiver's LSPDB. The key elements are:

- o Higher sequence numbers are newer
- o If sequence numbers are the same, an LSP with zero RemainingLifetime (a purged LSP) is newer than the same LSP w non-zero RemainingLifetime

- o If both the received and local copy of the LSP have non-zero RemainingLifetime they are considered the same even if the RemainingLifetimes differ

[ISO10589] Section 7.3.15.1.e(1) defines the actions to take on receipt of an LSP generated by another IS which is newer than the local copy and has a non-zero RemainingLifetime. An additional action is added:

- vi. If the RemainingLifetime of the new LSP is less than MaxAge it is set to MaxAge

This additional action insures that no matter what value of Remaining Lifetime is received a system other than the originator of an LSP will never purge the LSP until the LSP has existed in the database for at least MaxAge.

It is important to note that no change is proposed for handling the receipt of purged LSPs. The rules specified in [ISO10589] Section 7.3.15.1b still apply i.e., an LSP received with zero RemainingLifetime is still considered newer than a matching LSP with non-zero RemainingLifetime. Therefore the changes proposed here will not result in LSPDB inconsistency among routers in the network.

### 3. Deployment Considerations

This section discusses some possible deployment issues for this protocol extension.

#### 3.1. Inconsistent Values for MaxAge

[ISO10589] defines MaxAge (the maximum value for Remaining Lifetime in an LSP) as an architectural constant of 20 minutes (1200 seconds). However, in practice, implementations have supported allowing this value to be configurable. The common intent of a configurable value is to support longer lifetimes than the default - thus reducing the periodic regeneration of LSPs in the absence of topology changes. See a discussion of this point in [RFC3719]. It is therefore possible for the value of MaxAge on the IS which originates an LSP to be higher or lower than the value of MaxAge on the ISs which receive the LSP.

If the value of MaxAge of the IS which originated the LSP is smaller than the value of MaxAge of the receiver of an LSP, then setting the RemainingLifetime of the received LSP to the local value of MaxAge will insure that it is not purged prematurely. However, if the value of MaxAge on the receiver is less than that of the originator then it is still possible when using the extension defined in the previous

section to have an LSP purged prematurely. Implementors of this extension MAY wish to protect against this case by making the value to which RemainingLifetime is set under the conditions described in the previous section configurable. If that is done the configured value MUST be greater than or equal to the locally configured value of MaxAge.

### 3.2. Reporting Corrupted Lifetime

It may be useful for an IS to report reception of an LSP with a possible corrupt RemainingLifetime field. In order to minimize the reports of false positives the following algorithm SHOULD be used in determining whether the RemainingLifetime in the received LSP is possibly corrupt:

- o The LSP has passed all acceptance tests as specified in [ISO10589] Section 7.3.15.1
- o The LSP is newer than the copy in the local LSPDB (including the case of not being present in the LSPDB)
- o RemainingLifetime in the received LSP is less than ZeroAgeLifetime
- o The adjacency to the neighbor from which the LSP is received has been up for a minimum of ZeroAgeLifetime

In such a case an IS MAY generate a CorruptRemainingLifetime event.

Note that it is not possible to guarantee that all cases of corrupt RemainingLifetime will be detected using the above algorithm. It is also not possible to guarantee that all CorruptRemainingLifetime events reported using the above algorithm are valid. As a diagnostic aid an implementation MAY wish to retain the value of RemainingLifetime received when the LSP was added to the LSPDB.

### 3.3. Impact of Delayed LSP Purging

The extensions defined in this document may result in retaining an LSP longer than its original lifetime. In order for this to occur the scheduled refresh of the LSP by the originator of the LSP must fail to occur - which implies the originator is no longer reachable. In such a case its neighbors will update their own LSPs reporting the loss of connectivity to the originator. LSPs from a node which is unreachable (failure of the two-way-connectivity check) MUST NOT be used. Note this behavior applies to ALL information in the set of LSPs from such a node.

Retention of stale LSPs therefore has no negative side effects other than requiring additional memory for the LSPDB. In networks where a combination of pathological behaviors (LSP corruption, frequent resetting of nodes in the network) is seen this could lead to a large number of stale LSPs being retained - but such networks are already compromised.

#### 4. IANA Considerations

None.

#### 5. Security Considerations

The ability to introduce corrupt LSPs is not altered by the rules defined in this document. Use of authentication as defined in [RFC5304] and [RFC5310] prevents such LSPs from being intentionally introduced. A "man-in-the-middle" attack which modifies an existing LSP by changing the Remaining Lifetime to a small value can cause premature purges even in the presence of cryptographic authentication. The mechanisms defined in this document prevent such an attack from being effective.

#### 6. Acknowledgements

The problem statement in [LIFE-PROB] motivated this work.

#### 7. References

##### 7.1. Normative References

[ISO10589]

International Organization for Standardization, "Intermediate system to Intermediate system intra-domain 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, 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, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC5304]

Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<http://www.rfc-editor.org/info/rfc5310>>.

## 7.2. Informational References

[LIFE-PROB] "IS-IS LSP lifetime corruption - Problem Statement, draft-decraene-isis-lsp-lifetime-problem-statement-00(work in progress)", July 2015.

[RFC3719] Parker, J., Ed., "Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3719, DOI 10.17487/RFC3719, February 2004, <<http://www.rfc-editor.org/info/rfc3719>>.

## Authors' Addresses

Les Ginsberg  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas, CA 95035  
USA

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

Paul Wells  
Cisco Systems  
170 W Tasman Dr  
San Jose, Ca 95035  
USA

Email: [pauwells@cisco.com](mailto:pauwells@cisco.com)

Stefano Previdi  
Cisco Systems  
Via Del Serafico 200  
Rome 0144  
Italy

Email: [sprevidi@cisco.com](mailto:sprevidi@cisco.com)

Bruno Decraene  
Orange  
38 rue du General Leclerc  
Issy Moulineaux cedex 9 92794  
France

Email: [bruno.decraene@orange.com](mailto:bruno.decraene@orange.com)

Tony Przygienda  
Ericsson  
300 Holger Way  
San Jose, Ca 95134  
USA

Email: [antoni.przygienda@ericsson.com](mailto:antoni.przygienda@ericsson.com)

Hannes Gredler  
Juniper Networks, Inc  
1194 N. Matilda Ave  
Sunnyvale, Ca 94089  
USA

Email: [hannes@juniper.net](mailto:hannes@juniper.net)

Networking Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 9, 2016

L. Ginsberg  
S. Previdi  
Cisco Systems  
M. Chen  
Huawei Technologies Co., Ltd  
October 7, 2015

IS-IS Extensions for Advertising Router Info  
draft-ginsberg-isis-rfc4971bis-00.txt

Abstract

This document defines a new optional Intermediate System to Intermediate System (IS-IS) TLV named CAPABILITY, formed of multiple sub-TLVs, which allows a router to announce its capabilities within an IS-IS level or the entire routing domain.

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].

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 <http://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 April 9, 2016.

Copyright Notice

Copyright (c) 2015 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

(<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

#### Table of Contents

|  |   |
|--|---|
| 1. Introduction . . . . .  | 2 |
| 2. IS-IS Router CAPABILITY TLV . . . . .                                     | 3 |
| 3. Elements of Procedure . . . . .   | 4 |
| 4. Interoperability with Routers Not Supporting the Capability TLV . . . . . | 6 |
| 5. Security Considerations . . . . .   | 6 |
| 6. IANA Considerations . . . . .   | 7 |
| 7. Acknowledgements . . . . .  | 7 |
| 8. References . . . . .  | 7 |
| 8.1. Normative References . . . . .  | 7 |
| 8.2. Informational References . . . . .                                      | 8 |
| Authors' Addresses . . . . .   | 8 |

#### 1. Introduction

There are several situations where it is useful for the IS-IS [ISO10589] [RFC1195] routers to learn the capabilities of the other routers of their IS-IS level, area, or routing domain. For the sake of illustration, three examples related to MPLS Traffic Engineering (TE) are described here:

1. Mesh-group: the setting up of a mesh of TE Label Switched Paths (LSPs) [RFC5305] requires some significant configuration effort. [RFC4972] proposes an auto-discovery mechanism whereby every Label Switching Router (LSR) of a mesh advertises its mesh-group membership by means of IS-IS extensions.

2. Point to Multipoint TE LSP (RFC4875). A specific sub-TLV [RFC5073] allows an LSR to advertise its Point To Multipoint capabilities ([RFC4875] and [RFC4461]).
3. Inter-area traffic engineering: Advertisement of the IPv4 and/or the IPv6 Traffic Engineering Router IDs.

The use of IS-IS for Path Computation Element (PCE) discovery may also be considered and will be discussed in the PCE WG.

The capabilities mentioned above require the specification of new sub-TLVs carried within the CAPABILITY TLV defined in this document.

Note that the examples above are provided for the sake of illustration. This document proposes a generic capability advertising mechanism that is not limited to MPLS Traffic Engineering.

This document defines a new optional IS-IS TLV named CAPABILITY, formed of multiple sub-TLVs, which allows a router to announce its capabilities within an IS-IS level or the entire routing domain. The applications mentioned above require the specification of new sub-TLVs carried within the CAPABILITY TLV defined in this document.

Definition of these sub-TLVs is outside the scope of this document.

## 2. IS-IS Router CAPABILITY TLV

The IS-IS Router CAPABILITY TLV is composed of 1 octet for the type, 1 octet that specifies the number of bytes in the value field, and a variable length value field that starts with 4 octets of Router ID, indicating the source of the TLV, and followed by 1 octet of flags.

A set of optional sub-TLVs may follow the flag field. Sub-TLVs are formatted as described in [RFC5305].

TYPE: 242  
 LENGTH: from 5 to 255  
 VALUE:  
   Router ID (4 octets)  
   Flags (1 octet)  
   Set of optional sub-TLVs (0-250 octets)

#### Flags

```

0 1 2 3 4 5 6 7
+-----+-----+
| Reserved |D|S|
+-----+-----+
  
```

Currently two bit flags are defined.

S bit (0x01): If the S bit is set(1), the IS-IS Router CAPABILITY TLV MUST be flooded across the entire routing domain. If the S bit is not set(0), the TLV MUST NOT be leaked between levels. This bit MUST NOT be altered during the TLV leaking.

D bit (0x02): When the IS-IS Router CAPABILITY TLV is leaked from level-2 to level-1, the D bit MUST be set. Otherwise, this bit MUST be clear. IS-IS Router capability TLVs with the D bit set MUST NOT be leaked from level-1 to level-2. This is to prevent TLV looping.

The Router CAPABILITY TLV is OPTIONAL. As specified in Section 3, more than one Router CAPABILITY TLV from the same source MAY be present.

This document does not specify how an application may use the Router Capability TLV and such specification is outside the scope of this document.

### 3. Elements of Procedure

The Router ID SHOULD be identical to the value advertised in the Traffic Engineering Router ID TLV [RFC5305]. If no Traffic Engineering Router ID is assigned the Router ID SHOULD be identical to an IP Interface Address [RFC1195] advertised by the originating IS. If the originating node does not support IPv4, then the reserved value 0.0.0.0 MUST be used in the Router ID field and the IPv6 TE Router ID sub-TLV [RFC5316] MUST be present in the TLV. Router CAPABILITY TLVs which have a Router ID of 0.0.0.0 and do NOT have the IPv6 TE Router ID sub-TLV present MUST be ignored.

When advertising capabilities with different flooding scopes, a router MUST originate a minimum of two Router CAPABILITY TLVs, each TLV carrying the set of sub-TLVs with the same flooding scope. For instance, if a router advertises two sets of capabilities, C1 and C2, with an area/level scope and routing domain scope respectively, C1 and C2 being specified by their respective sub-TLV(s), the router will originate two Router CAPABILITY TLVs:

- One Router CAPABILITY TLV with the S flag cleared, carrying the sub-TLV(s) relative to C1. This Router CAPABILITY TLV will not be leaked into another level.
- One Router CAPABILITY TLV with the S flag set, carrying the sub-TLV(s) relative to C2. This Router CAPABILITY TLV will be leaked into other IS-IS levels. When the TLV is leaked from level-2 to level-1, the D bit will be set in the level-1 LSP advertisement.

In order to prevent the use of stale capabilities, a system MUST NOT use a Capability TLV present in an LSP of a system that is not currently reachable via Level-x paths, where "x" is the level (1 or 2) in which the sending system advertised the TLV. This requirement applies regardless of whether or not the sending system is the originator of the Capabilities TLV. Note that leaking a Capabilities TLV is one of the uses that is prohibited under these conditions.

Example: If Level-1 router A generates a Capability TLV and floods it to two L1/L2 routers, S and T, they will flood it into the Level-2 domain. Now suppose the Level-1 area partitions, such that A and S are in one partition and T is in another. IP routing will still continue to work, but if A now issues a revised version of the CAP TLV, or decides to stop advertising it, S will follow suit, but T will continue to advertise the old version until the LSP times out.

Routers in other areas have to choose whether to trust T's copy of A's capabilities or S's copy of A's information and, they have no reliable way to choose. By making sure that T stops leaking A's information, this removes the possibility that other routers will use stale information from A.

In IS-IS, the atomic unit of the update process is a TLV - or more precisely, in the case of TLVs that allow multiple entries to appear in the value field (e.g., IS-neighbors), the atomic unit is an entry in the value field of a TLV. If an update to an entry in a TLV is advertised in an LSP fragment different from the LSP fragment associated with the old advertisement, the possibility exists that other systems can temporarily have either 0 copies of a particular advertisement or 2 copies of a particular advertisement, depending on

the order in which new copies of the LSP fragment that had the old advertisement and the fragment that has the new advertisement arrive at other systems.

Wherever possible, an implementation SHOULD advertise the update to a capabilities TLV in the same LSP fragment as the advertisement that it replaces. Where this is not possible, the two affected LSP fragments should be flooded as an atomic action.

Systems that receive an update to an existing capability TLV can minimize the potential disruption associated with the update by employing a holddown time prior to processing the update so as to allow for the receipt of multiple LSP fragments associated with the same update prior to beginning processing.

Where a receiving system has two copies of a capabilities TLV from the same system that have different settings for a given attribute, the procedure used to choose which copy shall be used is undefined.

#### 4. Interoperability with Routers Not Supporting the Capability TLV

Routers that do not support the Router CAPABILITY TLV MUST silently ignore the TLV(s) and continue processing other TLVs in the same LSP. Routers that do not support specific sub-TLVs carried within a Router CAPABILITY TLV MUST silently ignore the unsupported sub-TLVs and continue processing those sub-TLVs that are supported in the Router CAPABILITY TLV. How partial support may impact the operation of the capabilities advertised within the Router CAPABILITY TLV is outside the scope of this document.

In order for Router CAPABILITY TLVs with domain-wide scope originated by L1 Routers to be flooded across the entire domain, at least one L1/L2 Router in every area of the domain MUST support the Router CAPABILITY TLV.

If leaking of the CAPABILITY TLV is required, the entire CAPABILITY TLV MUST be leaked into another level even though it may contain some of the unsupported sub-TLVs.

#### 5. Security Considerations

Any new security issues raised by the procedures in this document depend upon the opportunity for LSPs to be snooped and modified, the ease/difficulty of which has not been altered. As the LSPs may now contain additional information regarding router capabilities, this new information would also become available to an attacker. Specifications based on this mechanism need to describe the security considerations around the disclosure and modification of their

information. Note that an integrity mechanism, such as the one defined in [RFC5304] or [RFC5310], should be applied if there is high risk resulting from modification of capability information.

## 6. IANA Considerations

IANA assigned a new IS-IS TLV code-point for the newly defined IS-IS TLV type named the IS-IS Router CAPABILITY TLV and defined in this document. The assigned value is 242.

## 7. Acknowledgements

For the original version of RFC 4971 the authors thanked Jean-Louis Le Roux, Paul Mabey, Andrew Partan, and Adrian Farrel for their useful comments.

For this new version the authors would like to thank Kris Michielsen for calling the problem associated w an IPv6 only router to our attention.

## 8. References

### 8.1. Normative References

- [ISO10589] International Organization for Standardization, "Intermediate system to Intermediate system intra-domain 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, Nov 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<http://www.rfc-editor.org/info/rfc1195>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5073] Vasseur, J., Ed. and J. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, DOI 10.17487/RFC5073, December 2007, <<http://www.rfc-editor.org/info/rfc5073>>.

- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<http://www.rfc-editor.org/info/rfc5310>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, DOI 10.17487/RFC5316, December 2008, <<http://www.rfc-editor.org/info/rfc5316>>.

## 8.2. Informational References

- [RFC4461] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, DOI 10.17487/RFC4461, April 2006, <<http://www.rfc-editor.org/info/rfc4461>>.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<http://www.rfc-editor.org/info/rfc4875>>.
- [RFC4972] Vasseur, JP., Ed., Leroux, JL., Ed., Yasukawa, S., Previdi, S., Psenak, P., and P. Mabbey, "Routing Extensions for Discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) Mesh Membership", RFC 4972, DOI 10.17487/RFC4972, July 2007, <<http://www.rfc-editor.org/info/rfc4972>>.

## Authors' Addresses

Les Ginsberg  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas, CA 95035  
USA

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

Stefano Previdi  
Cisco Systems  
Via Del Serafico 200  
Rome 0144  
Italy

Email: sprevidi@cisco.com

Mach (Guoyi) Chen  
Huawei Technologies Co., Ltd  
KuiKe Building, No. 9 Xixi Rd. Hai-Dian District  
Beijing 100085  
P.R. China

Email: mach.chen@huawei.com

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: October 1, 2018

L. Ginsberg, Ed.  
Cisco Systems  
A. Przygienda  
Juniper Networks  
S. Aldrin  
Google  
J. Zhang  
Juniper Networks, Inc.  
March 30, 2018

BIER support via ISIS  
draft-ietf-bier-isis-extensions-11

Abstract

This document defines ISIS extensions to support multicast forwarding using the Bit Index Explicit Replication (BIER) architecture.

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.

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 October 1, 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 the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

|  |    |
|--|----|
| 1. Introduction                          | 2  |
| 2. Terminology                           | 3  |
| 3. IANA Considerations                   | 4  |
| 4. Concepts                              | 4  |
| 4.1. BIER Domains and Sub-Domains        | 5  |
| 4.2. Advertising BIER Information        | 5  |
| 5. Procedures                            | 5  |
| 5.1. Multi Topology and Sub-Domain       | 5  |
| 5.2. BFR-id Advertisements               | 6  |
| 5.3. Logging Misconfiguration            | 6  |
| 5.4. Flooding Reduction                  | 6  |
| 6. Packet Formats                        | 6  |
| 6.1. BIER Info sub-TLV                   | 7  |
| 6.2. BIER MPLS Encapsulation sub-sub-TLV | 8  |
| 7. Security Considerations               | 9  |
| 8. Acknowledgements                      | 9  |
| 9. References                            | 10 |
| 9.1. Normative References                | 10 |
| 9.2. Informative References              | 11 |
| Authors' Addresses                       | 11 |

## 1. Introduction

Bit Index Explicit Replication (BIER) [RFC8279] defines an architecture where all intended multicast receivers are encoded as bitmask in the Multicast packet header within different encapsulations such as [RFC8296]. A router that receives such a packet will forward the packet based on the Bit Position in the packet header towards the receiver(s), following a precomputed tree for each of the bits in the packet. Each receiver is represented by a unique bit in the bitmask.

This document presents necessary extensions to the currently deployed ISIS for IP [RFC1195] protocol to support distribution of information necessary for operation of BIER domains and sub-domains. This document defines a new TLV to be advertised by every router participating in BIER signaling.

This document defines support for MPLS encapsulation as specified in [RFC8296]. Support for other encapsulation types is outside the scope of this document. The use of multiple encapsulation types is outside the scope of this document.

## 2. Terminology

Some of the terminology specified in [RFC8279] is replicated here and extended by necessary definitions:

**BIER:** Bit Index Explicit Replication (The overall architecture of forwarding multicast using a Bit Position).

**BIER-OL:** BIER Overlay Signaling. (The method for the BFIR to learn about BFER's).

**BFR:** Bit Forwarding Router (A router that participates in Bit Index Multipoint Forwarding). A BFR is identified by a unique BFR-prefix in a BIER domain.

**BFIR:** Bit Forwarding Ingress Router (The ingress border router that inserts the BM into the packet). Each BFIR must have a valid BFR-id assigned.

**BFER:** Bit Forwarding Egress Router. A router that participates in Bit Index Forwarding as leaf. Each BFER must be a BFR. Each BFER must have a valid BFR-id assigned.

**BFT:** Bit Forwarding Tree used to reach all BFERs in a domain.

**BIER sub-domain:** A further distinction within a BIER domain identified by its unique sub-domain identifier. A BIER sub-domain can support multiple BitString Lengths.

**BFR-id:** An optional, unique identifier for a BFR within a BIER sub-domain.

**Invalid BFR-id:** Unassigned BFR-id. The special value 0 is reserved for this purpose.

**BAR** BIER Algorithm. Used to calculate underlay next hops.

IPA IGP Algorithm. May be used to modify, enhance or replace the calculation of underlay paths as defined by the BAR value

SPF Shortest Path First routing calculation based on IGP link metric

### 3. IANA Considerations

This document adds the following new sub-TLV to the registry of Sub-TLVs for TLVs 135, 235, 236, and 237.

Value: 32 (suggested - to be assigned by IANA)

Name: BIER Info

This document also introduces a new registry for sub-sub-TLVs for the BIER Info sub-TLV added above. The registration policy is Expert Review as defined in [RFC8126]. This registry is part of the "IS-IS TLV Codepoints" registry. The name of the registry is "sub-sub-TLVs for BIER Info sub-TLV". The defined values are:

| Type | Name                    |
|------|-------------------------|
| ---- | ----                    |
| 1    | BIER MPLS Encapsulation |

IANA is requested to set up a registry called "BIER Algorithm Registry" under category "Bit Index Explicit Replication". The registration policies [RFC8126] for this registry are:

"Standards Action" for values 0-127

"Specification Required" for values 128-240

"Experimental Use" for values 240-254"

The initial values in the BIER Algorithm Registry are:

0: No BIER specific algorithm is used

1-254: Unassigned

255: Reserved

### 4. Concepts

#### 4.1. BIER Domains and Sub-Domains

An ISIS signalled BIER domain is aligned with the scope of distribution of BFR-prefixes that identify the BFRs within ISIS. ISIS acts in such a case as the supporting BIER underlay.

Within such a domain, the extensions defined in this document advertise BIER information for one or more BIER sub-domains. Each sub-domain is uniquely identified by a subdomain-id (SD). Each subdomain is associated with a single ISIS topology (MT) [RFC5120], which may be any of the topologies supported by ISIS. Local configuration controls which <MT,SD> pairs are supported by a router. The mapping of sub-domains to topologies MUST be consistent within the IS-IS flooding domain used to advertise BIER information.

Each BIER sub-domain has as its unique attributes the encapsulation used and the type of tree it is using to forward BIER frames (currently always SPF). Additionally, per supported bitstring length in the sub-domain, each router will advertise the necessary label ranges to support it.

#### 4.2. Advertising BIER Information

BIER information advertisements are associated with a new sub-TLV in the extended reachability TLVs. BIER information is always associated with a host prefix which MUST be a node address for the advertising node. If this is not the case the advertisement MUST be ignored. Therefore the following restrictions apply:

- o Prefix length MUST be 32 for an IPv4 prefix or 128 for an IPv6 prefix
- o When the Prefix Attributes Flags sub-TLV is present N flag MUST be set and R flag MUST NOT be set. [RFC7794]
- o BIER sub-TLVs MUST be included when a prefix reachability advertisement is leaked between levels.

### 5. Procedures

#### 5.1. Multi Topology and Sub-Domain

A given sub-domain is supported within one and only one topology. All routers in the flooding scope of the BIER sub-TLVs MUST advertise the same sub-domain within the same multi-topology. A router receiving an <MT,SD> advertisement which does not match the locally configured pair MUST report a misconfiguration of the received <MT,SD> pair. All received BIER advertisements associated with the

conflicting <MT,SD> pair MUST be ignored. Note that in the presence of such a misconfiguration this will lead to partitioning of the sub-domain.

Example:

The following combination of advertisements are valid: <0,0> <0,1> <2,2>.

The following combination of advertisements are invalid: <0,0> <0,1> <2,0>. Advertisements associated with <0,0> and <2,0> must be ignored.

## 5.2. BFR-id Advertisements

If a BFER/BFIR is configured with a BFR-id then it advertises this value in its BIER advertisements. If no BFR-id is configured then the value "Invalid BFR-id" is advertised. A valid BFR-id MUST be unique within the flooding scope of the BIER advertisements. All BFERs/BFIRs MUST detect advertisement of duplicate valid BFR-IDs for a given <MT, SD>. When such duplication is detected all of the routers advertising duplicates MUST be treated as if they did not advertise a valid BFR-id. This implies they cannot act as BFER or BFIR in that <MT,SD>.

## 5.3. Logging Misconfiguration

Whenever an advertisement is received which violates any of the constraints defined in this document the receiving router MUST support logging this occurrence. Logging SHOULD be dampened to avoid excessive output.

## 5.4. Flooding Reduction

It is expected that changes in BIER domain information which is advertised by IS-IS occur infrequently. If this expectation is not met for an extended period of time (more than a few seconds of burstiness) changes will increase the number of Link State PDU (LSP) updates and negatively impact performance in the network. Implementations SHOULD protect against this possibility e.g., by dampening updates if they occur over an extended period of time.


## 6. Packet Formats

All ISIS BIER information is carried within the TLVs 235, 237 [RFC5120] or TLVs 135 [RFC5305], or TLV 236 [RFC5308].

6.1. BIER Info sub-TLV

This sub-TLV carries the information for the BIER sub-domains that the router participates in as BFR. This sub-TLV MAY appear multiple times in a given prefix-reachability TLV - once for each sub-domain supported in the associated topology.

The sub-TLV advertises a single <MT,SD> combination followed by optional sub-sub-TLVs as described in the following sections.



Type: as indicated in IANA section.

Length: variable

BAR BIER Algorithm. Specifies a BIER specific algorithm used to calculate underlay paths to reach BFRs. Values are allocated from the BIER Algorithm Registry. 1 octet

IPA IGP algorithm. Specifies an IGP Algorithm to either modify, enhance or replace the calculation of underlay paths to reach BFRs as defined by the BAR value. Values are from the IGP Algorithm registry. 1 octet

subdomain-id: Unique value identifying the BIER sub-domain. 1 octet

BFR-id: A 2 octet field encoding the BFR-id, as documented in [RFC8279]. If no BFR-id has been assigned the value of this field is set to "Invalid BFR-id", which is defined as illegal in [RFC8279] .

The use of non-zero values in either the BAR field or the IPA field is outside the scope of this document. If an implementation does not support the use of non-zero values in these fields, but receives a BIER Info sub-TLV containing non-zero values in these fields, it

SHOULD treat the advertising router as incapable of supporting BIER (one way of handling incapable routers is documented in section 6.9 of [RFC8279] and additional methods may be defined in the future).

6.2. BIER MPLS Encapsulation sub-sub-TLV

This sub-sub-TLV carries the information for the BIER MPLS encapsulation including the label range for a specific bitstring length for a certain <MT,SD>. It is advertised within the BIER Info sub-TLV (Section 6.1) . This sub-sub-TLV MAY appear multiple times within a single BIER info sub-TLV.

If the same Bitstring length is repeated in multiple sub-sub-TLVs inside the same BIER Info Sub-TLV, the BIER Info sub-TLV MUST be ignored.

Label ranges within all BIER MPLS Encapsulation sub-sub-TLVs across all BIER Info sub-TLVs advertised by the same BFR MUST NOT overlap. If overlap is detected, the advertising router MUST be treated as if it did not advertise any BIER sub-TLVs.

Label values MUST NOT match any of the reserved values defined in [RFC3032]

|        |   |   |   |   |   |   |   |   |   |        |   |   |   |   |   |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|--------|---|---|---|---|---|---|---|---|---|--------|---|---|---|---|---|---|---|---|---|-------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
|        |   |   |   |   |   |   |   |   |   | 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 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Type   |   |   |   |   |   |   |   |   |   | Length |   |   |   |   |   |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Max SI |   |   |   |   |   |   |   |   |   | BS Len |   |   |   |   |   |   |   |   |   | Label |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |

Type: value of 1 indicating MPLS encapsulation.

Length: 4

Max SI Maximum Set Identifier (section 1 of [RFC8279]) used in the encapsulation for this BIER sub-domain for this bitstring length, 1 octet. Each SI maps to a single label in the label range. The first label is for SI=0, the second label is for SI=1, etc. If the label associated with the Maximum Set Identifier exceeds the 20 bit range the sub-sub-TLV MUST be ignored.

Local BitString Length (BS Len): Encoded bitstring length as per [RFC8296]. 4 bits.

Label: First label of the range, 20 bits. The labels are as defined in [RFC8296].

## 7. Security Considerations

Security concerns for IS-IS are addressed in [RFC5304] and [RFC5310].

The Security Considerations section of [RFC8279] discusses the possibility of performing a Denial of Service (DoS) attack by setting too many bits in the BitString of a BIER-encapsulated packet. However, this sort of DoS attack cannot be initiated by modifying the ISIS BIER advertisements specified in this document. A BFIR decides which systems are to receive a BIER-encapsulated packet. In making this decision, it is not influenced by the ISIS control messages. When creating the encapsulation, the BFIR sets one bit in the encapsulation for each destination system. The information in the ISIS BIER advertisements is used to construct the forwarding tables that map each bit in the encapsulation into a set of next hops for the host that is identified by that bit, but is not used by the BFIR to decide which bits to set. Hence an attack on the ISIS control plane cannot be used to cause this sort of DoS attack.

While a BIER-encapsulated packet is traversing the network, a BFR that receives a BIER-encapsulated packet with  $n$  bits set in its BitString may have to replicate the packet and forward multiple copies. However, a given bit will only be set in one copy of the packet. That means that each transmitted replica of a received packet has fewer bits set (i.e., is targeted to fewer destinations) than the received packet. This is an essential property of the BIER forwarding process as defined in [RFC8279]. While a failure of this process might cause a DoS attack (as discussed in the Security Considerations of [RFC8279]), such a failure cannot be caused by an attack on the ISIS control plane.

Further discussion of BIER specific security considerations can be found in [RFC8279].

## 8. Acknowledgements

The RFC is aligned with the [I-D.ietf-bier-ospf-bier-extensions] draft as far as the protocol mechanisms overlap.

Many thanks for comments from (in no particular order) Hannes Gredler, Ijsbrand Wijnands, Peter Psenak and Chris Bowers.

Special thanks to Eric Rosen.

## 9. References

### 9.1. Normative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [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>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [RFC7794] Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability", RFC 7794, DOI 10.17487/RFC7794, March 2016, <<https://www.rfc-editor.org/info/rfc7794>>.

- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.

## 9.2. Informative References

- [I-D.ietf-bier-ospf-bier-extensions]  
Psenak, P., Kumar, N., Wijnands, I., Dolganow, A., Przygienda, T., Zhang, Z., and S. Aldrin, "OSPFv2 Extensions for BIER", draft-ietf-bier-ospf-bier-extensions-16 (work in progress), March 2018.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

## Authors' Addresses

Les Ginsberg (editor)  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas, CA 95035  
USA

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

Tony Przygienda  
Juniper Networks

Email: [prz@juniper.net](mailto:prz@juniper.net)

Sam Aldrin  
Google  
1600 Amphitheatre Parkway  
Mountain View, CA  
USA

Email: [aldrin.ietf@gmail.com](mailto:aldrin.ietf@gmail.com)

Jeffrey (Zhaohui) Zhang  
Juniper Networks, Inc.  
10 Technology Park Drive  
Westford, MA 01886  
USA

Email: [zzhang@juniper.net](mailto:zzhang@juniper.net)

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: October 19, 2016

C. Franke  
D. Lamparter  
NetDEF  
C. Hopps  
Deutsche Telekom  
April 17, 2016

IS-IS Point-to-Multipoint operation  
draft-lamparter-isis-p2mp-02

Abstract

This document describes a new mode operation for IS-IS. In addition to the existing LAN and point-to-point modes of operation, a point-to-multipoint mode is defined. This mode is useful for operation both on networks with more than two routers where multicast is expensive in comparison to unicast, as well on networks where creating a LAN pseudonode cannot adequately reflect metrics.

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 <http://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 October 19, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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 |
| 2. Point-To-Multipoint Pseudocircuits . . . . .          | 3 |
| 2.1. Pseudocircuit behaviour . . . . .                   | 3 |
| 2.1.1. Representation in LSPs . . . . .                  | 3 |
| 2.1.2. Forwarding . . . . .                              | 3 |
| 2.2. Neighbor IS discovery . . . . .                     | 3 |
| 2.2.1. Manual configuration . . . . .                    | 4 |
| 2.2.2. Lower layer autodiscovery . . . . .               | 4 |
| 2.2.3. Multicast autodiscovery . . . . .                 | 4 |
| 2.3. Adjacency formation . . . . .                       | 5 |
| 2.4. Pseudocircuit teardown . . . . .                    | 5 |
| 3. Configuration model . . . . .                         | 5 |
| 4. Security Considerations . . . . .                     | 5 |
| 5. Privacy Considerations . . . . .                      | 5 |
| 6. Acknowledgements . . . . .                            | 5 |
| 7. Change Log . . . . .                                  | 5 |
| 8. References . . . . .                                  | 6 |
| 8.1. Normative References . . . . .                      | 6 |
| 8.2. Informative References . . . . .                    | 6 |
| Appendix A. Misconfiguration With P2P over LAN . . . . . | 6 |
| Authors' Addresses . . . . .                             | 6 |

## 1. Introduction

The core functionality of the protocol outlined in this document consists of an additional Subnetwork dependent function per Section 8 of ISO/IEC 10589:2002 [IS-IS]. In that regard, the next section can be understood as new section "8.5 Point-to-multipoint networks".

The outlined protocol is remotely similar to the derelict section "8.3 ISO 8208 subnetworks" [X.25] in that multiple point-to-point adjacencies are established on an interface.

Point-to-multipoint operation of IS-IS is thus not a new idea; in fact section 6.2 already mentions "multipoint links." This document re-enables the concept.

## 2. Point-To-Multipoint Pseudocircuits

In place of ISO 8473 call management [CLNS] establishing sessions, this document establishes pairwise "pseudocircuits" between two IS on a common medium. Multiple such pseudocircuits can normally exist on one P2MP (Point-To-Multipoint) interface.

Each pseudocircuit operates, aside from subnetwork attachment procedures, exactly as a P2P (Point-to-Point) link.

It should be noted that while the Multicast autodiscovery mechanism requires broadcast capability, no other part of the protocol does. The P2MP interface type can be used on non-transitive and/or non-broadcast interfaces.

### 2.1. Pseudocircuit behaviour

An implementation maintains a set of pseudocircuits per P2MP interface. Each pseudocircuit is uniquely identified by the local interface and peer's SNPA address.

Each participating system MUST use a consistent SNPA address per local interface. A change in SNPA address results in all neighbors treating the interface as distinct from the previous one. A system MAY support multiple SNPA addresses per interface by treating them as distinct interfaces.

Unnumbered media are not supported by this protocol. Each participant on a link MUST have a unique SNPA address on that link.

#### 2.1.1. Representation in LSPs

Pseudocircuits are represented in LSPs as a regular P2P circuit would be. As a result, their treatment in SPF calculations is also identical to P2P circuits.

#### 2.1.2. Forwarding

In scenarios where pseudocircuits do not form a full mesh of all participants on a medium, the best path for a packet may be through the same interface as the one it was received on.

Systems implementing this specification MUST therefore support forwarding packets on the same interface that they were received on. This applies only to interfaces configured for P2MP operation.

## 2.2. Neighbor IS discovery

The discovery machinery produces as output a "candidate neighbor list," which is a list of possible neighbor's SNPAs and is maintained per P2MP interface. Adding and removing entries to the candidate neighbor list results in pseudocircuit creation and deletion.

A neighbors presence on the candidate list may be supported by multiple sources. These sources are described in the following sections

The IS-IS implementation SHOULD provide user configuration to disable or filter individual sources.

#### 2.2.1. Manual configuration

A list of neighbor IS MAY be configured by the user, providing possible neighbor's SNPAs on an interface.

#### 2.2.2. Lower layer autodiscovery

Lower protocol layers (VPLS, IEEE 802.11) may be able to provide a list of attached neighbors. This list MAY be fed into the candidate neighbor list.

#### 2.2.3. Multicast autodiscovery

For broadcast capable networks, this document defines an autodiscovery mechanism based on multicasting Hello packets. This mechanism MAY be disabled by the user, but MUST be implemented for all lower layer link types with (limited or full) broadcast capability.

A multicast autodiscovery packet is defined as a P2P IIH which is multicast over the LAN media as defined in [RFC5309]. Additionally it must include a Three-Way Adjacency TLV as defined in [RFC5303] indicating the circuit is in the DOWN state (i.e., 1 octet of value indicating the DOWN state). Receiving such a Hello places the sending IS on the candidate list for as long as the PDU's holdtime indicates.

A system MAY implement a receive-only passive multicast autodiscovery mode where it adds candidates (forms pseudocircuits) on receiving autodiscovery PDU, but does not send such PDUs itself.

If either passive or full multicast autodiscovery is enabled, receiving a unicast autodiscovery PDU also adds the neighbor to the candidate list.

### 2.3. Adjacency formation

Since there may be no underlying protocol layer partitioning the link into actual P2P circuits in this case, additional handling is required on bringing up the individual pseudocircuit adjacencies.

To prevent different pseudocircuits from "bleeding" into each other, implementations of this protocol MUST enable [RFC5303] on all P2MP pseudocircuits, with changes as follows:

- o Received IIH PDUs on P2MP pseudocircuits without the Point-to-Point Three-Way Adjacency option MUST be discarded.

### 2.4. Pseudocircuit teardown

Pseudocircuits are destroyed when their P2P state machine has transitioned into "Down" state and they are no longer supported as a candidate by any discovery mechanism.

As long as a pseudocircuit is present, its P2P state machine will, as expected for P2P circuits, trigger transmission of further Hello PDUs, even when it is in Down state.

## 3. Configuration model

TODO: YANG model

## 4. Security Considerations

TODO.

## 5. Privacy Considerations

TODO.

## 6. Acknowledgements

Acknowledge Les Ginsberg for pointing out that P2P Hellos rather than LAN hellos could and should be used.

## 7. Change Log

April 2016 [-02]: (no changes/keepalive)

October 2015 [-01]: Moved from new P2MP Hello PDU to using existing P2P PDU.

July 2015 [-00]: Initial Version

## 8. References

### 8.1. Normative References

- [IS-IS] ISO/IEC, "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, Second Edition, 2002.
- [RFC5303] Katz, D., Saluja, R., and D. Eastlake, "Three-Way Handshake for IS-IS Point-to-Point Adjacencies", RFC 5303, October 2008.
- [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, <<http://www.rfc-editor.org/info/rfc5309>>.

### 8.2. Informative References

- [CLNS] ISO/IEC, "Protocol for providing the connectionless-mode network service: Protocol specification", ISO/IEC 8473-1:1998, 1998.
- [RFC7176] Eastlake, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176, May 2014.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, September 2014.
- [X.25] ISO/IEC, "X.25 Packet Layer Protocol for Data Terminal Equipment", ISO/IEC 8208:2000, 2000.

## Appendix A. Misconfiguration With P2P over LAN

TODO.

### Authors' Addresses

Christian Franke  
NetDEF  
Leipzig  
DE

Email: [chris@opensourcerouting.org](mailto:chris@opensourcerouting.org)

David Lamparter  
NetDEF  
Leipzig 04229  
Germany

Email: david@opensourcerouting.org

Christian E. Hopps  
Deutsche Telekom

Email: chopps@chopps.org

isis  
Internet-Draft  
Intended status: Standards Track  
Expires: April 21, 2016

B. Liu, Ed.  
Huawei Technologies  
B. Decraene  
Orange  
I. Farrer  
Deutsche Telekom AG  
M. Abrahamsson  
T-Systems  
L. Ginsberg  
Cisco Systems  
October 19, 2015

ISIS Auto-Configuration  
draft-liu-isis-auto-conf-06

Abstract

This document specifies an IS-IS auto-configuration technology. The key mechanisms of this technology are IS-IS System ID self-generation, duplication detection and duplication resolution. This technology fits the environment where plug-and-play is expected.

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 <http://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 April 21, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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  |
| 2. Scope . . . . .   | 3  |
| 3. Protocol Specification . . . . .  | 3  |
| 3.1. IS-IS Default Configuration . . . . .                                       | 3  |
| 3.2. IS-IS NET Generation . . . . .  | 3  |
| 3.3. IS-IS System ID Duplication Detection and Resolution . .                    | 4  |
| 3.3.1. Router-Fingerprint TLV . . . . .  | 4  |
| 3.3.2. System ID Duplication Detection and Resolution<br>Procedures . . . . .    | 5  |
| 3.3.3. System ID and Router-Fingerprint Generation<br>Considerations . . . . .   | 9  |
| 3.3.4. Double-Duplication of both System ID and Router-<br>Fingerprint . . . . . | 10 |
| 3.4. IS-IS TLVs Usage . . . . .  | 11 |
| 3.4.1. Authentication TLV . . . . .  | 11 |
| 3.4.2. Wide Metric TLV . . . . .   | 11 |
| 3.4.3. Dynamic Host Name TLV . . . . .   | 11 |
| 3.5. Routing Behavior Considerations . . . . .                                   | 12 |
| 3.5.1. Adjacency Formation . . . . .   | 12 |
| 4. Security Considerations . . . . .   | 12 |
| 5. IANA Considerations . . . . .   | 12 |
| 6. Acknowledgements . . . . .  | 12 |
| 7. References . . . . .  | 13 |
| 7.1. Normative References . . . . .  | 13 |
| 7.2. Informative References . . . . .  | 13 |
| Authors' Addresses . . . . .   | 14 |

## 1. Introduction

This document describes mechanisms for IS-IS [RFC1195] [ISO\_IEC10589][RFC5308] to be auto-configuring. Such mechanisms could reduce the management burden to configure a network. Home networks and small or medium size enterprise networks where plug-and-play is expected can benefit from these mechanisms.

This document also defines mechanisms which prevent unintentional interoperation of autoconfigured routers with non-autoconfigured routers. See Section 3.3.1 .

IS-IS auto-configuration contains the following aspects:

1. IS-IS default configurations
2. IS-IS System ID self-generation
3. System ID duplication detection and resolution
4. ISIS TLVs utilization such as Authentication TLV, Wide Metric TLV etc.

## 2. Scope

The auto-configuring mechanisms support both IPv4 and IPv6 deployments.

This auto-configuration mechanism aims at simple case. The following advanced features are out of scope:

- o Multiple IS-IS instances
- o Multi-area and level-2 routing
- o Interworking with other routing protocols

## 3. Protocol Specification

### 3.1. IS-IS Default Configuration

- o IS-IS interfaces MUST be auto-configured to an interface type corresponding to their layer-2 capability. For example, Ethernet interfaces will be auto-configured as broadcast networks and Point-to-Point Protocol (PPP) interfaces will be auto-configured as Point-to-Point interfaces.
- o IS-IS auto-configuration instance MUST be configured with level-1, so that the interfaces operate at level-1 only.
- o IS-IS auto-configuration SHOULD allow P2P mode on Ethernet interfaces.

### 3.2. IS-IS NET Generation

In IS-IS, a router (known as an Intermediate System) is identified by an NET which is the address of a Network Service Access Point (NSAP) and represented with an IS-IS specific address format. The NSAP is a logical entity which represents an instance of the IS-IS protocol running on an Intermediate System.

The autoconfiguration mechanism generates the IS-IS NET as the following:

- o Area address

- This field is 1 to 13 octets in length. In IS-IS auto-configuration, this field MUST be 13 octets of all 0.

- o System ID

- This field follows the area address field, and is 6 octets in length. There are two basic requirements for the System ID generation:

- As specified in IS-IS protocol, this field must be unique among all routers in the same area.
    - In order to make the routing system stable, the System ID SHOULD remain the same after it is firstly generated. It SHOULD not be changed due to device status change (such as interface enable/disable, interface plug in/off, device reboot, firmware update etc.) or configuration change (such as changing system configurations or IS-IS configurations etc.); but it MUST allow be changed by collision resolution and SHOULD allow be cleared by user enforced system reset.


- More specific considerations for System ID generation are described in Section 3.3.3 .

### 3.3. IS-IS System ID Duplication Detection and Resolution

The System ID of each node MUST be unique. As described in Section 3.3.3, the System ID is generated based on entropies such as MAC address which are supposed to be unique, but in theory there is still possibility of duplication. This section defines how IS-IS detects and resolves System ID duplication.

#### 3.3.1. Router-Fingerprint TLV

The Router-Fingerprint TLV basically re-uses the design of Router-Hardware-Fingerprint TLV defined in [RFC7503]. However, there is one difference that one flag is added to indicate the node is in "start-up mode" which is defined in Section 3.3.2 .



The length of the Router-Fingerprint is variable but must be 32 octets or greater; and the content is also supposed to be unique among all the routers.

- o Type: to be assigned by IANA.
- o Length: the length of the value field.
- o S flag: indicates the router is in "start-up" mode as described below.
- o A flag: indicates the router is operating in autoconfiguration mode. This flag is in case the TLV gets used outside of autoconfiguration. If A flag setting does not match in hellos then no adjacency should be formed.
- o Reserved: these bits MUST be set to zero and MUST be ignored when received.
- o Router Fingerprint: uniquely identifies a router, variable length.

More specific considerations for Router-Fingerprint is described in Section 3.3.3 .

### 3.3.2. System ID Duplication Detection and Resolution Procedures

This section describes the System ID duplication detection and resolution between two neighbors and two non-neighbors respectively. This is because the routing messages between neighbors and non-neighbors are a bit different.

#### 3.3.2.1. Start-up Mode

While in startup-mode, an auto-configuration router forms adjacencies but generates only LSP #0 which contains only the Router-Fingerprint TLV. A router remains in startup-mode until it has successfully completed LSPDB synchronization with all neighbors or until 1 minute

has elapsed - whichever is longer. If duplicate system-ID is detected while in startup-mode the router MUST clear all adjacencies, select a new system-id (subject to rules defined in Section 3.3.2.2), and reenter Startup-mode.

The start-up mode is to minimize the occurrence of System ID changes for a router once it has become fully operational. It has minimal impact on a running network because the startup node is not yet being used for forwarding traffic. Once duplicate System ID has been resolved the router begins normal operation. If two routers are both in startup mode (or both NOT in startup mode) and duplicate system-id is detected then they determine which one changes its system-id based on fingerprint.

When an IS-IS auto-configuration router boots up, it MUST operate in start-up mode until duplicate system-id detection has successfully completed.

### 3.3.2.2. Duplication Between Neighbors

In case of System ID duplication occurs between neighbors, an IS-IS auto-configuration router MUST include the Router-Fingerprint TLV in the Hello messages, so that the duplication could be detected before adjacency forming.

Procedures of the nodes in Start-up Mode:

1. Boot up, advertise the Router-Fingerprint TLV in Hello message

The router sends Hellos which include the Router-Fingerprint TLV. Adjacencies are formed as normal but MUST NOT be advertised in LSPs until the router exits startup-mode.

2. Receive Hello message(s), and verifies System ID duplication

Received hellos are inspected for possible duplicate System ID. If duplication is detected, the router MUST check the S flag of the Router-Fingerprint TLV.

- + If the S flag is NOT set (which means the Hello was NOT generated by a neighbor also in Start-up mode), then the router MUST re-generate the System ID and reenter Startup-mode.
- + If the S flag is set (which means the neighbor is also in Startup-mode),

- the router which has a numerically smaller Router-Fingerprint MUST re-generate the System ID and reenter Startup-mode. Fingerprint comparison is performed octet by octet until octets are different. Then the smaller fingerprint is the one with the smaller octet (unsigned integer). If the fingerprints have different lengths, then the shorter length fingerprint MUST be padding with zero for comparison.
- If Router Fingerprints are identical, both routers MUST re-generate the System ID and the Router Fingerprint, and reenter Startup-mode.

### 3. Run in normal operation

After the System ID duplication procedure is done, the router begins to run in normal operation. The router MUST re-advertise the Router-Fingerprint TLV with the S flag off.

Procedures of the nodes NOT in Start-up Mode:

#### 1. Compare the System ID in received Hello messages

When receiving a Hello message, the router MUST check the System ID of the Hello. If the System ID is the same as its own, it indicates a System ID duplication occurs.

If there is no Router-Fingerprint TLV in the Hello message, it means a non-autoconfiguration router by accident connected to the auto-configuration domain or other unexpected bad behaviors. In this case, the auto-configuration router MUST NOT form adjacency with the non-autoconfiguration router.

#### 2. Duplication resolution

When System ID duplication occurs, the non-startup mode router MUST check the S flag of the duplicated Router-Fingerprint TLV:

- + If the S flag is NOT set, then the router with the numerically smaller or equal Router-Fingerprint MUST generate a new System ID. Note that, the router MUST compare the two Router-Fingerprint in terms of two numeric numbers.
- + If the S flag is set, then router does nothing, because it MUST be the node which is in start-up mode re-generates the System ID.

3. Re-join the network with the new System ID (if required)

The router with the smaller Router-Fingerprint advertise new Hellos based on the newly generated NET to re-join the IS-IS auto-configuration network. The router with the highest Router-Fingerprint MUST re-advertise its own LSP (after increasing the sequence number).

The newly generated System ID SHOULD take a duplication detection as well.

- 3.3.2.3. Duplication Between Non-neighbors

System ID duplication may also occur between non-neighbors, so an IS-IS auto-configuration router MUST also include the Router-Fingerprint TLV in the LSP messages. Specific procedures are as the following.

Procedures of the nodes in Start-up Mode:

1. Boot up, form adjacency
2. Acquire LSPDB and verifies System ID duplication

The router generates only LSP #0 which contains only the Fingerprint TLV; and that Fingerprint is only sent in LSP #0. A router remains in startup-mode until it has successfully completed LSPDB synchronization with all neighbors or until 1 minute has elapsed - whichever is longer. If duplicate system-ID is detected, the router MUST check the S flag of the Router-Fingerprint TLV of the LSP that contains the duplicated System ID.

- + If the S flag is not set, it means the LSP was not generated at the Start-up Mode, then the router itself MUST clear all adjacencies, re-generate a new system-id and reenter Startup-mode.
- + If the S flag is set, then the router which has a numerically smaller Router-Fingerprint MUST generate a new System ID and reenter Startup-mode.

3. Run in normal operation

After the System ID duplication procedure is done, the router begins to run in normal operation. The router MUST re-advertise the Router-Fingerprint TLV with the S flag off.

Procedures of the nodes not in Start-up Mode:

1. Compare the received Router-Fingerprint TLVs

When receiving a LSP containing its own System ID, the router MUST check the Router-Fingerprint TLV. If the Router-Fingerprint TLV is different from its own, it indicates a System ID duplication occurs.

2. Duplication resolution

When System ID duplication occurs, the non-startup mode router MUST check the S flag of the duplicated Router-Fingerprint TLV:

- + If the S flag is NOT set, then the router with the numerically smaller Router-Fingerprint MUST generate a new System ID. Note that, the router MUST compare the two Router-Fingerprint in terms of two numeric numbers.
- + If the S flag is set, then router does nothing, because according to the start-up mode procedure, the start-up node MUST re-generate the System ID.

3. Re-join the network with the new System ID

The router changing its system ID advertise new LSPs based on the newly generated System ID to re-join the IS-IS auto-configuration network. The router with the highest Router-Fingerprint MUST re-advertise its own LSP (after increasing the sequence number).

The newly generated System SHOULD take a duplication detection as well.

3.3.3. System ID and Router-Fingerprint Generation Considerations

As specified in this document, there are two distinguisher need to be self-generated, which is System ID and Router-Fingerprint. In a network device, normally there are resources which provide an extremely high probability of uniqueness thus could be used as seeds to derive distinguisher (e.g. hashing or generating pseudo-random numbers), such as:

- o MAC address(es)
- o Configured IP address(es)

- o Hardware IDs (e.g. CPU ID)
- o Device serial number(s)
- o System clock at a certain specific time
- o Arbitrary received packet

This document recommends to use an IEEE 802 48-bit MAC address associated with the router as the initial System ID. This document does not specify a specific method to re-generate the System ID when duplication happens.

This document also does not specify a specific method to generate the Router-Fingerprint. However, the generation of System ID and Router-Fingerprint MUST be based on different seeds so that the two distinguisher would not collide.

There is an important concern that the seeds listed above (except MAC address) might not be available in some small devices such as home routers. This is because of the hardware/software limitation and the lack of sufficient communication packets at the initial stage in the home routers when doing ISIS-autoconfiguration. In this case, this document suggests to use MAC address as System ID and generate a pseudo-random number based on another seed (such as the memory address of a certain variable in the program) as Router-Fingerprint. The pseudo-random number might not have a very high quality in this solution, but should be sufficient in home networks scenarios.

Note that, the Router-Fingerprint SHOULD also remain the same after it is firstly generated. It SHOULD not be changed due to device status change (such as interface enable/disable, interface plug in/off, device reboot, firmware update etc.) or configuration change (such as changing system configurations or IS-IS configurations etc.); but it MUST allow be changed by double-duplication resolution Section 3.3.4 and SHOULD allow be cleared by user enforced system reset.

#### 3.3.4. Double-Duplication of both System ID and Router-Fingerprint

As described above, the resources for generating the distinguisher might be very constrained at the initial stage. Hence, the double-duplication of both System ID and Router-Fingerprint needs to be considered.

ISIS-autoconfiguring routers SHOULD support detecting System ID duplication by LSP war. LSP war is a phenomenon that if a router receives a LSP originated with its System ID, but it doesn't find it

in the database, or it does not match the one the router has (e.g. It advertises IP prefixes that the router doesn't own, or IS neighbors that the router doesn't see), then per ISIS specification, the router must re-originate its LSP with an increased sequence number. If double-duplication happens, the duplicated two routers will both continuously have the above behavior. After multiples iterations, the program should be able to deduce that double-duplication happens.

At the point when double-duplication happens, routers should have much more entropies available. Thus, the router is to extend or re-generate its Router-Fingerprint (one simple way is just adding the LSP sequence number of the next LSP it will send to the Router-Fingerprint). (Optimized solution TBD.)

### 3.4. IS-IS TLVs Usage

This section describes several TLVs that are utilized by IS-IS auto-configuration.

#### 3.4.1. Authentication TLV

It is RECOMMENDED that IS-IS routers supporting this specification minimally offer an option to explicitly configure a single password for HMAC-MD5 authentication, which is Type 54 authentication mode of [RFC5304]. In this case, the Authentication TLV (TLV 10) is needed.

#### 3.4.2. Wide Metric TLV

IS-IS auto-configuration routers MUST support TLVs using wide metric as defined in [RFC5305]).

It is recommended that IS-IS auto-configuration routers use a high metric value (e.g. 1000000) as default in order to typically prefer the manually configured adjacencies rather than the auto-configuring ones.

#### 3.4.3. Dynamic Host Name TLV

IS-IS auto-configuration routers MAY advertise their Dynamic Host Names TLV (TLV 137, [RFC5301]). The host names could be provisioned by an IT system, or just use the name of vendor, device type or serial number etc. Note that, the hostname needs to be unique so that it could be useful.

### 3.5. Routing Behavior Considerations

#### 3.5.1. Adjacency Formation

Since ISIS does not require strict hold timers matching to form adjacency, this document does not specify specific hold timers. However, the timers should be within a reasonable range based on current practise in the industry. (For example, the defaults defined in [ISO\_IEC10589] .)

### 4. Security Considerations

In general, auto-configuration is mutually incompatible with authentication. This is a common problem that IS-IS auto-configuration can not avoid.

For wired deployment, the wired line itself could be considered as an implicit authentication that normally unwanted routers are not able to connect to the wire line; for wireless deployment, the authentication could be achieve at the lower wireless link layer.

Malicious router could modify the System ID field to keep causing System ID duplication detection and resolution thus cause the routing system oscillate. However, this is not a new attack vector as without this document the consequences would be higher as other routers would not try to adapt.

### 5. IANA Considerations

The Router-Fingerprint TLV type code needs an assignment by IANA.

### 6. Acknowledgements

This document was heavily inspired by [RFC7503].

Martin Winter, Christian Franke and David Lamparter gave essential feedback to improve the technical design based on their implementation experience.

Many useful comments were made by Acee Lindem, Karsten Thomannby, Hannes Gredler, Peter Lothberg, Uma Chundury, Qin Wu, Sheng Jiang and Nan Wu, etc.

This document was produced using the xml2rfc tool [RFC2629].  
(initially prepared using 2-Word-v2.0.template.dot. )

## 7. References

### 7.1. Normative References

- [ISO\_IEC10589]  
"Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589", November 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<http://www.rfc-editor.org/info/rfc1195>>.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <<http://www.rfc-editor.org/info/rfc2629>>.
- [RFC5301] McPherson, D. and N. Shen, "Dynamic Hostname Exchange Mechanism for IS-IS", RFC 5301, DOI 10.17487/RFC5301, October 2008, <<http://www.rfc-editor.org/info/rfc5301>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<http://www.rfc-editor.org/info/rfc5308>>.
- [RFC6232] Wei, F., Qin, Y., Li, Z., Li, T., and J. Dong, "Purge Originator Identification TLV for IS-IS", RFC 6232, DOI 10.17487/RFC6232, May 2011, <<http://www.rfc-editor.org/info/rfc6232>>.

### 7.2. Informative References

- [I-D.ietf-homenet-hnccp]  
Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", draft-ietf-homenet-hnccp-09 (work in progress), August 2015.

[RFC7503] Lindem, A. and J. Arkko, "OSPFv3 Autoconfiguration",  
RFC 7503, DOI 10.17487/RFC7503, April 2015,  
<<http://www.rfc-editor.org/info/rfc7503>>.

Authors' Addresses

Bing Liu  
Huawei Technologies  
Q14, Huawei Campus, No.156 Beiqing Road  
Hai-Dian District, Beijing, 100095  
P.R. China

Email: [leo.liubing@huawei.com](mailto:leo.liubing@huawei.com)

Bruno Decraene  
Orange  
38 rue du General Leclerc  
Issy-les-Moulineaux FR  
FR

Email: [bruno.decraene@orange.com](mailto:bruno.decraene@orange.com)

Ian Farrer  
Deutsche Telekom AG  
Bonn  
Germany

Email: [ian.farrer@telekom.de](mailto:ian.farrer@telekom.de)

Mikael Abrahamsson  
T-Systems  
Stockholm  
Sweden

Email: [mikael.abrahamsson@t-systems.se](mailto:mikael.abrahamsson@t-systems.se)

Les Ginsberg  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas CA 95035  
USA

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

ISIS Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: August 17, 2016

J. You  
Q. Liang  
Huawei Technologies  
K. Patel  
Cisco Systems  
P. Fan

Z. Li  
China Mobile  
February 14, 2016

ISIS-ISIS Extensions for Flow Specification  
draft-you-isis-flowspec-extensions-04

Abstract

Dissemination of the Traffic flow information was first introduced in the BGP protocol [RFC5575]. FlowSpec rules are used to distribute traffic filtering rules that are used to filter Denial-of-Service (DoS) attacks. For the networks that only deploy IS-IS or IS-IS variant, it is required that IS-IS is extended to distribute Flow Specification or FlowSpec rules.

This document discusses the use cases for distributing flow specification (FlowSpec) routes using IS-IS. Furthermore, this document defines a new IS-IS FlowSpec Reachability TLV encoding format that can be used to distribute FlowSpec rules, its validation procedures for imposing the filtering information on the routers, and a capability to indicate the support of FlowSpec functionality.

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 [RFC2119].

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 <http://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 August 17, 2016.

#### Copyright Notice

Copyright (c) 2016 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 (<http://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 . . . . .                                    | 3  |
| 2. Terminology . . . . .                                     | 3  |
| 3. Use Cases for IS-IS based FlowSpec Distribution . . . . . | 3  |
| 3.1. Anti-DDOS . . . . .                                     | 3  |
| 4. IS-IS Extensions for FlowSpec Rules . . . . .             | 4  |
| 4.1. FlowSpec Filters sub-TLV . . . . .                      | 5  |
| 4.1.1. Order of Traffic Filtering Rules . . . . .            | 7  |
| 4.1.2. Validation Procedure . . . . .                        | 7  |
| 4.2. FlowSpec Action sub-TLV . . . . .                       | 8  |
| 4.2.1. Traffic-rate . . . . .                                | 9  |
| 4.2.2. Traffic-action . . . . .                              | 9  |
| 4.2.3. Traffic-marking . . . . .                             | 9  |
| 4.2.4. Redirect-to-IP . . . . .                              | 10 |
| 5. Redistribution of FlowSpec Rules . . . . .                | 10 |
| 6. IANA Considerations . . . . .                             | 11 |
| 6.1. FlowSpec Reachability TLV . . . . .                     | 11 |
| 6.2. FlowSpec Filters sub-TLVs . . . . .                     | 11 |
| 6.3. FlowSpec Filter Component Types . . . . .               | 11 |
| 6.4. FlowSpec Action sub-TLVs . . . . .                      | 12 |
| 7. Security Considerations . . . . .                         | 13 |
| 8. Acknowledgement . . . . .                                 | 13 |
| 9. References . . . . .                                      | 13 |
| 9.1. Normative References . . . . .                          | 13 |
| 9.2. Informative References . . . . .                        | 13 |

|                              |    |
|------------------------------|----|
| Authors' Addresses . . . . . | 14 |
|------------------------------|----|

## 1. Introduction

[RFC5575] defines Border Gateway Protocol protocol extensions that can be used to distribute traffic flow specifications. One application of this encoding format is to automate inter-domain coordination of traffic filtering, such as what is required in order to mitigate (distributed) denial-of-service attacks.

For the networks deploying only IS-IS or IS-IS variant, it is expected to extend IS-IS to distribute FlowSpec rules. This document discusses the use cases for distributing FlowSpec rules using IS-IS. Furthermore, this document also defines a new IS-IS FlowSpec Reachability TLV encoding format that can be used to distribute FlowSpec entries to the specific routers in the campus network, its validation procedures for imposing the filtering information on the routers, and a capability to indicate the support of FlowSpec functionality.

The semantic content of the FlowSpec extensions defined in this document are identical to the corresponding extensions to BGP ([RFC5575] and [I-D.ietf-idr-flow-spec-v6]). In order to avoid repetition, this document only concentrates on those parts of specification where IS-IS is different from BGP. The IS-IS FlowSpec extensions defined in this document can be used to mitigate the impacts of DoS attacks.

## 2. Terminology

This section contains definitions for terms used frequently throughout this document. However, many additional definitions can be found in [ISO-10589] and [RFC5575].

**Flow Specification (FlowSpec):** A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. Each FlowSpec consists of a set of filters and a set of actions.

## 3. Use Cases for IS-IS based FlowSpec Distribution

### 3.1. Anti-DDOS

For the networks using IS-IS or IS-IS variant, for example, the campus network or DC network, it is expected to extend IS-IS to distribute FlowSpec rules as shown in Figure 1. In this network, the traffic analyzer could be deployed to inject the FlowSpec rules into Router A. Router A creates FlowSpec entries according to the

FlowSpec rules, then the FlowSpec entries would be distributed to the other routers in this domain using IS-IS. Consequently, the attack traffic could be blocked or the suspicious traffic could be limited to a low rate as early as possible.

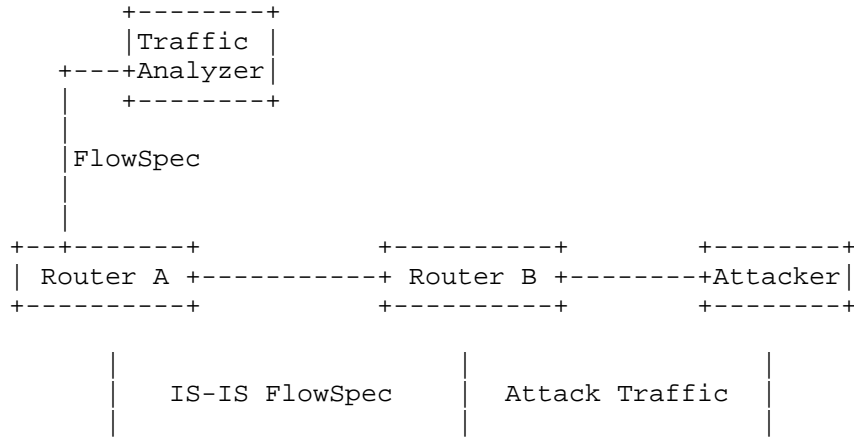


Figure 1: Anti-DDOS in IS-IS Network

#### 4. IS-IS Extensions for FlowSpec Rules

This document defines a new IS-IS TLV, i.e. the FlowSpec reachability TLV (TLV type: TBD1), to describe the FlowSpec rules. An LSP (Link State Protocol) Data Unit [ISO-10589] can carry one or more FlowSpec reachability TLVs.

Each FlowSpec Reachability TLV carries a FlowSpec entry. The FlowSpec entry consists of a FlowSpec Filters sub-TLV and one or more corresponding FlowSpec Action sub-TLVs.

The FlowSpec Reachability TLV is defined below in Figure 2:





Figure 2: FlowSpec Reachability TLV

Type: 1 octet. Type code is TBD1.

Length: 1 octet. The length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of 0).

Value: variable. The value field contains a "Flags" field and a FlowSpec entry, which consists of a FlowSpec filters sub-TLV and one or more corresponding FlowSpec action sub-TLVs. The size of the FlowSpec entry cannot be greater than 253. In most scenarios, using one FlowSpec entry is sufficient. If the injected FlowSpec rule is too complex that the IS-IS router has to use more than 253 octets to encode it into a FlowSpec entry, the IS-IS router should reject it. It is strongly recommended that the FlowSpec rule provider should split or revise the complex FlowSpec rule to a suitable one for the IS-IS routers.

Flags: One octet Field identifying Flags



The least significant bit L is defined as a Leaking enable bit. If set, the FlowSpec Reachability TLV SHOULD be flooded across the entire routing domain. If the L flag is not set, the FlowSpec Reachability TLV MUST NOT be leaked between levels. This bit MUST NOT be altered during the TLV leaking. This Flags may be modified by the IS-IS Speaker according to a local policy.

4.1. FlowSpec Filters sub-TLV

IS-IS FlowSpec filters sub-TLV is one component of FlowSpec entry, carried in the FlowSpec reachability TLV. It is defined below in Figure 3.





Figure 3: IS-IS FlowSpec Filters sub-TLV

Type: the TLV type (Type Code: TBD2 for IPv4 FlowSpec filters, TBD3 for IPv6 FlowSpec filters)

Length: the size of the value field in octets, it cannot be greater than 253.

Flags: One octet Field identifying Flags



The least significant bit S is defined as a strict filter check bit. If set, strict validation rules outlined in the validation section Section 4.1.2 need to be enforced.

Filters: the same as "flow-spec filter components" defined in [RFC5575] and [I-D.ietf-idr-flow-spec-v6].

Table 1: IS-IS Supported FlowSpec Filter Component Types

| Type | Description  | RFC/ WG draft                        |
|------|--|--------------------------------------|
| 1    | Destination IPv4 Prefix<br>Destination IPv6 Prefix | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 2    | Source IPv4 Prefix<br>Source IPv6 Prefix           | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 3    | IP Protocol<br>Next Header                         | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 4    | Port   | RFC5575                              |
| 5    | Destination port                                   | RFC5575                              |
| 6    | Source port  | RFC5575                              |
| 7    | ICMP type  | RFC5575                              |
| 8    | ICMP code  | RFC5575                              |
| 9    | TCP flags  | RFC5575                              |
| 10   | Packet length                                      | RFC5575                              |
| 11   | DSCP   | RFC5575                              |
| 12   | Fragment   | RFC5575                              |
| 13   | Flow Label   | I-D.ietf-idr-flow-spec-v6            |

#### 4.1.1.1. Order of Traffic Filtering Rules

With traffic filtering rules, more than one rule may match a particular traffic flow. The order of applying the traffic filter rules is the same as described in Section 5.1 of [RFC5575] and in Section 3.1 of [I-D.ietf-idr-flow-spec-v6].

#### 4.1.1.2. Validation Procedure

[RFC5575] defines a validation procedure for BGP FlowSpec rules, and [I-D.ietf-idr-bgp-flowspec-oid] describes a modification to the validation procedure defined in [RFC5575] for the dissemination of BGP flow specifications. The IS-IS FlowSpec should support similar features to mitigate the unnecessary or invalid application of

traffic filter rules. The IS-IS FlowSpec validation procedure is described as follows.

When a router receives a FlowSpec rule including a destination prefix filter from its neighbor router, it should consider the prefix filter as a valid filter unless the S bit in the flags field of Filter TLV is set. If the S bit is set, then the FlowSpec rule is considered valid if and only if:

The originator of the FlowSpec rule matches the originator of the best-match unicast route for the destination prefix embedded in the FlowSpec.

The former rule allows any centralized controller to originate the prefix filter and advertise it within a given IS-IS network. The latter rule, also known as a Strict Validation rule, allows strict checking and enforces that the originator of the FlowSpec filter is also the originator of the destination prefix.

When multiple equal-cost paths exist in the routing table entry, each path could end up having a separate set of FlowSpec rules.

When a router receives a FlowSpec rule not including a destination prefix filter from its neighbor router, the validation procedure described above is not applicable.

The FlowSpec filter validation state is used by an IS-IS speaker when the filter is considered for an installation in its FIB. An IS-IS speaker MUST flood IS-IS LSP containing a FlowSpec Reachability TLV as per the entries defined in [ISO-10589] regardless of the validation state of the prefix filters.

#### 4.2. FlowSpec Action sub-TLV

There are one or more FlowSpec Action TLVs associated with a FlowSpec Filters TLV. Different FlowSpec Filters TLV could have the same FlowSpec Action TLVs. The following IS-IS FlowSpec action TLVs, except Redirect, are same as defined in [RFC5575].


Redirect: IPv4 or IPv6 address. This target IP address MUST correspond to a tunnel in the current IS-IS router, if not, the "redirect to IP" action is invalid, and if the flowspec entry has no other action, the flowspec entry is invalid and wouldn't be installed. If the IS-IS router doesn't have a valid route for the target IP, the "redirect to IP" action is also invalid.

Table 2: BGP FlowSpec Actions

| type   | FlowSpec Action  | RFC/WG draft                          |
|--------|------------------|---------------------------------------|
| 0x8006 | traffic-rate     | RFC5575                               |
| 0x8007 | traffic-action   | RFC5575                               |
| 0x8108 | redirect-to-IPv4 | I-D.ietf-idr-flowspec-redirect-rt-bis |
| 0x800b | redirect-to-IPv6 | I-D.ietf-idr-flow-spec-v6             |
| 0x8009 | traffic-marking  | RFC5575                               |

4.2.1. Traffic-rate


Traffic-rate TLV is encoded as:



Traffic-rate: the same as defined in [RFC5575].

4.2.2. Traffic-action


Traffic-action TLV is encoded as:



S flag and T flag: the same as defined in [RFC5575].

4.2.3. Traffic-marking


Traffic-marking TLV is encoded as:




DSCP value: the same as defined in [RFC5575].

4.2.4. Redirect-to-IP

Redirect-to-IPv4 is encoded as:



Redirect to IPv6 TLV is encoded as:



IPv4/6 Address: the redirection target IP address.

'C' (or copy) bit: when the 'C' bit is set, the redirection applies to copies of the matching packets and not to the original traffic stream [I-D.ietf-idr-flowspec-redirect-ip].

5. Redistribution of FlowSpec Rules

An implementation MAY provide an option for an IS-IS speaker to announce a redistributed FlowSpec route within an IS-IS domain regardless of being installed in its local FIB. An implementation MAY impose an upper bound on number of FlowSpec entries that an IS-IS router MAY advertise.

## 6. IANA Considerations

This document defines the following new IS-IS TLV types, which need to be reflected in the IS-IS TLV codepoint registry.

### 6.1. FlowSpec Reachability TLV

| Type | Description                   | IIH | LSP | SNP |
|------|-------------------------------|-----|-----|-----|
| TBD1 | The FlowSpec Reachability TLV | n   | y   | n   |

### 6.2. FlowSpec Filters sub-TLVs

| Type | Description                   |
|------|-------------------------------|
| TBD2 | IPv4 FlowSpec filters sub-TLV |
| TBD3 | IPv6 FlowSpec filters sub-TLV |

### 6.3. FlowSpec Filter Component Types

| Type | Description  | RFC/ WG draft                        |
|------|--|--------------------------------------|
| 1    | Destination IPv4 Prefix<br>Destination IPv6 Prefix | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 2    | Source IPv4 Prefix<br>Source IPv6 Prefix           | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 3    | IP Protocol<br>Next Header                         | RFC5575<br>I-D.ietf-idr-flow-spec-v6 |
| 4    | Port   | RFC5575                              |
| 5    | Destination port                                   | RFC5575                              |
| 6    | Source port  | RFC5575                              |
| 7    | ICMP type  | RFC5575                              |
| 8    | ICMP code  | RFC5575                              |
| 9    | TCP flags  | RFC5575                              |
| 10   | Packet length                                      | RFC5575                              |
| 11   | DSCP   | RFC5575                              |
| 12   | Fragment   | RFC5575                              |
| 13   | Flow Label   | I-D.ietf-idr-flow-spec-v6            |

#### 6.4. FlowSpec Action sub-TLVs

This document defines a group of FlowSpec actions. The following TLV types need to be assigned:

Type TBD4 - traffic-rate

Type TBD5 - traffic-action

Type TBD6 - traffic-marking

Type TBD7 - redirect to IPv4

Type TBD8 - redirect to IPv6

## 7. Security Considerations

This extension to IS-IS does not change the underlying security issues inherent in the existing IS-IS. Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard IS-IS failures.

## 8. Acknowledgement

The authors would like to thank Jeff Haas for his useful comments.

## 9. References

### 9.1. Normative References

[ISO-10589]

ISO, "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", International Standard 10589: 2002, Second Edition, 2002.

[RFC2119]

Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC4360]

Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.

[RFC5575]

Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<http://www.rfc-editor.org/info/rfc5575>>.

### 9.2. Informative References

[I-D.ietf-idr-bgp-flowspec-oid]

Uttaro, J., Filsfils, C., Smith, D., Alcaide, J., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", draft-ietf-idr-bgp-flowspec-oid-02 (work in progress), January 2014.

## [I-D.ietf-idr-flow-spec-v6]

Raszuk, R., Pithawala, B., McPherson, D., and A. Andy,  
"Dissemination of Flow Specification Rules for IPv6",  
draft-ietf-idr-flow-spec-v6-06 (work in progress),  
November 2014.

## [I-D.ietf-idr-flowspec-redirect-ip]

Uttaro, J., Haas, J., Texier, M., Andy, A., Ray, S.,  
Simpson, A., and W. Henderickx, "BGP Flow-Spec Redirect to  
IP Action", draft-ietf-idr-flowspec-redirect-ip-02 (work  
in progress), February 2015.

## [I-D.ietf-idr-flowspec-redirect-rt-bis]

Haas, J., "Clarification of the Flowspec Redirect Extended  
Community", draft-ietf-idr-flowspec-redirect-rt-bis-05  
(work in progress), July 2015.

## Authors' Addresses

Jianjie You  
Huawei Technologies  
101 Software Avenue, Yuhuatai District  
Nanjing 210012  
China

Email: youjianjie@huawei.com

Qiandeng Liang  
Huawei Technologies  
101 Software Avenue, Yuhuatai District  
Nanjing 210012  
China

Email: liangqiandeng@huawei.com

Keyur Patel  
Cisco Systems  
170 W. Tasman Drive  
San Jose CA 95124 95134  
US

Email: keyupate@cisco.com

Peng Fan  
Beijing  
China

Email: peng.fan@139.com

Zhenqiang Li  
China Mobile  
Beijing  
China

Email: li\_zhenqiang@hotmail.com