Workgroup: Internet Engineering Task Force
Internet-Draft:
draft-ietf-lsr-isis-rfc5316bis-07
Obsoletes: 5316 (if approved)
Published: 28 September 2022
Intended Status: Standards Track
Expires: 1 April 2023
Authors: M. Chen L. Ginsberg S. Previdi
Huawei Cisco Systems Huawei Technologies
D. Xiaodong
China Mobile

IS-IS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering

Abstract

This document describes extensions to the Intermediate System to Intermediate System (IS-IS) protocol to support Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) for multiple Autonomous Systems (ASs). It defines IS-IS extensions for the flooding of TE information about inter-AS links, which can be used to perform inter-AS TE path computation.

No support for flooding information from within one AS to another AS is proposed or defined in this document.

This document builds on RFC 5316 by adding support for IPv6-only operation.

This document obsoletes RFC 5316.

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

Copyright Notice

Table of Contents

1. Introduction

[RFC5305] defines extensions to the IS-IS protocol [RFC1195] to support intra-area Traffic Engineering (TE). The extensions provide a way of encoding the TE information for TE-enabled links within the network (TE links) and flooding this information within an area. The extended IS reachability TLV and traffic engineering router ID TLV, which are defined in [RFC5305], are used to carry such TE information. The extended IS reachability TLV has several nested sub-TLVs that describe the TE attributes for a TE link.

[RFC6119] and [RFC5307] define similar extensions to IS-IS in support of IPv6 and Generalized Multiprotocol Label Switching (GMPLS) TE respectively.

Requirements for establishing Multiprotocol Label Switching (MPLS) TE Label Switched Paths (LSPs) that cross multiple Autonomous Systems (ASes) are described in [RFC4216]. As described in [RFC4216], a method SHOULD provide the ability to compute a path spanning multiple ASes. So a path computation entity that may be the head-end Label Switching Router (LSR), an AS Border Router (ASBR), or a Path Computation Element (PCE) [RFC4655] needs to know the TE information not only of the links within an AS, but also of the links that connect to other ASes.

In this document, a new TLV, which is referred to as the inter-AS reachability TLV, is defined to advertise inter-AS TE information, and three new sub-TLVs are defined for inclusion in the inter-AS reachability TLV to carry the information about the remote AS number and remote ASBR ID. The sub-TLVs defined in [RFC5305][RFC6119] and other documents for inclusion in the extended IS reachability TLV for describing the TE properties of a TE link are applicable to be included in the Inter-AS Reachability TLV for describing the TE properties of an inter-AS TE link as well. Also, two more new sub-TLVs are defined for inclusion in the IS-IS router capability TLV to carry the TE Router ID when the TE Router ID is needed to reach all routers within an entire IS-IS routing domain. The extensions are equally applicable to IPv4 and IPv6 as identical extensions to [RFC5305] and [RFC6119]. Detailed definitions and procedures are discussed in the following sections.

This document does not propose or define any mechanisms to advertise any other extra-AS TE information within IS-IS. See Section 2.1 for a full list of non-objectives for this work.

2. Problem Statement

As described in [RFC4216], in the case of establishing an inter-AS TE LSP that traverses multiple ASes, the Path message [RFC3209] may include the following elements in the Explicit Route Object (ERO) in order to describe the path of the LSP:

* a set of AS numbers as loose hops; and/or

* a set of LSRs including ASBRs as loose hops.

Two methods for determining inter-AS paths have been described elsewhere. The per-domain method [RFC5152] determines the path one domain at a time. The backward recursive method [RFC5441] uses cooperation between PCEs to determine an optimum inter-domain path. The sections that follow examine how inter-AS TE link information could be useful in both cases.

2.1. A Note on Non-Objectives

It is important to note that this document does not make any change to the confidentiality and scaling assumptions surrounding the use of ASes in the Internet. In particular, this document is conformant to the requirements set out in [RFC4216].

The following features are explicitly excluded:

* There is no attempt to distribute TE information from within one AS to another AS.

* There is no mechanism proposed to distribute any form of TE reachability information for destinations outside the AS.

* There is no proposed change to the PCE architecture or usage.

* TE aggregation is not supported or recommended.

* There is no exchange of private information between ASes.

* No IS-IS adjacencies are formed on the inter-AS link.

2.2. Per-Domain Path Determination

In the per-domain method of determining an inter-AS path for an MPLS-TE LSP, when an LSR that is an entry-point to an AS receives a Path message from an upstream AS with an ERO containing a next hop that is an AS number, it needs to find which LSRs (ASBRs) within the local AS are connected to the downstream AS. That way, it can compute a TE LSP segment across the local AS to one of those LSRs and forward the Path message to that LSR and hence into the next AS. See Figure 1 for an example.

```
R1------R3----R5-----R7------R9-----R11
        |     | \    |      / |
        |     |  \   |  ----  |
        |     |   \  | /      |
R2------R4----R6   --R8------R10----R12
           :              :
<-- AS1 -->:<---- AS2 --->:<--- AS3 --->
```

Figure 1: Inter-AS Reference Model

The figure shows three ASes (AS1, AS2, and AS3) and twelve LSRs (R1 through R12). R3 and R4 are ASBRs in AS1. R5, R6, R7, and R8 are ASBRs in AS2. R9 and R10 are ASBRs in AS3.

If an inter-AS TE LSP is planned to be established from R1 to R12, the AS sequence will be: AS1, AS2, AS3.

Suppose that the Path message enters AS2 from R3. The next hop in the ERO shows AS3, and R5 must determine a path segment across AS2 to reach AS3. It has a choice of three exit points from AS2 (R6, R7, and R8), and it needs to know which of these provide TE connectivity to AS3, and whether the TE connectivity (for example, available bandwidth) is adequate for the requested LSP.

Alternatively, if the next hop in the ERO is an entry ASBR for AS3 (say R9), R5 needs to know which of its exit ASBRs has a TE link that connects to R9. Since there may be multiple ASBRs that are connected to R9 (both R7 and R8 in this example), R5 also needs to know the TE properties of the inter-AS TE links so that it can select the correct exit ASBR.

Once the Path message reaches the exit ASBR, any choice of inter-AS TE link can be made by the ASBR if not already made by the entry ASBR that computed the segment.

More details can be found in Section 4 of [RFC5152], which clearly points out why advertising of inter-AS links is desired.

To enable R5 to make the correct choice of exit ASBR, the following information is needed:

* List of all inter-AS TE links for the local AS.

* TE properties of each inter-AS TE link.

* AS number of the neighboring AS connected to by each inter-AS TE link.

* Identity (TE Router ID) of the neighboring ASBR connected to by each inter-AS TE link.

In GMPLS networks, further information may also be required to select the correct TE links as defined in [RFC5307].

The example above shows how this information is needed at the entry-point ASBRs for each AS (or the PCEs that provide computation services for the ASBRs). However, this information is also needed throughout the local AS if path computation functionality is fully distributed among LSRs in the local AS, for example to support LSPs that have start points (ingress nodes) within the AS.

2.3. Backward Recursive Path Computation

Another scenario using PCE techniques has the same problem. [RFC5441] defines a PCE-based TE LSP computation method (called Backward Recursive Path Computation) to compute optimal inter-domain constrained MPLS-TE or GMPLS LSPs. In this path computation method, a specific set of traversed domains (ASes) are assumed to be selected before computation starts. Each downstream PCE in domain(i) returns to its upstream neighbor PCE in domain(i-1) a multipoint-to-point tree of potential paths. Each tree consists of the set of paths from all boundary nodes located in domain(i) to the destination where each path satisfies the set of required constraints for the TE LSP (bandwidth, affinities, etc.).

So a PCE needs to select boundary nodes (that is, ASBRs) that provide connectivity from the upstream AS. In order for the tree of paths provided by one PCE to its neighbor to be correlated, the identities of the ASBRs for each path need to be referenced. Thus, the PCE must know the identities of the ASBRs in the remote AS that are reached by any inter-AS TE link, and, in order to provide only suitable paths in the tree, the PCE must know the TE properties of the inter-AS TE links. See the following figure as an example.

```
             PCE1<------>PCE2<-------->PCE3
             /       :             :
            /        :             :
          R1------R3----R5-----R7------R9-----R11
                  |     | \    |      / |
                  |     |  \   |  ----  |
                  |     |   \  | /      |
          R2------R4----R6   --R8------R10----R12
                     :              :
          <-- AS1 -->:<---- AS2 --->:<--- AS3 --->
```

Figure 2: BRPC for Inter-AS Reference Model

The figure shows three ASes (AS1, AS2, and AS3), three PCEs (PCE1, PCE2, and PCE3), and twelve LSRs (R1 through R12). R3 and R4 are ASBRs in AS1. R5, R6, R7, and R8 are ASBRs in AS2. R9 and R10 are ASBRs in AS3. PCE1, PCE2, and PCE3 cooperate to perform inter-AS path computation and are responsible for path segment computation within their own domain(s).

If an inter-AS TE LSP is planned to be established from R1 to R12, the traversed domains are assumed to be selected: AS1->AS2->AS3, and the PCE chain is: PCE1->PCE2->PCE3. First, the path computation request originated from the PCC (R1) is relayed by PCE1 and PCE2 along the PCE chain to PCE3. Then, PCE3 begins to compute the path segments from the entry boundary nodes that provide connection from AS2 to the destination (R12). But, to provide suitable path segments, PCE3 must determine which entry boundary nodes provide connectivity to its upstream neighbor AS (identified by its AS number), and must know the TE properties of the inter-AS TE links. In the same way, PCE2 also needs to determine the entry boundary nodes according to its upstream neighbor AS and the inter-AS TE link capabilities.

Thus, to support Backward Recursive Path Computation, the same information listed in Section 2.2 is required. The AS number of the neighboring AS connected to by each inter-AS TE link is particularly important.

3. Extensions to ISIS-TE

Note that this document does not define mechanisms for distribution of TE information from one AS to another, does not distribute any form of TE reachability information for destinations outside the AS, does not change the PCE architecture or usage, does not suggest or recommend any form of TE aggregation, and does not feed private information between ASes. See Section 2.1.

In this document, for the advertisement of inter-AS TE links, a new TLV, which is referred to as the inter-AS reachability TLV, is defined. Three new sub-TLVs are also defined for inclusion in the inter-AS reachability TLV to carry the information about the neighboring AS number and the remote ASBR ID of an inter-AS link. The sub-TLVs defined in [RFC5305], [RFC6119], and other documents for inclusion in the extended IS reachability TLV are applicable to be included in the inter-AS reachability TLV for inter-AS TE links advertisement.

This document also defines two new sub-TLVs for inclusion in the IS-IS router capability TLV to carry the TE Router ID when the TE Router ID is needed to reach all routers within an entire IS-IS routing domain.

While some of the TE information of an inter-AS TE link may be available within the AS from other protocols, in order to avoid any dependency on where such protocols are processed, this mechanism carries all the information needed for the required TE operations.

3.1. Choosing the TE Router ID Value

Subsequent sections specify advertisement of a TE Router ID value for IPv4 and/or IPv6. This section defines how this value is chosen.

A TE Router ID MUST be an address which is unique within the IS-IS domain and stable i.e., it can always be referenced in a path that will be reachable from multiple hops away, regardless of the state of the node's interfaces.

When advertising an IPv4 address as a TE Router ID, if the Traffic Engineering Router ID TLV [RFC5305] is being advertised, then the address SHOULD be identical to the address in the Traffic Engineering Router ID TLV. The TE Router ID MAY be identical to an IP Interface Address [RFC1195] advertised by the originating IS so long as the address meets the requirements specified above.

When advertising an IPv6 address as a TE Router ID, if the IPv6 TE Router ID TLV [RFC6119] is being advertised, then the address SHOULD be identical to the address in the IPv6 TE Router ID TLV. The TE Router ID MAY be identical to a non-link-local IPv6 Interface Address advertised by the originating IS in a Link State PDU using the IPv6 Intf. Addr TLV [RFC5308] so long as the address meets the requirements specified above.

3.2. Inter-AS Reachability TLV

The inter-AS reachability TLV has type 141 (see Section 6.1) and contains a data structure consisting of:

```
 0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Router ID                                     (4 octets)    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   default metric                              | (3 octets)
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Flags     |                                 (1 octet)
+-+-+-+-+-+-+-+-+
|sub-TLVs length|                                 (1 octet)
+-+-+-+-+-+-+-+-+-+-+-+-+-
| sub-TLVs ...                                    (0-246 octets)
+-+-+-+-+-+-+-+-+-+-+-+-+-
```

```
Flags consists of the following:
    0 1 2 3 4 5 6 7
   +-+-+-+-+-+-+-+-+
   |S|D| Rsvd      |
   +-+-+-+-+-+-+-+-+
```

where:

S bit: If the S bit is set(1), the Inter-AS Reachability 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: When the Inter-AS Reachability TLV is leaked from Level 2 (L2) to Level 1 (L1), the D bit MUST be set. Otherwise, this bit MUST be clear. Inter-AS Reachability TLVs with the D bit set MUST NOT be leaked from Level 1 to Level 2. This is to prevent TLV looping.

Reserved(Rsvd) bits MUST be zero when originated and ignored when received.

Compared to the extended reachability TLV which is defined in [RFC5305], the inter-AS reachability TLV replaces the "7 octets of System ID and Pseudonode Number" field with a "4 octets of Router ID" field and introduces an extra "control information" field, which consists of a flooding-scope bit (S bit), an up/down bit (D bit), and 6 reserved bits.

The Router ID field of the inter-AS reachability TLV is 4 octets in length and has a value as defined in Section 3.1. 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 Router ID sub-TLV MUST be present in the inter-AS reachability TLV. The Router ID could be used to indicate the source of the inter-AS reachability TLV.

The flooding procedures for inter-AS reachability TLV are identical to the flooding procedures for the GENINFO TLV, which are defined in Section 4 of [RFC6823]. These procedures have been previously discussed in [RFC7981]. The flooding-scope bit (S bit) SHOULD be set to 0 if the flooding scope is to be limited to within the single IGP area to which the ASBR belongs. It MAY be set to 1 if the information is intended to reach all routers (including area border routers, ASBRs, and PCEs) in the entire IS-IS routing domain. The choice between the use of 0 or 1 is an AS-wide policy choice, and configuration control SHOULD be provided in ASBR implementations that support the advertisement of inter-AS TE links.

The sub-TLVs defined in [RFC5305], [RFC6119], and other documents for describing the TE properties of a TE link are also applicable to the inter-AS reachability TLV for describing the TE properties of an Inter-AS TE link. Apart from these sub-TLVs, four new sub-TLVs are defined for inclusion in the inter-AS reachability TLV defined in this document:

Sub-TLV type	Length	Name
24	4	remote AS number
25	4	IPv4 remote ASBR identifier
26	16	IPv6 remote ASBR identifier
TBD1	16	IPv6 local ASBR identifier

Detailed definitions of the four new sub-TLVs are described in Sections 3.3.1, 3.3.2, 3.3.3, and 3.3.4.

3.3. TE Router ID

The Traffic Engineering router ID TLV and IPv6 TE Router ID TLV, which are defined in [RFC5305] and [RFC6119] respectively, only have area flooding-scope. When performing inter-AS TE, the TE Router ID MAY be needed to reach all routers within an entire IS-IS routing domain and it MUST have the same flooding scope as the Inter-AS Reachability TLV does.

[RFC7981] defines a generic advertisement mechanism for IS-IS which allows a router to advertise its capabilities within an IS-IS area or an entire IS-IS routing domain. [RFC7981] also points out that the TE Router ID is a candidate to be carried in the IS-IS router capability TLV when performing inter-area TE.

This document uses such mechanism for TE Router ID advertisement when the TE Router ID is needed to reach all routers within an entire IS-IS Routing domain. Two new sub-TLVs are defined for inclusion in the IS-IS Router Capability TLV to carry the TE Router IDs.

```
Sub-TLV type    Length  Name
------------    ------  -----------------
          11        4   IPv4 TE Router ID
          12       16   IPv6 TE Router ID
```

Detailed definitions of the new sub-TLVs are described in Section 3.4.1 and 3.4.2.

3.4. Sub-TLVs for Inter-AS Reachability TLV

3.4.1. Remote AS Number Sub-TLV

A new sub-TLV, the remote AS number sub-TLV, is defined for inclusion in the inter-AS reachability TLV when advertising inter-AS links. The remote AS number sub-TLV specifies the AS number of the neighboring AS to which the advertised link connects.

The remote AS number sub-TLV is TLV type 24 (see Section 6.2) and is 4 octets in length. The format is as follows:

```
 0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote AS Number                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The remote AS number field has 4 octets. When only 2 octets are used for the AS number, the left (high-order) 2 octets MUST be set to 0. The remote AS number sub-TLV MUST be included when a router advertises an inter-AS TE link.

3.4.2. IPv4 Remote ASBR ID Sub-TLV

A new sub-TLV, which is referred to as the IPv4 remote ASBR ID sub-TLV, is defined for inclusion in the inter-AS reachability TLV when advertising inter-AS links. The IPv4 remote ASBR ID sub-TLV specifies the IPv4 identifier of the remote ASBR to which the advertised inter-AS link connects. The value advertised is selected as defined in Section 3.1.

The IPv4 remote ASBR ID sub-TLV is TLV type 25 (see Section 6.2) and is 4 octets in length. The format of the IPv4 remote ASBR ID sub-TLV is as follows:

```
 0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote ASBR ID                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The IPv4 remote ASBR ID sub-TLV MUST be included if the neighboring ASBR has an IPv4 address. The value advertised is selected as defined in Section 3.1. If the neighboring ASBR does not have an IPv4 address, the IPv6 remote ASBR ID sub-TLV MUST be included instead. An IPv4 remote ASBR ID sub-TLV and IPv6 remote ASBR ID sub-TLV MAY both be present in an extended IS reachability TLV.

3.4.3. IPv6 Remote ASBR ID Sub-TLV

A new sub-TLV, which is referred to as the IPv6 remote ASBR ID sub-TLV, is defined for inclusion in the inter-AS reachability TLV when advertising inter-AS links. The IPv6 remote ASBR ID sub-TLV specifies the IPv6 identifier of the remote ASBR to which the advertised inter-AS link connects. The value advertised is selected as defined in Section 3.1.

The IPv6 remote ASBR ID sub-TLV is TLV type 26 (see Section 6.2) and is 16 octets in length. The format of the IPv6 remote ASBR ID sub-TLV is as follows:

```
 0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote ASBR ID                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote ASBR ID (continued)             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote ASBR ID (continued)             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Remote ASBR ID (continued)             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The IPv6 remote ASBR ID sub-TLV MUST be included if the neighboring ASBR has an IPv6 address. If the neighboring ASBR does not have an IPv6 address, the IPv4 remote ASBR ID sub-TLV MUST be included instead. An IPv4 remote ASBR ID sub-TLV and IPv6 remote ASBR ID sub-TLV MAY both be present in an extended IS reachability TLV.

3.4.4. IPv6 Local ASBR ID sub-TLV

The IPv6 Local ASBR ID sub-TLV is TLV type TBD1 (see Section 6.3) and is 16 octets in length. The format of the IPv6 Local ASBR ID sub-TLV is as follows:

```
 0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |              Length           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Local ASBR ID                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Local ASBR ID (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Local ASBR ID (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Local ASBR ID (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The value advertised is selected as defined in Section 3.1.

If the originating node does not support IPv4, the IPv6 Local ASBR ID sub-TLV MUST be present in the inter-AS reachability TLV. Inter-AS reachability TLVs which have a Router ID of 0.0.0.0 and do not have the IPv6 Local ASBR ID sub-TLV present MUST be ignored.

3.5. Sub-TLVs for IS-IS Router Capability TLV

3.5.1. IPv4 TE Router ID sub-TLV

The IPv4 TE Router ID sub-TLV is TLV type 11 (see Section 6.3) and is 4 octets in length. The format of the IPv4 TE Router ID sub-TLV is as follows:

```
0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |              Length           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        TE Router ID                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The value advertised is selected as defined in Section 3.1.

When the TE Router ID is needed to reach all routers within an entire IS-IS routing domain, the IS-IS Router capability TLV MUST be included in its LSP. If an ASBR supports Traffic Engineering for IPv4 and if the ASBR has an IPv4 TE Router ID, the IPv4 TE Router ID sub-TLV MUST be included. If the ASBR does not have an IPv4 TE Router ID, the IPv6 TE Router sub-TLV MUST be included instead. An IPv4 TE Router ID sub-TLV and IPv6 TE Router ID sub-TLV MAY both be present in an IS-IS router capability TLV.

3.5.2. IPv6 TE Router ID sub-TLV

The IPv6 TE Router ID sub-TLV is TLV type 12 (see Section 6.3) and is 16 octets in length. The format of the IPv6 TE Router ID sub-TLV is as follows:

```
0                   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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       TE Router ID                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       TE Router ID   (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       TE Router ID   (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       TE Router ID   (continued)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The value advertised is selected as defined in Section 3.1.

When the TE Router ID is needed to reach all routers within an entire IS-IS routing domain, the IS-IS router capability TLV MUST be included in its LSP. If an ASBR supports Traffic Engineering for IPv6 and if the ASBR has an IPv6 TE Router ID, the IPv6 TE Router ID sub-TLV MUST be included. If the ASBR does not have an IPv6 TE Router ID, the IPv4 TE Router sub-TLV MUST be included instead. An IPv4 TE Router ID sub-TLV and IPv6 TE Router ID sub-TLV MAY both be present in an IS-IS router capability TLV.

4. Procedure for Inter-AS TE Links

When TE is enabled on an inter-AS link and the link is up, the ASBR SHOULD advertise this link using the normal procedures for [RFC5305]. When either the link is down or TE is disabled on the link, the ASBR SHOULD withdraw the advertisement. When there are changes to the TE parameters for the link (for example, when the available bandwidth changes), the ASBR SHOULD re-advertise the link but MUST take precautions against excessive re-advertisements.

Hellos MUST NOT be exchanged over the inter-AS link, and consequently, an IS-IS adjacency MUST NOT be formed.

The information advertised comes from the ASBR's knowledge of the TE capabilities of the link, the ASBR's knowledge of the current status and usage of the link, and configuration at the ASBR of the remote AS number and remote ASBR TE Router ID.

Legacy routers receiving an advertisement for an inter-AS TE link are able to ignore it because they do not know the new TLV and sub-TLVs that are defined in Section 3 of this document. They will continue to flood the LSP, but will not attempt to use the information received.

In the current operation of ISIS-TE, the LSRs at each end of a TE link emit LSPs describing the link. The databases in the LSRs then have two entries (one locally generated, the other from the peer) that describe the different 'directions' of the link. This enables Constrained Shortest Path First (CSPF) to do a two-way check on the link when performing path computation and eliminate it from consideration unless both directions of the link satisfy the required constraints.

In the case we are considering here (i.e., of a TE link to another AS), there is, by definition, no IGP peering and hence no bidirectional TE link information. In order for the CSPF route computation entity to include the link as a candidate path, we have to find a way to get LSPs describing its (bidirectional) TE properties into the TE database.

This is achieved by the ASBR advertising, internally to its AS, information about both directions of the TE link to the next AS. The ASBR will normally generate an LSP describing its own side of a link; here we have it 'proxy' for the ASBR at the edge of the other AS and generate an additional LSP that describes that device's 'view' of the link.

Only some essential TE information for the link needs to be advertised; i.e., the Interface Address, the remote AS number, and the remote ASBR ID of an inter-AS TE link.

Routers or PCEs that are capable of processing advertisements of inter-AS TE links SHOULD NOT use such links to compute paths that exit an AS to a remote ASBR and then immediately re-enter the AS through another TE link. Such paths would constitute extremely rare occurrences and SHOULD NOT be allowed except as the result of specific policy configurations at the router or PCE computing the path.

4.1. Origin of Proxied TE Information

Section 4 describes how an ASBR advertises TE link information as a proxy for its neighbor ASBR, but does not describe where this information comes from.

Although the source of the information described in Section 4 is outside the scope of this document, it is possible that it will be a configuration requirement at the ASBR, as are other local properties of the TE link. Further, where BGP is used to exchange IP routing information between the ASBRs, a certain amount of additional local configuration about the link and the remote ASBR is likely to be available.

We note further that it is possible, and may be operationally advantageous, to obtain some of the required configuration information from BGP. Whether and how to utilize these possibilities is an implementation matter.

5. Security Considerations

The protocol extensions defined in this document are relatively minor and can be secured within the AS in which they are used by the existing IS-IS security mechanisms (e.g., using the cleartext passwords or Hashed Message Authentication Codes, which are defined in [RFC1195], [RFC5304], and [RFC5310] separately).

There is no exchange of information between ASes, and no change to the IS-IS security relationship between the ASes. In particular, since no IS-IS adjacency is formed on the inter-AS links, there is no requirement for IS-IS security between the ASes.

Some of the information included in these new advertisements (e.g., the remote AS number and the remote ASBR ID) is obtained manually from a neighboring administration as part of a commercial relationship. The source and content of this information should be carefully checked before it is entered as configuration information at the ASBR responsible for advertising the inter-AS TE links.

It is worth noting that in the scenario we are considering, a Border Gateway Protocol (BGP) peering may exist between the two ASBRs and that this could be used to detect inconsistencies in configuration (e.g., the administration that originally supplied the information may provide incorrect information, or some manual mis-configurations or mistakes may be made by the operators). For example, if a different remote AS number is received in a BGP OPEN [RFC4271] from that locally configured to ISIS-TE, as we describe here, then local policy SHOULD be applied to determine whether to alert the operator to a potential mis-configuration or to suppress the IS-IS advertisement of the inter-AS TE link. Advertisement of incorrect information could result in an inter-AS TE LSP that traverses an unintended AS. Note further that if BGP is used to exchange TE information as described in Section 4.1, the inter-AS BGP session SHOULD be secured using mechanisms such as described in [RFC5925] to provide authentication and integrity checks.

For a discussion of general security considerations for IS-IS, see [RFC5304].

6. IANA Considerations

IANA is requested to make the following allocations from registries under its control.

6.1. Inter-AS Reachability TLV

This document defines the following new IS-IS TLV type, described in Section 3.1, which has been registered in the IS-IS TLV codepoint registry:

Type	Description	IIH	LSP	SNP	Purge	Reference
141	inter-AS reachability information	n	y	n	n	[This.I-D]

6.2. Sub-TLVs for the Inter-AS Reachability TLV

This document defines the following new sub-TLV types (described in Sections 3.3.1, 3.3.2, 3.3.3, and, 3.3.4) of top-level TLV 141 (see Section 6.1 above). Three of these sub-TLVs have been registered in the IS-IS Sub-TLVs for TLVs Advertising Neighbor Information registry by [RFC5316]. One additional sub-TLV (IPv6 local ASBR identifier) is introduced by this document and needs to be added to the same registry.

Type	Description	22	23	25	141	222	223	Reference
24	remote AS number	n	n	n	y	n	n	[This.I-D]
25	IPv4 remote ASBR identifier	n	n	n	y	n	n	[This.I-D]
26	IPv6 remote ASBR identifier	n	n	n	y	n	n	[This.I-D]
TBD1	IPv6 local ASBR identifier	n	n	n	y	n	n	[This.I-D]

As described above in Section 3.1, the sub-TLVs which are defined in [RFC5305], [RFC6119] and other documents for describing the TE properties of a TE link are applicable to describe an inter-AS TE link and MAY be included in the inter-AS reachability TLV when adverting inter-AS TE links.

6.3. Sub-TLVs for the IS-IS Router Capability TLV

This document defines the following new sub-TLV types, described in Sections 3.4.1 and 3.4.2, of top-level TLV 242 (which is defined in [RFC7981]) that have been registered in the IS-IS Sub-TLVs for IS-IS Router CAPABILITY TLV registry:

Type	Description	Reference
11	IPv4 TE Router ID	[This.I-D]
12	IPv6 TE Router ID	[This.I-D]

7. Acknowledgements

For the original version of [RFC5316] the authors thanked Adrian Farrel, Jean-Louis Le Roux, Christian Hopps, and Hannes Gredler for their review and comments on this document.

8. References

8.1. 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>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI

10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[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>.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <https://www.rfc-editor.org/info/rfc5925>.

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

[RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <https://www.rfc-editor.org/info/rfc7981>.

[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>.

8.2. Informative References

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC4216] Zhang, R., Ed. and J.-P. Vasseur, Ed., "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, DOI 10.17487/RFC4216, November 2005, <https://www.rfc-editor.org/info/rfc4216>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, DOI 10.17487/RFC5152, February 2008, <https://www.rfc-editor.org/info/rfc5152>.

[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>.

[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, <https://www.rfc-editor.org/info/rfc5307>.

[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>.

[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, <https://www.rfc-editor.org/info/rfc5316>.

[RFC5441] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, DOI 10.17487/RFC5441, April 2009, <https://www.rfc-editor.org/info/rfc5441>.

[RFC6823] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", RFC 6823, DOI 10.17487/RFC6823, December 2012, <https://www.rfc-editor.org/info/rfc6823>.

Appendix A. Changes to RFC 5316

The following is a summary of the substantive changes this document makes to RFC 5316. Some editorial changes were also made.

RFC 5316 only allowed a 32 bit Router ID in the fixed header of TLV 141. This is problematic in an IPv6-only deployment where an IPv4 address may not be available. This document specifies:

1. The Router ID should be identical to the value advertised in the Traffic Engineering Router ID TLV (134) if available.

2. 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.

3. 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 new IPv6 Local ASBR identifier sub-TLV must be present in the TLV.

Authors' Addresses

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Les Ginsberg
Cisco Systems

Email: ginsberg@cisco.com

Stefano Previdi
Huawei Technologies
Italy

Email: stefano@previdi.net

Xiaodong Duan
China Mobile

Email: duanxiaodong@chinamobile.com