

Network Working Group
Request for Comments: 5392
Category: Standards Track

M. Chen
R. Zhang
Huawei Technologies Co., Ltd.
X. Duan
China Mobile
January 2009

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

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 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.

Abstract

This document describes extensions to the OSPF version 2 and 3 protocols to support Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) for multiple Autonomous Systems (ASes). OSPF-TE v2 and v3 extensions are defined for the flooding of TE information about inter-AS links that 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.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	3
2. Problem Statement	3
2.1. A Note on Non-Objectives	4
2.2. Per-Domain Path Determination	4
2.3. Backward Recursive Path Computation	6
3. Extensions to OSPF	7
3.1. LSA Definitions	8
3.1.1. Inter-AS-TE-v2 LSA	8
3.1.2. Inter-AS-TE-v3 LSA	8
3.2. LSA Payload	9
3.2.1. Link TLV	9
3.3. Sub-TLV Details	10
3.3.1. Remote AS Number Sub-TLV	10
3.3.2. IPv4 Remote ASBR ID Sub-TLV	11
3.3.3. IPv6 Remote ASBR ID Sub-TLV	11
4. Procedure for Inter-AS TE Links	12
4.1. Origin of Proxied TE Information	13
5. Security Considerations	14
6. IANA Considerations	14
6.1. Inter-AS TE OSPF LSA	14
6.1.1. Inter-AS-TE-v2 LSA	14
6.1.2. Inter-AS-TE-v3 LSA	14
6.2. OSPF LSA Sub-TLVs Type	15
7. Acknowledgments	15
8. References	15
8.1. Normative References	15
8.2. Informative References	16

1. Introduction

[OSPF-TE] defines extensions to the OSPF protocol [OSPF] 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. Type 10 Opaque Link State Advertisements (LSAs) [RFC5250] are used to carry such TE information. Two top-level Type Length Values (TLVs) are defined in [OSPF-TE]: Router Address TLV and Link TLV. The Link TLV has several nested sub-TLVs that describe the TE attributes for a TE link.

[OSPF-V3-TE] defines similar extensions to OSPFv3 [OSPFV3]. It defines a new LSA, which is referred to as the Intra-Area-TE LSA, to advertise TE information. [OSPF-V3-TE] uses "Traffic Engineering Extensions to OSPF" [OSPF-TE] as a base for TLV definitions and defines some new TLVs and sub-TLVs to extend TE capabilities to IPv6 networks.

Requirements for establishing Multiprotocol Label Switching Traffic Engineering (MPLS-TE) Label Switched Paths (LSPs) that cross multiple Autonomous Systems (ASes) are described in [INTER-AS-TE-REQ]. As described in [INTER-AS-TE-REQ], 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] 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, two new separate LSAs are defined to advertise inter-AS TE information for OSPFv2 and OSPFv3, respectively, and three new sub-TLVs are added to the existing Link TLV to extend TE capabilities for inter-AS Traffic Engineering. The 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 OSPF. See Section 2.1 for a full list of non-objectives for this work.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Problem Statement

As described in [INTER-AS-TE-REQ], in the case of establishing an inter-AS TE LSP traversing 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 are currently being discussed. The per-domain method [PD-PATH] determines the path one domain at a time. The backward recursive method [BRPC] 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 [INTER-AS-TE-REQ].

The following features are explicitly excluded:

- o There is no attempt to distribute TE information from within one AS to another AS.
- o There is no mechanism proposed to distribute any form of TE reachability information for destinations outside the AS.
- o There is no proposed change to the PCE architecture or usage.
- o TE aggregation is not supported or recommended.
- o There is no exchange of private information between ASes.
- o No OSPF adjacencies are formed on the inter-AS link.

Note also that the extensions proposed in this document are used only to advertise information about inter-AS TE links. As such these extensions address an entirely different problem from L1VPN Auto-Discovery [L1VPN-OSPF-AD], which defines how TE information about links between Customer Edge (CE) equipment and Provider Edge (PE) equipment can be advertised in OSPF-TE alongside the auto-discovery information for the CE-PE links. There is no overlap between this document and [L1VPN-OSPF-AD].

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 so that it can compute a TE LSP segment across the local AS to one of those LSRs and forward the Path message to it and hence into the next AS. See Figure 1 for an example:




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 the 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 [PD-PATH], which clearly points out why the advertising of inter-AS links is desired.

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

- o List of all inter-AS TE links for the local AS.
- o TE properties of each inter-AS TE link.
- o AS number of the neighboring AS to which each inter-AS TE link is connected.

- o Identity (TE Router ID) of the neighboring ASBR to which each inter-AS TE link is connected.

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

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), but this information is also needed throughout the local AS if path computation function 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. [BRPC] 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 that the tree of paths provided by one PCE to its neighbor can be correlated, the identities of the ASBRs for each path need to be referenced, so the PCE must know the identities of the ASBRs in the remote AS reached by any inter-AS TE link, and, in order that it provides 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:




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 Path Computation Client (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 to which each inter-AS TE link is connected is particularly important.

3. Extensions to OSPF

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.

The extensions defined in this document allow an inter-AS TE link advertisement to be easily identified as such by the use of two new types of LSA, which are referred to as Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA. Three new sub-TLVs are added to the Link TLV to carry the information about the neighboring AS and the remote ASBR.

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. LSA Definitions

3.1.1. Inter-AS-TE-v2 LSA

For the advertisement of OSPFv2 inter-AS TE links, a new Opaque LSA, the Inter-AS-TE-v2 LSA, is defined in this document. The Inter-AS-TE-v2 LSA has the same format as "Traffic Engineering LSA", which is defined in [OSPF-TE].

The inter-AS TE link advertisement SHOULD be carried in a Type 10 Opaque LSA [RFC5250] if the flooding scope is to be limited to within the single IGP area to which the ASBR belongs, or MAY be carried in a Type 11 Opaque LSA [RFC5250] if the information is intended to reach all routers (including area border routers, ASBRs, and PCEs) in the AS. The choice between the use of a Type 10 (area-scoped) or Type 11 (AS-scoped) Opaque LSA is an AS-wide policy choice, and configuration control of it SHOULD be provided in ASBR implementations that support the advertisement of inter-AS TE links.

The Link State ID of an Opaque LSA as defined in [RFC5250] is divided into two parts. One of them is the Opaque type (8-bit), the other is the Opaque ID (24-bit). The value for the Opaque type of Inter-AS-TE-v2 LSA is 6 and has been assigned by IANA (see Section 6.1). The Opaque ID of the Inter-AS-TE-v2 LSA is an arbitrary value used to uniquely identify Traffic Engineering LSAs. The Link State ID has no topological significance.

The TLVs within the body of an Inter-AS-TE-v2 LSA have the same format as used in OSPF-TE. The payload of the TLVs consists of one or more nested Type/Length/Value triplets. New sub-TLVs specifically for inter-AS TE Link advertisement are described in Section 3.2.

3.1.2. Inter-AS-TE-v3 LSA

In this document, a new LS type is defined for OSPFv3 inter-AS TE link advertisement. The new LS type function code is 13 (see Section 6.1).

The format of an Inter-AS-TE-v3 LSA follows the standard definition of an OSPFv3 LSA as defined in [OSPFV3].

The high-order three bits of the LS type field of the OSPFv3 LSA header encode generic properties of the LSA and are termed the U-bit, S2-bit, and S1-bit [OSPFV3]. The remainder of the LS type carries the LSA function code.

For the Inter-AS-TE-v3-LSA, the bits are set as follows:

The U-bit is always set to 1 to indicate that an OSPFv3 router **MUST** flood the LSA at its defined flooding scope even if it does not recognize the LS type.

The S2 and S1 bits indicate the flooding scope of an LSA. For the Inter-AS-TE-v3-LSA, the S2 and S1 bits **SHOULD** be set to 01 to indicate that the flooding scope is to be limited to within the single IGP area to which the ASBR belongs, but **MAY** be set to 10 if the information should reach all routers (including area border routers, ASBRs, and PCEs) in the AS. The choice between the use of 01 or 10 is a network-wide policy choice, and configuration control **SHOULD** be provided in ASBR implementations that support the advertisement of inter-AS TE links.

The Link State ID of the Inter-AS-TE-v3 LSA is an arbitrary value used to uniquely identify Traffic Engineering LSAs. The LSA ID has no topological significance.

The TLVs within the body of an Inter-AS-TE-v3 LSA have the same format and semantics as those defined in [OSPF-V3-TE]. New sub-TLVs specifically for inter-AS TE Link advertisement are described in Section 3.2.

3.2. LSA Payload

Both the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA contain one top level TLV:

2 - Link TLV

For the Inter-AS-TE-v2 LSA, this TLV is defined in [OSPF-TE], and for the Inter-AS-TE-v3 LSA, this TLV is defined in [OSPF-V3-TE]. The sub-TLVs carried in this TLV are described in the following sections.

3.2.1. Link TLV

The Link TLV describes a single link and consists a set of sub-TLVs. The sub-TLVs for inclusion in the Link TLV of the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA are defined, respectively, in [OSPF-TE] and [OSPF-V3-TE], and the list of sub-TLVs may be extended by other documents. However, this document defines the following exceptions.

The Link ID sub-TLV [OSPF-TE] MUST NOT be used in the Link TLV of an Inter-AS-TE-v2 LSA, and the Neighbor ID sub-TLV [OSPF-V3-TE] MUST NOT be used in the Link TLV of an Inter-AS-TE-v3 LSA. Given that OSPF is an IGP and should only be utilized between routers in the same routing domain, the OSPF specific Link ID and Neighbor ID sub-TLVs are not applicable to inter-AS links.

Instead, the remote ASBR is identified by the inclusion of the following new sub-TLVs defined in this document and described in the subsequent sections.

21 - Remote AS Number sub-TLV

22 - IPv4 Remote ASBR ID sub-TLV

23 - IPv6 Remote ASBR ID sub-TLV

The Remote-AS-Number sub-TLV MUST be included in the Link TLV of both the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA. At least one of the IPv4-Remote-ASBR-ID sub-TLV and the IPv6-Remote-ASBR-ID sub-TLV SHOULD be included in the Link TLV of the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA. Note that it is possible to include the IPv6-Remote-ASBR-ID sub-TLV in the Link TLV of the Inter-AS-TE-v2 LSA, and to include the IPv4-Remote-ASBR-ID sub-TLV in the Link TLV of the Inter-AS-TE-v3 LSA because the sub-TLVs refer to ASBRs that are in a different addressing scope (that is, a different AS) from that where the OSPF LSA is used.

3.3. Sub-TLV Details

3.3.1. Remote AS Number Sub-TLV

A new sub-TLV, the Remote AS Number sub-TLV is defined for inclusion in the Link 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 REQUIRED in a Link TLV that advertises an inter-AS TE link.

The Remote AS Number sub-TLV is TLV type 21 (see Section 6.2), and is four 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 two octets are used for the AS number, as in current deployments, the left (high-order) two octets MUST be set to zero.

3.3.2. IPv4 Remote ASBR ID Sub-TLV

A new sub-TLV, which is referred to as the IPv4 Remote ASBR ID sub-TLV, can be included in the Link 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. This could be any stable and routable IPv4 address of the remote ASBR. Use of the TE Router Address TE Router ID as specified in the Router Address TLV [OSPF-TE] is RECOMMENDED.

The IPv4 Remote ASBR ID sub-TLV is TLV type 22 (see Section 6.2), and is four octets in length. Its 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 ASBR ID       |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

In OSPFv2 advertisements, the IPv4 Remote ASBR ID sub-TLV MUST be included if the neighboring ASBR has an IPv4 address. If the neighboring ASBR does not have an IPv4 address (not even an IPv4 TE Router ID), 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 a Link TLV in OSPFv2 or OSPFv3.

3.3.3. IPv6 Remote ASBR ID Sub-TLV

A new sub-TLV, which is referred to as the IPv6 Remote ASBR ID sub-TLV, can be included in the Link TLV when advertising inter-AS links. The IPv6 Remote ASBR ID sub-TLV specifies the identifier of the remote ASBR to which the advertised inter-AS link connects. This could be any stable, routable, and global IPv6 address of the remote ASBR. Use of the TE Router IPv6 Address IPv6 TE Router ID as specified in the IPv6 Router Address, which is specified in the IPv6 Router Address TLV [OSPF-V3-TE], is RECOMMENDED.

The IPv6 Remote ASBR ID sub-TLV is TLV type 24 (see Section 6.2), and is sixteen octets in length. Its 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 ASBR ID																																							
Remote ASBR ID (continued)																																							
Remote ASBR ID (continued)																																							
Remote ASBR ID (continued)																																							

In OSPFv3 advertisements, 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 a Link TLV in OSPFv2 or OSPFv3.

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 OSPF-TE [OSPF-TE]. 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 the ASBR MUST take precautions against excessive re-advertisements as described in [OSPF-TE].

Hellos MUST NOT be exchanged over the inter-AS link, and consequently, an OSPF 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 the Link Type carries an unknown value. They will continue to flood the LSA, but will not attempt to use the information received as if the link were an intra-AS TE link.

In the current operation of TE OSPF, the LSRs at each end of a TE link emit LSAs 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 LSAs 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 LSA 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 LSA that describes that device's 'view' of the link.

Only some essential TE information for the link needs to be advertised; i.e., the Link Type, the Remote AS number, and the Remote ASBR ID. 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 this information 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 OSPF security mechanisms.

There is no exchange of information between ASes, and no change to the OSPF security relationship between the ASes. In particular, since no OSPF adjacency is formed on the inter-AS links, there is no requirement for OSPF 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 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 this could be used to detect inconsistencies in configuration (e.g., the administration that originally supplied the information may be lying, or some manual misconfigurations or mistakes are made by the operators). For example, if a different remote AS number is received in a BGP OPEN [BGP] from that locally configured into OSPF-TE, as we describe here, then local policy SHOULD be applied to determine whether to alert the operator to a potential misconfiguration or to suppress the OSPF advertisement of the inter-AS TE link. 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 as described in [BGP] to provide authentication and integrity checks.

6. IANA Considerations

IANA has made the following allocations from registries under its control.

6.1. Inter-AS TE OSPF LSA

6.1.1. Inter-AS-TE-v2 LSA

IANA has assigned a new Opaque LSA type (6) to Inter-AS-TE-v2 LSA.

6.1.2. Inter-AS-TE-v3 LSA

IANA has assigned a new OSPFv3 LSA type function code (13) to Inter-AS-TE-v3 LSA.

6.2. OSPF LSA Sub-TLVs Type

IANA maintains the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registry with sub-registry "Types for sub-TLVs in a TE Link TLV". IANA has assigned three new sub-TLVs as follows (see Section 3.3 for details):

Value	Meaning
21	Remote AS Number sub-TLV
22	IPv4 Remote ASBR ID sub-TLV
24	IPv6 Remote ASBR ID sub-TLV

7. Acknowledgments

The authors would like to thank Adrian Farrel, Acee Lindem, JP Vasseur, Dean Cheng, and Jean-Louis Le Roux for their review and comments to this document.

8. References

8.1. Normative References

- [GMPLS-TE] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [OSPF] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [OSPF-TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [OSPF-V3-TE] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008.
- [OSPFV3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

8.2. Informative References

- [BGP] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [BRPC] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Inter-Domain Traffic Engineering Label Switched Paths", Work in Progress, April 2008.
- [INTER-AS-TE-REQ] Zhang, R., Ed., and J.-P. Vasseur, Ed., "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005.
- [L1VPN-OSPF-AD] Bryskin, I. and L. Berger, "OSPF-Based Layer 1 VPN Auto-Discovery", RFC 5252, July 2008.
- [PCE] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [PD-PATH] 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, February 2008.

Authors' Addresses

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

EMail: mach@huawei.com

Renhai Zhang
Huawei Technologies Co., Ltd.
KuiKe Building, No.9 Xinxu Rd.
Hai-Dian District
Beijing, 100085
P.R. China

EMail: zhangrenhai@huawei.com

Xiaodong Duan
China Mobile
53A,Xibianmennei Ave,Xunwu District
Beijing, China

EMail: duanxiaodong@chinamobile.com

