Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 30, 2013

P. Psenak, Ed.
S. Previdi, Ed.
C. Filsfils
Cisco Systems, Inc.

June 28, 2013

OSPF Extensions for Segment Routing
draft-psenak-ospf-segment-routing-extensions-00

Abstract

Segment Routing (SR) allows for a flexible definition of end-to-end paths within IGP topologies by encoding paths as sequences of topological sub-paths, called "segments". These segments are advertised by the link-state routing protocols (IS-IS and OSPF).

This draft describes the necessary OSPF extensions that need to be introduced for Segment Routing.

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].

Table of Contents

1. Introduction

Segment Routing (SR) allows for a flexible definition of end-to-end paths within IGP topologies by encoding paths as sequences of topological sub-paths, called "segments". These segments are advertised by the link-state routing protocols (IS-IS and OSPF). Two types of segments are defined, Prefix segments and Adjacency segments. Prefix segments represent an ecmp-aware shortest-path to a prefix (or a node), as per the state of the IGP topology. Adjacency segments represent a hop over a specific adjacency between two nodes in the IGP. A prefix segment is typically a multi-hop path while an adjacency segment, in most of the cases, is a one-hop path. SR's control-plane can be applied to both IPv6 and MPLS data-planes, and do not require any additional signaling (other than the regular IGP). For example, when used in MPLS networks, SR paths do not require any LDP or RSVP-TE signaling. Still, SR can interoperate in the presence of LSPs established with RSVP or LDP .

This draft describes the necessary OSPF extensions that need to be introduced for Segment Routing.

Segment Routing architecture is described in [draft-filsfils-rtgwg-segment-routing-00].

Segment Routing use cases are described in [draft-filsfils-rtgwg-segment-routing-use-cases-00].

2. OSPFv2 Segment Routing Capability TLV

Segment Routing requires each router to advertise various capabilities associated with Segment Routing throughout the autonomous system.

For the purposes of Segment Routing, a new TLV is defined in Router Information Opaque LSA (defined in [RFC4970]: the Segment Routing Capability TLV (SR-Cap TLV).

If the SR-Cap TLV appear in Router Information Opaque LSA, it MUST appear only once and has following format:

```
 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               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       SR Capability Flags                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+-                        Sub-TLVs (variable)                  -+
|                              ...                              |
```

where:

Type: TBA.

Length: A 16-bit field that indicates the length of the value portion in octets. Set to 12.

SR Capabilities Flags: 2 octets field of following flags:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|R|X|                                                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

R-flag: SR Control-Plane flag. If set, then the advertising router is capable of SR control-plane.

X-Flag: Index flag. If set, then each SID with the G-Flag set that is advertised by this router represents an index in the SID space also advertised by this router (using the encodings defined below). Indexed SID values are described in [draft-filsfils-rtgwg-segment-routing-00]. If the X-Flag is not set, then SID values 0-63 are reserved and MUST NOT be used by the SR control plane for either node or adjacency segments. Any SID which is received with a value 0-63 and the X-Flag unset MUST be ignored and the router SHOULD log an error.

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

RI LSA can be advertised at any of the defined flooding scopes (link, area, or autonomous system (AS)). For the purpose of the SR-Cap TLV propagation, autonomous system scope flooding is required.

Within the SR-Cap TLV, new Sub-TLVs are defined:

SID Range Sub-TLV

Segment Routing Mirroring Context Sub-TLV (SRMC)

2.1. SID Range Sub-TLV

The SID Range Sub-TLV allows a router to advertise its SID space to other routers in the area.

The SID Range Sub-TLV is a Sub-TLV of the SR-Cap-TLV. It MUST appear only once and has following format:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        First SID Value                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Last  SID Value                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

First SID Value and Last SID Value represent the local SID range of the router. The semantic and procedures of these values are described in [draft-filsfils-rtgwg-segment-routing-use-cases-00].

2.2. Segment Routing Mirroring Context Sub-TLV (SRMC)

Segment Routing allows the use of a context SID that determines a specific use case or applicability. Use cases are described in [draft-filsfils-rtgwg-segment-routing-use-cases-00]. One of such use cases is the service mirroring, for which the new sub-TLV is defined.

Another example of mirroring use case is described in [I-D.minto-rsvp-lsp-egress-fast-protection].

The SRMC Sub-TLV is an optional sub-TLV of SR-Cap-TLV. It MAY appear in SR-Cap-TLV more than once and has following format:

```
 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    |          SRMC Flags           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Mirrored Address (4 octets)            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Context-SID (4 octets)                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type: TBA.

Length: 2 octets.

SRMC Flags: 2 octets field of flags. Following flags are defined:

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|R|                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

R-flag: Segment Routing Identifier. If set, the Mirrored Address contains a SID.

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

Mirrored Address: IPv4 address or SID value representing the mirrored element.

Context-SID: 32 bit value of Segment Identifier for the Mirroring Service.

3. Segment Routing Identifiers

Segment Routing defines two types of Segment Identifiers: Prefix-SID and Adjacency-SID.

For the purpose of the Prefix-SID and Adjacency-SID advertisement new Opaque LSAs (defined in [RFC5250]) are defined. These new LSAs are defined as generic containers that can be used in order to advertise any additional attributes associated with the prefix or link. These new Opaque LSAs are complementary to the existing LSAs and are not aimed to replace any of the existing LSAs.

3.1. OSPFv2 Extended Prefix Opaque LSA type

A new Opaque LSA (defined in [RFC5250]) is defined in OSPFv2 in order to advertise additional prefix attributes: OSPFv2 Extended Prefix Opaque LSA.

Multiple OSPFv2 Extended Prefix Opaque LSAs can be advertised by a single router. Flooding scope of the OSPFv2 Extended Prefix Opaque LSA depends on the content inside the LSA and is in control of the originating router.

The format of the OSPFv2 Extended Prefix Opaque LSA 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            LS age             |     Options   |  9, 10, or 11 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Opaque type  |                  Instance                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Advertising Router                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     LS sequence number                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         LS checksum           |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+-                            TLVs                             -+
|                             ...                               |
```

The format of the TLVs within the body of the LSA is the same as the format used by the Traffic Engineering Extensions to OSPF defined in [RFC3630]. The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets. The format of each TLV is:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Value...                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Length field defines the length of the value portion in octets. The TLV is padded to 4-octet alignment; padding is not included in the length field. Nested TLVs are also 32-bit aligned. Unrecognized types are ignored.

3.1.1. OSPF Extended Prefix TLV

The OSPF Extended Prefix TLV is used in order to advertise additional attributes associated with the prefix. Multiple OSPF Extended Prefix TLVs MAY be carried in each OSPFv2 Extended Prefix Opaque LSA, however all prefixes included in the single OSPFv2 Extended Prefix Opaque LSA MUST have the same flooding scope. The structure of the OSPF Extended Prefix 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Route Type   | Prefix Length |     AF        |   Reserved    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Address Prefix (variable)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Sub-TLVs (variable)                     |
+-                                                             -+
|                                                               |
```

Type is TBA.

Length is variable

Route type: type of the OSPF route. Supported types are:

1 - intra-area
3 - inter-area
5 - external
7 - NSSA external

Prefix length: length of the prefix

AF:
0 - IPv4 unicast

Address Prefix: the prefix itself encoded as an even multiple of 32-bit words, padded with zeroed bits as necessary. This encoding consumes ((PrefixLength + 31) / 32) 32-bit words. The default route is represented by a prefix of length 0.

3.1.2. Prefix SID Sub-TLV

The Prefix SID Sub-TLV is a Sub-TLV of the OSPF Extended Prefix TLV. It may MAY appear more than once for the same prefix and has following format:

```
 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           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Flags             |    MT-ID      |   Reserved    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Prefix SID                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type is TBA.

Length: A 16-bit field that indicates the length of the value portion in octets. Set to 8.

Flags: 2 octets field. The following flags are defined:

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|N|P|G|M|                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

N-Flag: Node-SID flag. If set, then the Prefix-SID refers to the router identified by the prefix. Typically, the N-Flag is set on Prefix-SIDs attached to a router loopback address. The N-Flag is set when the Prefix-SID is a Node- SID as described in [draft-filsfils-rtgwg-segment-routing-00].

P-Flag: no-PHP flag. If set, then the penultimate hop MUST NOT pop the Prefix-SID before delivering the packet to the node that advertised the Prefix-SID.

G-Flag: Global flag. When set, the SID value has global significance which means the SID has been allocated from the Segment Routing Global Block (SRGB) as described in [draft-filsfils-rtgwg-segment-routing-00]. When unset, the SID value has local (within the router) significance which means the SID has NOT been allocated from the SRGB. When the Prefix-SID Sub-TLV carries an IGP Prefix SID, then the G-flag MUST be set.

M-Flag: Mapping Server Flag. If set, the SID is advertised from the Segment Routing Mapping Server functionality as described in [draft-filsfils-rtgwg-segment-routing-use-cases-00].

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

Prefix Segment Identifier (SID): 32 bits of Prefix-SID.

If multiple Prefix-SIDs are advertised for the same prefix, the receiving router MUST use the first encoded SID and MAY use the subsequent ones.

The No-PHP flag MUST be set on the Prefix-SIDs allocated to inter-area prefixes that are originated by the router based on intra-area reachability.

3.2. Adjacency Segment Identifier (Adj-SID)

An Adjacency Segment Identifier (Adj-SID) represents a router adjacency in Segment Routing. At the current stage of Segment Routing architecture it is assumed that the Adj-SID value has local significance (to the router). The Adj-SID value is encoded as a 32 bit number.

3.2.1. OSPFv2 Extended Link Opaque LSA

A new Opaque LSA (defined in [RFC5250] is defined in OSPFv2 in order to advertise additional link attributes: the OSPFv2 Extended Link Opaque LSA.

The OSPFv2 Extended Link Opaque LSA has an area flooding scope. Multiple OSPFv2 Extended Link Opaque LSAs can be advertised by a single router in an area.

The format of the OSPFv2 Extended Link Opaque LSA 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            LS age             |     Options   |       10      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Opaque type  |                   Instance                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Advertising Router                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      LS sequence number                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         LS checksum           |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+-                            TLVs                             -+
|                             ...                               |
```

The format of the TLVs within the body of LSA is the same as the format used by the Traffic Engineering Extensions to OSPF defined in [RFC3630]. The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets. The format of each TLV is:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Value...                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Length field defines the length of the value portion in octets. The TLV is padded to 4-octet alignment; padding is not included in the length field. Nested TLVs are also 32-bit aligned. Unrecognized types are ignored.

3.2.2. OSPFv2 Extended Link TLV

OSPFv2 Extended Link TLV is used in order to advertise various attributes of the link. It describes a single link and is constructed of a set of Sub-TLVs. There are no ordering requirements for the Sub-TLVs. Only one Extended Link TLV SHALL be carried in each Extended Link Opaque LSA, allowing for fine granularity changes in the topology.

The Extended Link TLV has following format:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Link-Type |                 Reserved                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Link ID                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Link Data                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Sub-TLVs (variable)                      |
+-                                                             -+
|                                                               |
```

where:

Type is TBA.

Length is variable.

Link-Type: as defined in section A.4.2 of [RFC2328].

Link-ID: as defined in section A.4.2 of [RFC2328].

Link Data: as defined in section A.4.2 of [RFC2328].

3.2.3. Adj-SID sub-TLV

Adj-SID is an optional Sub-TLV of the Extended Link TLV. It MAY appear multiple times in Extended Link TLV. Examples where more than one Adj-SID may be used per neighbor are described in [draft-filsfils-rtgwg-segment-routing-use-cases-00]. The structure of the Adj-SID 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   MT-ID       | Reserved      |             Flags             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Adjacency SID                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Neighbor ID   (optional)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Adjacency SID (optional)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

                                  ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Neighbor ID   (optional)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Adjacency SID (optional)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type is TBA.

Length is variable.

MT-ID: Multi-Topology ID (as defined in [RFC4915].

Flags: 2 octets field of following flags:

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|F|G|S|T|                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

F-Flag: FA flag. If set, then Adj-SID refers to a Forwarding Adjacency.

G-Flag: Global flag. When set, the SID value has global significance which means the SID has been allocated from the Segment Routing Global Block (SRGB) as described in [draft-filsfils-rtgwg-segment-routing-00]. When unset, the SID value has local (within the router) significance which means the SID has NOT been allocated from the SRGB. When the Adj-SID Sub-TLV carries the IGP Adjacency SID, then the G-Flag MUST be unset.

S-Flag: Secondary SID. If set the SID represents the secondary SID. One case where secondary SID may be useful is the parallel adjacency case as described in [draft-filsfils-rtgwg-segment-routing-00].

T-Flag: Protected-flag: set if the Adj-SID refer to an adjacency being protected (e.g.: using IPFRR or MPLS-FRR) as described in [draft-filsfils-rtgwg-segment-routing-use-cases-00].

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

Adj-SID: the 32 bits of Adjacency Segment Identifier.

Neighbor ID: the Neighbor ID on broadcast segment to which the SID belongs. Only used if more then one SID is advertised in the Adj-SID Sub-TLV.

The first Adjacency SID value in TLV is mandatory and advertises the Adjacency Segment Identifier for the neighbor that is identified by the Extended Link TLV itself. Subsequent Neighbor ID/Adjacency SID tuples are optional and represent Adj-SIDs for other neighbors on the broadcast segment.

A SR capable router SHOULD allocate an Adj-SID for each of its adjacencies and set the T-Flag when the adjacency is protected by a FRR mechanism (IP or MPLS) as described in [draft-filsfils-rtgwg-segment-routing-use-cases-00].

3.2.4. Explicit Path sub-TLV

Explicit Path Sub-TLV is a Sub-TLV of the Extended Link TLV. The structure of the Explicit Path 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          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Flags    |    MT-ID      |       Reserved                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Explicit Route Hop #1                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Explicit Route Hop #...                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type is TBA.

Length is variable.

Flags: 16-bit flag field where following flags are defined:

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|R|             |
+-+-+-+-+-+-+-+-+
```

where:

R-Flag: Segment Routing flag. When set, the content of the Explicit Route Object contains SIDs.

MT-ID: the Multi-Topology ID (as defined in [RFC4915].

Explicit Route Hop: the address of the explicit route hop. IPv4 and SR hops are encoded using a 32 bit value.

4. Elements of Procedure

4.1. Intra-area Segment routing in OSPFv2

The OSPFv2 node that supports segment routing MAY advertise Prefix-SIDs for any prefix that is advertising reachability to (e.g. loopback IP address) as described in Section 3.1. If multiple routers advertise Prefix-SID for the same prefix, then the Prefix-SID MUST be the same. This is required in order to allow traffic load-balancing if multiple equal cost paths to the destination exist in the network.

Prefix-SID can also be advertised by the SR Mapping Servers functionality (as described in [draft-filsfils-rtgwg-segment-routing-use-cases-00]) that advertise Prefix-SID for remote prefixes which exist in the network. Multiple SR Mapping Servers can advertise Prefix-SID for the same prefix, in which case the same Prefix-SID MUST be advertised by all of them. Flooding scope of the OSPF Extended Prefix Opaque LSA that is generated by the SR Mapping Server could be either area scope or autonomous system scope and is decided based on the configuration of the SR Mapping Server.

Prefix-SID advertisements do not contribute to the prefix reachability.

4.2. Inter-area Segment routing in OSPFv2

In order to support SR in a multi-area environment, OSPFv2 must propagate SR information between areas. The following procedure is used in order to propagate Prefix SIDs between areas.

When an OSPF ABR advertises a Type-3 Summary LSA from an intra-area route to all its connected areas, it will also originate an Extended Prefix Opaque LSA, as described in Section 3.1. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-scope. The route-type in OSPF Extended Prefix TLV is set to inter-area. The Prefix-SID Sub-TLV will be included in this LSA and the Prefix-SID value will be set as follows:

The ABR will look at it's best path to the prefix in the source area and find out the advertising router associated with its best path to that prefix.

If no Prefix-SID was advertised for the prefix in the source area by the router that contributes to the best path to the prefix, then the ABR will use the Prefix-SID advertised by any other router (e.g.: a Prefix-SID coming from an SR Mapping Server as defined in [draft-filsfils-rtgwg-segment-routing-use-cases-00]) when propagating Prefix-SID for the prefix to other areas.

When an OSPF ABR advertises Type-3 Summary LSAs from an inter-area route to all its connected areas it will also originate an Extended Prefix Opaque LSA, as described in Section 3.1. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-scope. The route-type in OSPF Extended Prefix TLV is set to inter-area. The Prefix-SID Sub-TLV will be included in this LSA and the Prefix-SID will be set as follows:

The ABR will look at its best path to the prefix in the source area and find out the advertising router associated with its best path to that prefix.

The ABR will then look if such router advertised a Prefix-SID for the prefix and use it when advertising the Prefix-SID to other connected areas.

If no Prefix-SID was advertised for the prefix in the source area by the ABR that contributes to the best path to the prefix, the originating ABR will use the Prefix-SID advertised by any other router (e.g.: a Prefix-SID coming from an SR Mapping Server as defined in [draft-filsfils-rtgwg-segment-routing-use-cases-00]) when propagating Prefix-SID for the prefix to other areas.

4.3. SID for External Prefixes

Type-5 LSAs are flooded domain wide. When an ASBR, which supports SR, generates Type-5 LSAs, it should also originate Extended Prefix Opaque LSAs, as described in Section 3.1. The flooding scope of the Extended Prefix Opaque LSA type is set to AS-scope. The route-type in OSPF Extended Prefix TLV is set to external. Prefix-SID Sub-TLV is included in this LSA and the Prefix-SID value will be set to the SID that has been reserved for that prefix.

When a NSSA ASBR translates Type-7 LSAs into Type-5 LSAs, it should also advertise the Prefix-SID for the prefix. The NSSA ABR determines its best path to the prefix advertised in the translated Type-7 LSA and finds the advertising router associated with such path. If such advertising router has advertised a Prefix-SID for the prefix, then the NSSA ASBR uses it when advertising the Prefix-SID for the Type-5 prefix. Otherwise the Prefix-SID advertised by any other router will be used (e.g.: a Prefix-SID coming from an SR Mapping Server as defined in [draft-filsfils-rtgwg-segment-routing-use-cases-00]).

4.4. Advertisement of Adj-SID

The Adjacency Segment Routing Identifier (Adj-SID) is advertised using the Adj-SID Sub-TLV as described in Section 3.2.

4.4.1. Advertisement of Adj-SID on Point-to-Point Links

The Adj-SID MUST be advertised for any adjacency on a p2p link that is in a FULL state. Adj-SID MAY be advertised for any adjacency on p2p link that is in a state 2-Way or higher. If the adjacency on a p2p link transitions from the FULL state, then the Adj-SID for that adjacency MAY be removed from the area. If the adjacency transitions to a state lower then 2-Way, then the Adj-SID MUST be removed from the area.

4.4.2. Adjacency SID on Broadcast Interfaces

Broadcast networks in OSPF are represented by a star topology where the Designated Router (DR) is the central point all other routers on the broadcast network connect to. As a result, the routers on the broadcast network advertise only their adjacency to DR and BDR. Routers that are neither DR nor BDR do not form, and do not advertise, adjacencies between them. They, however, maintain a 2-Way state between them.

When Segment Routing is used, each router in the broadcast segment must advertise the Adj-SID for each of its neighbors. To satisfy

this requirement Adj-SID on broadcast networks are advertised as follows:

If the router is acting as DR on a broadcast segment, the Adj-SID MUST be advertised for each adjacency in a FULL state and MAY be advertised for any adjacency in a state 2-way or higher. When the DR advertises the Adj-SID Sub-TLV for the broadcast link, the first Adj-SID in the Adj-SID sub-TLV corresponds to the adjacency of DR to itself and as such MUST be set to zero and ignored.

If the router is acting as BDR or as DR-Other, then the Adj-SID MUST be advertised for each adjacency in a state 2-Way or higher.

5. IANA Considerations

TBD

6. Manageability Considerations

TBD

7. Security Considerations

TBD

8. Acknowledgements

We would like to thank Anton Smirnov for his contribution to the content of this document.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, June 2007.

[RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.

[RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

9.2. Informative References

[I-D.minto-rsvp-lsp-egress-fast-protection]
Jeganathan, J., Gredler, H., and Y. Shen, "RSVP-TE LSP egress fast-protection", draft-minto-rsvp-lsp-egress-fast-protection-02 (work in progress), April 2013.

[draft-filsfils-rtgwg-segment-routing-00]
Previdi, S. and C. Filsfils, "Segment Routing", May 2013.

[draft-filsfils-rtgwg-segment-routing-use-cases-00]
Filsfils, C., "Segment Routing Use Cases", May 2013.

Authors' Addresses

Peter Psenak (editor)
Cisco Systems, Inc.
Apollo Business Center
Mlynske nivy 43
Bratislava 821 09
Slovakia

Email: ppsenak@cisco.com

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico, 200
Rome 00142
Italy

Email: sprevidi@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels,
Belgium

Email: cfilsfil@cisco.com