Network Working Group
Internet Draft
Category: Standard Track
Expires: March 2006

J.L. Le Roux (Editor)
France Telecom
J.P. Vasseur (Editor)
Cisco System Inc.
Yuichi Ikejiri
NTT

October 2005

IGP protocol extensions for Path Computation Element (PCE) Discovery

draft-leroux-pce-disco-proto-igp-00.txt

Status of this Memo

Abstract

In various situations it would be highly useful for a Path Computation Client (PCC) to be able to dynamically and automatically discover a set of Path Computation Element(s) (PCE), along with some of information relevant for PCE selection. When the PCE is an LSR participating to the IGP, or even a server participating passively to the IGP, a simple and efficient way for PCE discovery, consists of relying on IGP flooding. For that purpose this document defines simple OSPF and ISIS extensions for the advertisement of PCE Discovery information within and across IGP areas.

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.

Table of Contents

1. Note

This document specifies new TLVs and sub-TLVs to be carried within the OSPF Router information LSA ([OSPF-CAP]) and ISIS Router Capability TLV ([ISIS-CAP]). Because this document does not introduce any new element of procedure it will be discussed within the PCE Working Group with a review of the OSPF and ISIS Working Groups. Furthermore, once stabilized, it may be decided to split the document in two documents addressing the OSPF and ISIS aspects respectively.

2. Terminology

Terminology used in this document

LSR: Label Switch Router

TE LSP: Traffic Engineered Label Switched Path

PCE: Path Computation Element: an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph, and applying computational constraints.

PCC: Path Computation Client: any client application requesting a path computation to be performed by a Path Computation Element.

IGP Area: OSPF Area or ISIS level

ABR: IGP Area Border Router (OSPF ABR or ISIS L1L2 router)

AS: Autonomous System

ASBR: AS Border Router

Intra-area TE LSP: A TE LSP whose path does not cross IGP area boundaries.

Inter-area TE LSP: A TE LSP whose path transits through two or more IGP areas.

Inter-AS MPLS TE LSP: A TE LSP whose path transits through two or more ASes or sub-ASes (BGP confederations).

Domain: any collection of network elements within a common sphere of address management or path computational responsibility. Examples of domains include IGP areas and Autonomous Systems.

3. Introduction

[PCE-ARCH] describes the motivations and architecture for a PCE-based path computation model for MPLS and GMPLS TE LSPs. The model allows the separation of PCE from PCC (also referred to as non co-located PCE) and allows cooperation between PCEs. This relies on a communication protocol between PCC and PCE, and between PCEs, whose generic requirements are listed in [PCE-COM-REQ].

The PCE architecture requires, of course, that a PCC be aware of the location of one or more PCEs in its domain, and also potentially of some PCEs in other domains, e.g. in case of inter-domain path computation.

A network may comprise a large number of PCEs with potentially distinct capabilities. In such context it would be highly desirable to have a mechanism for automatic and dynamic PCE discovery, which would allow PCCs to automatically discover a set of PCEs, including information required for PCE selection, and to dynamically detect new PCEs or any modification of PCE information. This includes the discovery by a PCC of a set of one or more PCEs in its domain, and potentially in some other domains. The latter is a desirable function in the case of inter-domain path computation for example. Detailed requirements for such a PCE discovery mechanism are described in [PCE-DISCO-REQ].

When PCCs are LSRs participating to the IGP (OSPF, ISIS), and PCEs are LSRs participating to the IGP or any network servers participating to the IGP, an efficient mechanism for PCE discovery consists of relying on IGP advertisement.

This document defines OSPF and ISIS extensions allowing a PCE participating to the IGP to advertise its location along with some information useful for PCE selection so as to satisfy dynamic PCE discovery requirements set forth in [PCE-DISC-REQ].

Generic capability mechanisms have been defined in [OSPF-CAP] and [ISIS-CAP] for OSPF and ISIS respectively the purpose of which is to allow a router to advertise its capability within an IGP area or an entire routing domain. This perfectly fits with the aforementioned dynamic PCE discovery requirements. Thus, a new TLV (named the PCE Discovery (PCED) TLV) is defined for ISIS and OSPF to be carried within the ISIS Capability TLV ([ISIS-CAP]) for ISIS and the OSPF Router Information LSA ([OSPF-CAP]).

The PCE discovery information is detailed in section 3. Protocol extensions and procedures are defined in section 4 and 5 for ISIS and OSPF respectively. This document does not define any new OSPF or ISIS

element of procedure but how the procedures defined in [OSPF-CAP] and [ISIS-CAP] should be used.

The routing extensions defined in this document allow for PCE discovery within and across IGP areas. Solutions for PCE discovery across AS boundaries are for further study.

4. PCE Discovery Information

4.1. Description

The PCE Discovery information allows for dynamic discovery of PCEs. This information is advertised by means of IGP advertisements by PCE(s) participating to the IGP. It allows all PCCs participating to the IGP to dynamically discover PCEs along with information useful for PCE selection.

Such dynamic PCE discovery has various advantages:

- Simplified PCC configuration,
- Reduces risk of misconfiguration,
- Dynamic detection of a new PCE,
- Dynamic detection of any change in PCE information
- Dynamic detection of PCE aliveness (PCE liveness may require additional mechanisms provided by the PCC-PCE communication protocol)

4.2. Mandatory versus optional PCE information

The PCE Discovery information is comprised of:

- The PCE location: This is a set of one ore more IPv4 and or IPv6 addresses used to reach the PCE. These are basically loopback addresses always reachable provided that the PCE is still alive.

- The PCE inter-domain functions: this refers to the PCE path computation scope (i.e. inter-area, inter-AS, inter-layer).

- The PCE domain(s): This is the set domain(s) where the PCE has visibility and can compute paths.

- The set of destination domain(s) towards which a PCE can compute paths.

- A set of general PCE capabilities (e.g. support for request prioritization) and path computation specific PCE capabilities (e.g. supported constraints, supported objective functions) These are two variable length sets of bits flags, where each bit represent a given PCE capability.

4.3. Flooding scope

The PCE Discovery information can be flooded locally within the IGP

area the PCE belongs to, or globally across the entire routing domain. Note that some PCEs may belong to multiple areas, in which case the flooding scope can correspond to these areas. This could be

the case of an ABR for instance that can advertise its PCE information within the backbone area and/or a subset of its attached IGP area(s).

4.4. Frequency of change

The rate at which PCE information is advertised must be controlled so as to not impact by any mean the IGP scalability. Changes in PCE information may occur as result of PCE configuration updates, PCE deployment/activation or PCE deactivation/suppression. Hence, this information is not expected to change frequently.

5. OSPF extensions

5.1. The OSPF PCED TLV

The OSPF PCE Discovery TLV (PCED TLV) is made of a set of non-ordered sub-TLVs.

The format of the OSPF PCED TLV and its sub-TLVs is the identical as the TLV format used by the Traffic Engineering Extensions to OSPF [OSPF-TE]. That is, the TLV is composed of 2 octets for the type, 2 octets specifying the TLV length and a value field. The Length field defines the length of the value portion in octets.
The TLV is padded to four-octet alignment; padding is not included in the length field (so a three octet value would have a length of three, but the total size of the TLV would be eight octets). Nested TLVs are also 32-bit aligned. Unrecognized types are ignored. All types between 32768 and 65535 are reserved for vendor-specific extensions. All other undefined type codes are reserved for future assignment by IANA.

The OSPF PCED TLV has the following format:

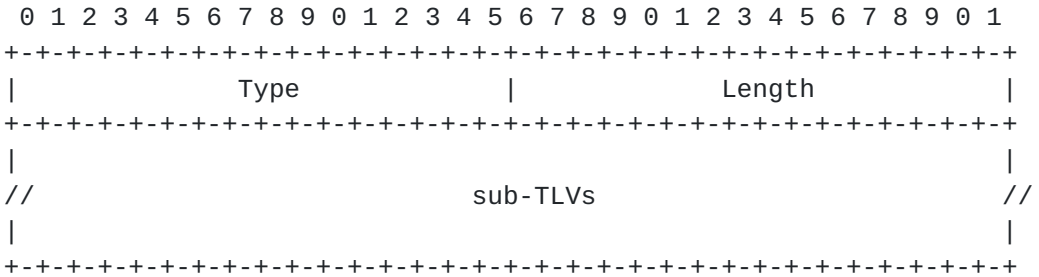

Type To be defined by IANA (suggested value=2)
Length Variable
Value This comprises one or more sub-TLVs

Sub-TLVs types are under IANA control.

Currently five sub-TLVs are defined (type values to be assigned by IANA):

Sub-TLV type	Length	Name
1	variable	PCE-ADDRESS sub-TLV
2	4	PATH-SCOPE sub-TLV
3	variable	PCE-DOMAINS sub-TLV
4	variable	PCE-DEST-DOMAINS sub-TLV
5	variable	GENERAL-CAP sub-TLV
6	variable	PATH-COMP-CAP sub-TLV

The sub-TLVs PCE-ADDRESS and PATH SCOPE MUST always be present within the PCED TLV.

The sub-TLVs PCE-DOMAINS and DEST-DOMAINS, MUST only be present in some particular inter-domain cases.

The GENERAL-CAP and PATH-COMP-CAP are optional and MAY be present in the PCED TLV to facilitate PCE selection.

Any non recognized sub-TLV MUST be silently ignored.

Additional sub-TLVs could be added in the future to advertise additional information.

The PCED TLV is carried within an OSPF Router Information LSA which is defined in [OSPF-CAP].

5.1.1. PCE-ADDRESS sub-TLV

The PCE-ADDRESS sub-TLV specifies the IP address to be used to reach the PCE. This address will typically be a loop-back address that is always reachable, provided that the PCE is alive.
The PCE-ADDRESS sub-TLV is mandatory; it MUST be present within the PCED TLV.

The format of the PCE-ADDRESS sub-TLV is as follows:

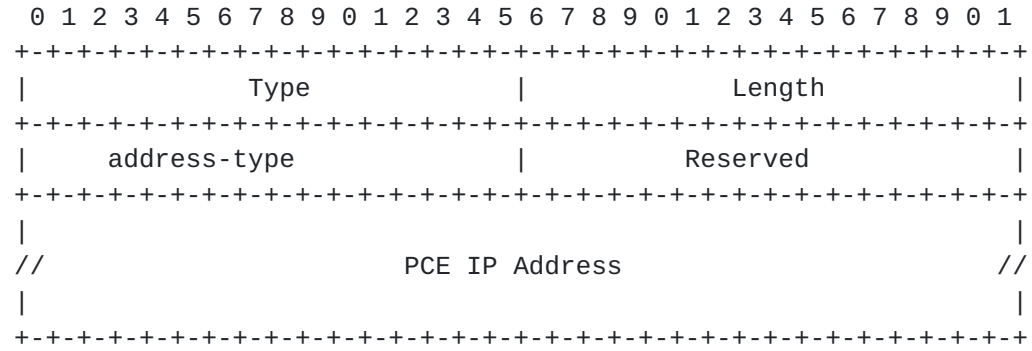

PCE-ADDRESS sub-TLV format

Type   To be assigned by IANA (suggested value =1)
Length   8 or 20

```
        Address-type:
                     1   IPv4
                     2   IPv6
        PCE IP Address: The IP address to be used to reach the PCE.
                        This is the address that will be used for
                        setting up PCC-PCE communication sessions.
```

The PCE-ADDRESS sub-TLV MUST appear at least once in the PCED sub-TLV originated by a PCE. It MAY appear multiple times, for instance when the PCE has both an IPv4 and IPv6 address.

5.1.2. PATH-SCOPE sub-TLV

The PATH-SCOPE sub-TLV indicates the PCE path computation scope; in other words the ability of the PCE to compute or take part into the computation of intra-area, inter-area, inter-AS or inter-layer_TE LSP(s).

The PATH-SCOPE sub-TLV is mandatory; it MUST be present within the PCED TLV. There MUST be exactly one PATH-SCOPE sub-TLV within each PCED TLV.

The PATH-SCOPE sub-TLV contains a set of bit flags indicating the supported path scopes (intra-area, inter-area, inter-AS, inter-layer) and two fields indicating PCE preferences.

The PATH-SCOPE sub-TLV has the following format:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|1|2|3|4|5|   Reserved        |PrefR|PrefS|       Reserved    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
        Type     To be defined by IANA (suggested value =3)
        Length   Variable
        Value    This comprises a 2 bytes flag where each bit
                 represents a supported path scope, as well as two
                 preference fields allowing to specify PCE preferences.

        The following bits are defined:

        Bit      Path Scope

         0      L bit: Can compute intra-area path
         1      R bit: Can act as PCE for inter-area TE LSPs
```

computation

2 Rd bit: Can act as a default PCE for inter-area TE LSPs computation

```
        3      S bit: Can act as PCE for inter-AS TE LSPs computation
        4      Sd bit: Can act as a default PCE for inter-AS TE LSPs
                       computation
        5      Y bit: Can compute or take part into the computation of
                      paths across layers.
```

Pref-R field: PCE s preference for inter-area path computation

Pref-S field: PCE s preference for inter-AS path computation

The bits L, R, S and Y bits are set when the PCE can act as a PCE for intra-area, inter-area, inter-AS and inter-layer TE LSPs computation respectively. These bits are non exclusive.

When set the Rd bit indicates that the PCE can act as a default PCE for inter-area TE LSPs computation, it means that it can compute path for any destination area. Similarly, when set the Sd bit indicates that the PCE can act as a default PCE for inter-AS TE LSPs computation, it means that it can compute path for any destination AS.

When the Rd bit is set the PCE-DEST-DOMAIN TLV (see 5.1.4) does not comprise any Area ID DOMAIN sub-TLV.
When the Sd bit is set the PCE-DEST-DOMAIN TLV does not comprise any AS DOMAIN sub-TLV.

The PrefR and PrefS fields are 3-bit long and allow the PCE to specify a preference where 7 reflects the highest preference. For the sake of illustration, consider the situation where N ABRs act as PCEs for inter-area TE LSPs computation. An operator may decide to configure a preference to each PCE that could be used to load balance the path computation load with respect to their respective CPU capacity. The algorithms used by a PCC to load balance its path computation requests according to such PCE s preference is out of the scope of this document.

5.1.3. PCE-DOMAINS sub-TLV

The PCE-DOMAINS sub-TLV specifies the set of domains (areas, AS) where the PCE has topology visibility and can compute paths. It contains a set of one or more sub-TLVs where each sub-TLV identifies a domain.

The PCE-DOMAINS sub-TLV MUST only be present when PCE domains cannot be inferred by other IGP information, for instance when the PCE is inter-domain capable (i.e. when the R bit or S bit is set) and the flooding scope the entire routing domain.

The PCE-DOMAINS sub-TLV has the following format:

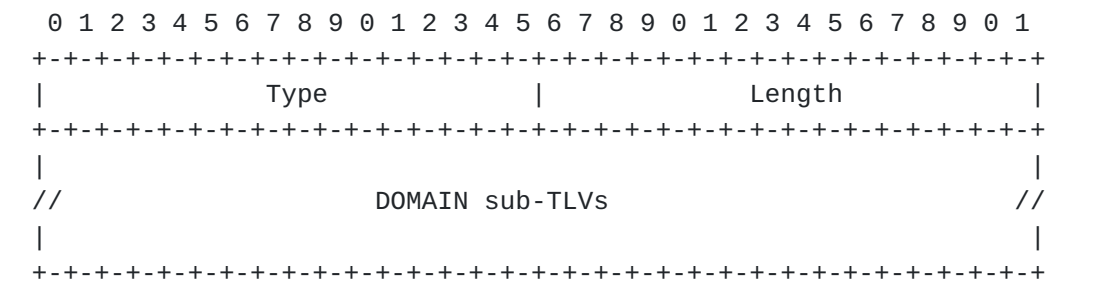

Type To be defined by IANA (suggested value =3)
Length Variable
Value This comprises a set of one or more DOMAIN sub-TLVs where each DOMAIN sub-TLV identifies a domain where the PCE has topology visibility and can compute TE LSP paths.

Sub-TLVs types are under IANA control.

Currently three DOMAIN sub-TLVs are defined (suggested type values to be assigned by IANA):

Sub-TLV type	Length	Name
1	variable	IPv4 area ID sub-TLV
2	variable	IPv6 area ID sub-TLV
3	variable	AS number sub-TLV

The PCE-DOMAINS sub-TLV MUST include at least one DOMAIN sub-TLV. Note than when the PCE visibility is an entire AS, the PCE-DOMAINS sub-TLV MUST uniquely include one AS number sub-TLV.

5.1.3.1. IPv4 area ID DOMAIN sub-TLV

The IPv4 area ID DOMAIN sub-TLV carries an IPv4 OSPF area identifier. It has the following format:

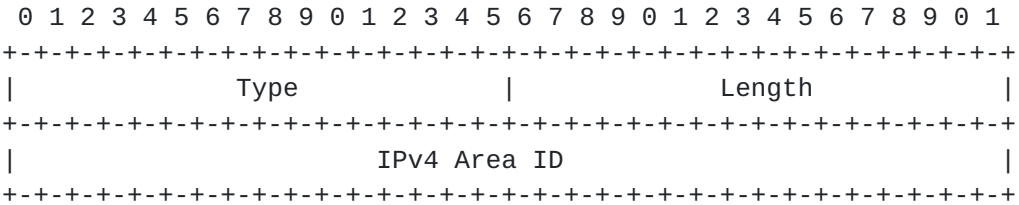

Type To be assigned by IANA (suggested value =1)
Length 4
IPv4 OSPF area ID: The IPv4 identifier of the OSPF area

5.1.3.2. IPv6 area ID DOMAIN sub-TLV

The IPv6 area ID sub-TLV carries an IPv6 OSPF area identifier. It has the following format:

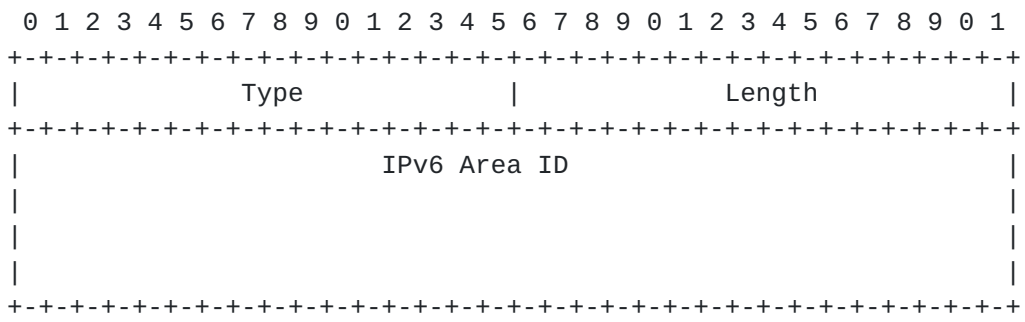

Type To be assigned by IANA (suggested value =2)
Length 16
IPv6 OSPF area ID: The IPv6 identifier of the OSPF area

5.1.3.3. AS Number sub-TLV

The AS Number sub-TLV carries an AS number. It has the following format:

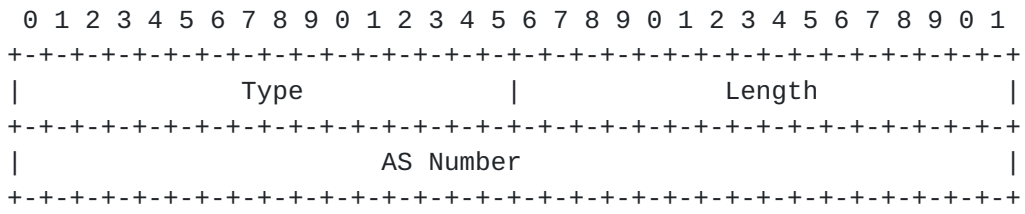

Type To be assigned by IANA (suggested value =3)
Length 4
AS Number: AS number identifying an AS. When coded on two bytes (which is the current defined format as the time of writing this document), the AS Number field MUST have its left two bytes set to 0.

5.1.4. PCE-DEST-DOMAINS sub-TLV

The PCE-DEST-DOMAINS sub-TLV specifies the set of destination domains (areas, AS) toward which a PCE can compute path. It means that the PCE can compute or take part in the computation of inter-domain LSPs whose destination is located in one of these domains. It contains a set of one or more sub-TLVs where each sub-TLV identifies a domain.

The PCE-DEST-DOMAINS sub-TLV has the following format:

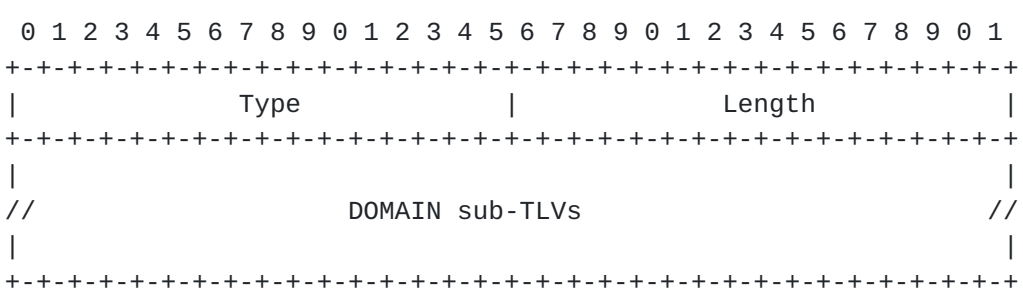

Type	To be defined by IANA (suggested value =3)
Length	Variable
Value	This comprises a set of one or more Area and/or AS DOMAIN sub-TLVs where each DOMAIN sub-TLV identifies a domain toward which a PCE can compute paths.

The PCE-DEST-DOMAINS sub-TLV MUST include at least one DOMAIN sub-TLV.

The PCE-DEST-DOMAINS sub-TLV is optional. It MUST be present only if the R bit is set and the Rd bit is cleared, and/or, if the S bit is set and the Sd bit is cleared.

The PCE-DEST-DOMAINS sub-TLV MUST include at least one area ID sub-TLV, if the R bit of the PATH-SCOPE TLV is set and the Rd bit of the PATH-SCOPE TLV is cleared.
Similarly, the PCE-DEST- DOMAINS sub-TLV MUST include at least one AS number sub-TLV if the S bit of the PATH-SCOPE TLV is set and the Sd bit of the PATH-SCOPE TLV is cleared.

5.1.5. The GENERAL-CAP sub-TLV

The GENERAL-CAP sub-TLV is used to indicate general PCE capabilities.
The GENERAL-CAP sub-TLV is optional. It MAY be carried within the PCED TLV.
The value field of the GENERAL-CAP sub-TLV is made of bit flags, where each bit corresponds to a general PCE capability.

The format of the GENERAL-CAP sub-TLV is as follows:

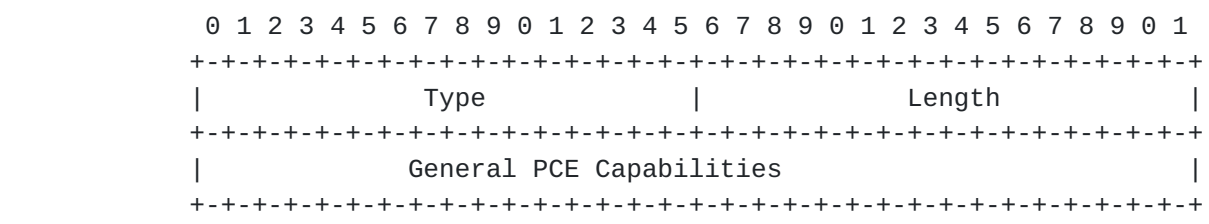

// Optional sub-TLVs //

Type To be assigned by IANA (suggested value =1)
Length It is set to N x 4 octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag.
Value This comprises one or more capability flags. For each 4 octets, the bits are indexed from the most significant to the least significant, where each bit represents one general PCE capability. When the first 32 capabilities are defined, a new capability flag will be used to accommodate the next capability. Optional TLVs may be defined to specify more complex capabilities: there is no optional TLVs currently defined.

IANA is requested to manage the space of general PCE capability bit flags.

The following bits in the first capability flag are to be assigned by IANA:

Bit	Capabilities
0	P bit: Support for Request prioritization.
1	M bit: Support for multiple messages within the same request message.
2-31	Reserved for future assignments by IANA.

5.1.6. The PATH-COMP-CAP sub-TLV

The PATH-COMP-CAP sub-TLV is used to indicate path computation specific capabilities of a PCE.
The PATH-COMP-CAP sub-TLV is optional. It MAY be carried within the PCED TLV.
This is a series of bit flags, where each bit correspond to a path computation capability.

The format of the PATH-COMP-CAP sub-TLV is as follows:

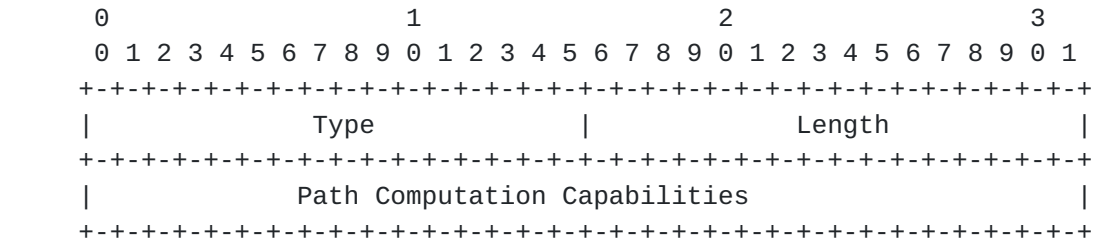

// Optional sub-TLVs //

+-+

Type	To be assigned by IANA (suggested value =1)
Length	It is set to N x 4 octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag.
Value	This comprises one or more capability flags. For each 4 octets, the bits are indexed from the most significant to the least significant, where each bit represents one path computation PCE capability. When the first 32 capabilities are defined, a new capability flag will be used to accommodate the next capability. Optional TLVs may be defined to specify more complex capabilities: there is no optional TLVs currently defined.

IANA is requested to manage the space of PCE path commutation capability bit flags.

The following bits in the first capability flag are to be assigned by IANA:

Bit	Capabilities
0	M bit: Capability to compute P2P paths in MPLS-TE networks
1	G bit: Capability to compute P2P paths in GMPLS networks
2	D bit: Capability to compute link/node/SRLG diverse paths
3	L bit: Capability to compute load-balanced paths
4	S bit: Capability to compute a set of paths in a synchronized Manner
5	O bit: Support for multiple objective functions
6-31	Reserved for future assignments by IANA.

The M, G, D, L, S and O bits are not exclusive.

5.2. Elements of Procedure

The PCED TLV is carried within an OSPF Router information opaque LSA (opaque type of 4, opaque ID of 0) which is defined in [OSPF-CAP].

A router MUST originate a new OSPF router information LSA whenever the content of any of the carried TLVs changes or whenever required by the regular OSPF procedure (LSA refresh (every LSRefreshTime)).

The PCED TLV may be carried within a type 10 or 11 router information LSA depending on the flooding scope of the PCE information.

If the flooding scope is local to an area then it MUST be carried within a type 10 router information LSA.
If the flooding scope is the entire domain then it MUST be carried within type 11 router information LSA.
Note that when the L bit of the PATH-SCOPE TLV is set and the R bit and S bit are cleared, the flooding scope MUST be local, and the PCED TLV MUST be carried within a type 10 Router Information LSA.

PCED TLVs are OPTIONAL. When an OSPF LSA does not contain any PCED TLV, this means that the PCE information of that node is unknown.

Note that a change in PCED information MUST not trigger normal SPF computation.

6. ISIS extensions

6.1. IS-IS PCED TLV format

The IS-IS PCED TLV is made of various non ordered sub-TLVs.

The format of the IS-IS PCED TLV and its sub-TLVs is the same as the TLV format used by the Traffic Engineering Extensions to IS-IS [ISIS-TE]. That is, the TLV is composed of 1 octet for the type, 1 octet specifying the TLV length and a value field.

The IS-IS PCED TLV has the following format:

TYPE: To be assigned by IANA
LENGTH: Variable
VALUE: set of sub-TLVs

Sub-TLVs types are under IANA control.

Currently five sub-TLVs are defined (suggested type values to be assigned by IANA):

Sub-TLV type	Length	Name
1	variable	PCE-ADDRESS sub-TLV
2	2	PATH-SCOPE sub-TLV
3	variable	PCE-DOMAINS sub-TLV
4	variable	PCE-DEST-DOMAINS sub-TLV
5	variable	GENERAL-CAP sub-TLV
6	variable	PATH-COMP-CAP sub-TLV

The sub-TLVs PCE-ADDRESS and PATH-SCOPE MUST always be present within the PCED TLV.

The sub-TLVs PCE-DOMAINS and DEST-DOMAINS, MUST only be present in some particular inter-domain cases.

The GENERAL-CAP and PATH-COMP-CAP are optional and MAY be present in the PCED TLV to facilitate PCE selection.

Any non recognized sub-TLV MUST be silently ignored.

Additional sub-TLVs could be added in the future to advertise additional PCE information.

The PCED TLV is carried within an ISIS CAPABILITY TLV which is defined in [ISIS-CAP].

6.1.1. PCE-ADDRESS sub-TLV

The PCE-ADDRESS sub-TLV specifies the IP address to be used to reach the PCE. This address will typically be a loop-back address that is always reachable, provided the PCE is alive.
The PCE-ADDRESS sub-TLV is mandatory; it MUST be present within the PCED TLV.

The PCE-ADDRESS sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =1)
LENGTH: 4 for IPv4 address and 16 for IPv6 address
VALUE: This comprises one octet indicating the address-type and 4 or 16 octets encoding the IPv4 or IPv6 address to be used to reach the PCE

Address-type:

1	IPv4
2	IPv6

The PCE-ADDRESS sub-TLV MUST appear at least once in the PCED sub-LTV originated by a PCE. It MAY appear multiple times, for instance when the PCE has both an IPv4 and IPv6 address.

6.1.2. The PATH-SCOPE sub-TLV

The PATH-SCOPE sub-TLV indicates the PCE path computation scope; in other words the ability of the PCE to compute or take part into the computation of intra-area, inter-area, inter-AS or inter-layer_TE LSP(s).

The PATH-SCOPE sub-TLV is mandatory; it MUST be present within the PCED TLV. There MUST be exactly one PATH-SCOPE sub-TLV within each PCED TLV.

The PATH-SCOPE sub-TLV contains a set of bit flags indicating the supported path scopes (intra-area, inter-area, inter-AS, inter-layer) and two fields indicating PCE preferences.

The PATH-SCOPE sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =2)
LENGTH: Variable
VALUE: This comprises a one-byte flag of bits where each bit represents a supported path scope, followed by a one-byte preference field indicating PCE preferences.

Here is the structure of the bits flag:

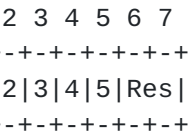

Bit	Path Scope
0	L bit: Can compute intra-area path
1	R bit: Can act as PCE for inter-area TE LSPs computation
2	Rd bit: Can act as a default PCE for inter-area TE LSPs computation
3	S bit: Can act as PCE for inter-AS TE LSPs computation
4	Sd bit: Can act as a default PCE for inter-AS TE LSPs computation
5	Y bit: Can compute or take part into the computation of paths across layers

Here is the structure of the preference field

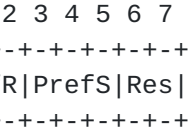

PrefR field: PCE s preference for inter-area TE LSPs computation

PrefS field: PCE s preference for inter-AS TE LSPs computation

The bits L, R, S and Y bits are set when the PCE can act as a PCE for intra-area, inter-area, inter-AS and inter-layer TE LSPs computation respectively. These bits are non exclusive.

When set the Rd bit indicates that the PCE can act as a default PCE for inter-area TE LSPs computation, it means that it can compute path for any destination area. Similarly, when set the Sd bit indicates that the PCE can act as a default PCE for inter-AS TE LSPs computation, it means that it can compute path for any destination AS.

When the Rd bit is set the PCE-DEST-DOMAINS TLV (see 6.1.4) does not comprise any Area ID DOMAIN sub-TLV. When the Sd bit is set the PCE-DEST-DOMAINS TLV does not comprise any AS DOMAIN sub-TLV.

The PrefR and PrefS fields are 3-bit long and allow the PCE to specify a preference where 7 reflects the highest preference. For the sake of illustration, consider the situation where N ABRs act as PCEs

for inter-area TE LSPs computation. An operator may decide to configure a preference to each PCE that could be used to load balance the path computation load with respect to their respective CPU capacity. The algorithms used by a PCC to load balance its path computation requests according to such PCE s preference is out of the scope of this document

6.1.3. PCE-DOMAINS sub-TLV

The PCE-DOMAINS sub-TLV specifies the set of domains (areas, AS) where the PCE has topology visibility and can compute paths. It contains a set of one or more sub-TLVs where each sub-TLV identifies a domain.

The PCE-DOMAINS sub-TLV MUST only be present when PCE domains cannot be inferred by other IGP information, for instance when the PCE is inter-domain capable (i.e. when the R bit or S bit is set) and the flooding scope is the entire routing domain.

The PCE-DOMAINS sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =2)
LENGTH: Variable
VALUE: This comprises a set of one or more DOMAIN sub-TLVs where each DOMAIN sub-TLV identifies a domain where the PCE has topology visibility and can compute paths

DOMAIN Sub-TLVs types are under IANA control.

Currently two DOMAIN sub-TLVs are defined (suggested type values to be assigned by IANA):

Sub-TLV type	Length	Name
1	variable	Area ID sub-TLV
2	variable	AS number sub-TLV

At least one DOMAIN sub-TLV MUST be present in the PCE-DOMAINS sub-TLV.

6.1.3.1. Area ID DOMAIN sub-TLV

This sub-TLV carries an ISIS area ID. It has the following format

TYPE: To be assigned by IANA (Suggested value =1)
LENGTH: Variable
VALUE: This comprises a variable length ISIS area ID. This is the combination of an Initial Domain Part (IDP) and High Order part of the Domain Specific part (HO-DPS)

6.1.3.2. AS Number DOMAIN sub-TLV

The AS Number sub-TLV carries an AS number. It has the following format:

TYPE: To be assigned by IANA (Suggested value =2)
LENGTH: 4
VALUE: AS number identifying an AS. When coded on two
bytes (which is the current defined format as the
time of writing this document), the AS Number field
MUST have its left two bytes set to 0.

6.1.4. PCE-DEST-DOMAINS sub-TLV

The PCE-DEST-DOMAINS sub-TLV specifies the set of destination domains (areas, AS) toward which a PCE can compute path. It means that the PCE can compute or take part in the computation of inter-domain LSPs whose destination is located in one of these domains. It contains a set of one or more DOMAIN sub-TLVs where each DOMAIN sub-TLV identifies a domain.

The PCE-DEST-DOMAINS sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =2)
LENGTH: Variable
VALUE: This comprises a set of one or more Area or/and AS DOMAIN sub-TLVs where each sub-TLV identifies a destination domain toward which a PCE can compute path.

At least one DOMAIN sub-TLV MUST be present in the PCE-DEST-DOMAINS sub-TLV.

The PCE-DEST-DOMAINS sub-TLV is optional. It MUST be present only if the R bit is set and the Rd bit is cleared, and/or, if the S bit is set and the Sd bit is cleared.

The PCE-DEST-DOMAINS sub-TLV MUST include at least one area ID sub-TLV, if the R bit of the PATH-SCOPE TLV is set and the Rd bit of the PATH-SCOPE TLV is cleared. Similarly, the PCE-DEST-DOMAINS sub-TLV MUST include at least one AS number sub-TLV if the S bit of the PATH-SCOPE TLV is set and the Sd bit of the PATH-SCOPE TLV is cleared.

6.1.5. GENERAL-CAP sub-TLV

The GENERAL-CAP sub-TLV is used to indicate general PCE capabilities. The GENERAL-CAP sub-TLV is optional. It MAY be carried within the PCED TLV.
This is a series of bits flags, where each bit corresponds to a general PCE capability.

The GENERAL-CAP sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =4)

LENGTH: It is set to N. N starts from 1 and can be increased when there is a need. Each octet is referred to as a capability flag.

VALUE: This comprises one or more general PCE capability flags.

The following bits in the first capability flag are to be assigned by IANA:

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

P bit: Support for request prioritization.
M bit: Support for multiple messages within the same request message.

Reserved bits are for future assignment by IANA.

6.1.6. The PATH-COMP-CAP sub-TLV

The PATH-COMP-CAP sub-TLV is used to indicate path computation specific capabilities of a PCE.
The PATH-COMP-CAP sub-TLV is optional. It MAY be carried within the PCED TLV.
This is a series of bit flags, where each bit correspond to a path computation capability.

The PATH-COMP-CAP sub-TLV has the following format:

TYPE: To be assigned by IANA (suggested value = 5)
LENGTH: It is set to N. N starts from 1 and can be increased when there is a need. Each octet is referred to as a capability flag.
VALUE: This comprises one or more Path Computation specific PCE capability flags.

The following bits in the first capability flag are to be assigned by IANA.

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|M|G|D|L|S|0|Res|
+-+-+-+-+-+-+-+-+
```

M bit: Capability to compute P2P paths in MPLS-TE networks
G bit: Capability to compute P2P paths in GMPLS networks
D bit: Capability to compute link/node/SRLG diverse paths
L bit: Capability to compute load-balanced paths
S bit: Capability to compute a set of paths in a synchronized Manner
O bit: Support for multiple objective functions

Reserved bits are for future assignment by IANA.

6.2. Elements of procedure

The PCED TLV is carried within an IS-IS CAPABILITY TLV defined in [IS-IS-CAP].

An IS-IS router MUST originate a new IS-IS LSP whenever the content of any of the PCED TLV changes or whenever required by the regular IS-IS procedure (LSP refresh).

When the scope of the PCED TLV is area local it MUST be carried within an ISIS CAPABILITY TLV having the S bit cleared.
When the scope of the PCED TLV is the entire domain, the PCED TLV MUST be carried within an ISIS CAPABILITY TLV having the S bit set.
Note that when the L bit of the PATH-SCOPE sub-TLV is set and the R and S bits are cleared the flooding scope MUST be local.

PCED TLVs are OPTIONAL. When an IS-IS LSP does not contain any PCED TLV, this means that the PCE information of that node is unknown.

Note that a change in PCED information MUST not trigger normal SPF computation.

7. Backward compatibility

The PCED TLVs defined in this document do not introduce any interoperability issue. For OSPF, a router not supporting the PCED TLV SHOULD just silently ignore the TLV as specified in RFC2370. For IS-IS a router not supporting the PCED TLV SHOULD just silently ignore the TLV.

8. Security Considerations

No new security issues are raised in this document.

9. IANA considerations

9.1. OSPF TLVs

IANA will assign a new codepoint for the OSPF PCED TLV defined in this document and carried within the Router Information LSA.

Five sub-TLVs types are defined for this TLV and should be assigned by IANA:

- -PCE-ADDRESS sub-TLV (suggested value = 1)
- -PATH-SCOPE sub-TLV (suggested value = 2)
- -PCE-DOMAINS sub-TLV (suggested value = 3)
- -PCE-DEST-DOMAINS sub-TLV (suggested value =4)
- -GENERAL-CAP sub-TLV (suggested value = 5)

-PATH-COMP-CAP sub-TLV (suggested value = 6)

Three sub-TLVs types are defined for the PCE-DOMAINS and PCE-DEST-DOMAINS TLVs and should be assigned by IANA:

- -IPv4 area ID sub-TLV (suggested value = 1)
- -IPv6 area ID sub-TLV (suggested value = 2)
- -AS number sub-TLV (suggested value = 3)

9.2. ISIS TLVs

IANA will assign a new codepoint for the PCED TLV defined in this document and carried within the ISIS CAPABILITY TLV.

Five sub-TLVs types are defined for the PCED TLV and should be assigned by IANA:

- -PCE-ADDRESS sub-TLV (suggested value = 1)
- -PATH-SCOPE sub-TLV (suggested value = 2)
- -PCE-DOMAINS sub-TLV (suggested value = 3)
- -GENERAL-CAP sub-TLV (suggested value = 4)
- -PATH-COMP-CAP sub-TLV (suggested value = 5)

Two sub-TLVs types are defined for the PCE-DOMAINS and PCE-DEST-DOMAINS TLVs and should be assigned by IANA:

- -Area ID sub-TLV (suggested value = 1)
- -AS number sub-TLV (suggested value = 2)

9.3. Capability bits

IANA is requested to manage the space of general and path computation specific PCE capability bits flags, numbering them in the usual IETF notation starting at zero and continuing at least through 31.
New bit numbers may be allocated only by an IETF Consensus action.
Each bit should be tracked with the following qualities:

- Bit number
- Defining RFC
- Name of bit

Currently two bits are defined in the first general PCE capability flag. Here are the suggested values:

- -0: Support for Request prioritization.
- -1: Support for multiple messages within the same request message

Currently six bits are defined in the first path computation specific PCE capability flag. Here are the suggested values:

- -0: Capability to compute P2P paths in MPLS-TE networks
- -1: Capability to compute P2P paths in GMPLS networks
- -2: Capability to compute link/node/SRLG diverse paths
- -3: Capability to compute load-balanced paths

-4: Capability to compute a set of paths in a
    synchronized Manner
-5: Support for multiple objective functions

10. Security Considerations

No new security issues are raised in this document.

11. References

11.1. Normative references

[OSPF-v2] Moy, J., "OSPF Version 2", RFC 2328, April 1998.

[IS-IS] "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol " ISO 10589.

[IS-IS-IP] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.

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

[IS-IS-TE] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", RFC 3784, June 2004.

[OSPF-CAP] Lindem, A., Shen, N., Aggarwal, R., Shaffer, S., Vasseur, J.P., "Extensions to OSPF for advertising Optional Router Capabilities", draft-ietf-ospf-cap, work in progress.

[IS-IS-CAP] Vasseur, J.P. et al., "IS-IS extensions for advertising router information", draft-ietf-isis-caps, work in progress.

[PCE-ARCH] Farrel, A., Vasseur, J.P., Ash, J., "Path Computation Element (PCE) Architecture", draft-ietf-pce-architecture, work in progress.

[PCE-DISCO-REQ] Le Roux, J.L., et al. "Requirements for PCE discovery", draft-ietf-pce-discovery-reqs, work in progress

11.2. Informative references

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

[RFC3667] Bradner, S., "IETF Rights in Contributions", BCP 78, RFC 3667, February 2004.

[RFC3668] Bradner, S., "Intellectual Property Rights in IETF Technology", BCP 79, RFC 3668, February 2004.

[PCE-COM-REQ] Ash, J., Le Roux, J.L., " PCE Communication Protocol Generic Requirements", draft-ietf-pce-comm-protocol-gen-reqs, work in progress.

12. Authors' Addresses:

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@francetelecom.com

Jean-Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MA - 01719
USA
Email: jpv@cisco.com

Yuichi Ikejiri
NTT Communications Corporation
1-1-6, Uchisaiwai-cho, Chiyoda-ku
Tokyo 100-8019
JAPAN
Email: y.ikejiri@ntt.com

13. Intellectual Property Statement

rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

Copyright Statement