

Network Working Group
Internet Draft
Intended Status: Standard Track
Expires: April 2008

J.L. Le Roux (Editor)
France Telecom

J.P. Vasseur (Editor)
Cisco System Inc.

Yuichi Ikejiri
NTT Communications

Raymond Zhang
BT Infonet

October 2007

OSPF Protocol Extensions for Path Computation Element (PCE) Discovery

[draft-ietf-pce-disco-proto-ospf-08.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The IETF Trust (2007). All rights reserved.

Internet Draft [draft-ietf-pce-disco-proto-ospf-08.txt](#)

October 2007

Abstract

There are various circumstances where it is highly desirable for a Path Computation Client (PCC) to be able to dynamically and automatically discover a set of Path Computation Elements (PCEs), along with information that can be used by the PCC for PCE selection. When the PCE is a Label Switching Router (LSR) participating in the Interior Gateway Protocol (IGP), or even a server participating passively in the IGP, a simple and efficient way to announce PCEs consists of using IGP flooding. For that purpose, this document defines extensions to the Open Shortest Path First (OSPF) routing protocol for the advertisement of PCE Discovery information within an OSPF area or within the entire OSPF routing domain.

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 [[RFC2119](#)].

Table of Contents

1.	Terminology.....	3
2.	Introduction.....	4
3.	Overview.....	5
3.1.	PCE Discovery Information.....	5
3.2.	Flooding Scope.....	5
4.	The OSPF PCED TLV.....	6
4.1.	PCE-ADDRESS Sub-TLV.....	7
4.2.	PATH-SCOPE Sub-TLV.....	8
4.3.	PCE-DOMAIN Sub-TLV.....	10
4.4.	NEIG-PCE-DOMAIN Sub-TLV.....	11
4.5.	PCE-CAP-FLAGS Sub-TLV.....	11
5.	Elements of Procedure.....	13
6.	Backward Compatibility.....	13
7.	IANA Considerations.....	14
7.1.	OSPF TLV.....	14
7.2.	PCE Capability Flags registry.....	14
8.	Security Considerations.....	15
9.	Manageability Considerations.....	15
9.1.	Control of Policy and Functions.....	15
9.2.	Information and Data Model.....	15

9.3.	Liveness Detection and Monitoring.....	15
9.4.	Verify Correct Operations.....	16
9.5.	Requirements on Other Protocols and Functional Components.....	16
9.6.	Impact on Network Operations.....	16
10.	Acknowledgments.....	16
11.	References.....	16
11.1.	Normative References.....	16
11.2.	Informative References.....	17
12.	Editor's Addresses.....	17

13.	Contributors' Addresses.....	18
14.	Intellectual Property Statement.....	18

[1.](#) Terminology

ABR: OSPF Area Border Router.

AS: Autonomous System.

IGP: Interior Gateway Protocol. Either of the two routing protocols Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (ISIS).

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

Intra-AS TE LSP: A TE LSP whose path does not cross an AS boundary.

Inter-area TE LSP: A TE LSP whose path transits two or more IGP areas. That is a TE LSP that crosses at least one IGP area boundary.

Inter-AS TE LSP: A TE LSP whose path transits two or more ASes or sub-ASes (BGP confederations). That is a TE LSP that crosses at least one AS boundary.

LSA: Link State Advertisement.

LSR: Label Switching Router.

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

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.

PCE-Domain: In a PCE context this refers to any collection of network elements within a common sphere of address management or path computational responsibility (referred to as a "domain" in [[RFC4655](#)]). Examples of PCE-Domains include IGP areas and ASes. This should be distinguished from an OSPF routing domain.

PCEP: Path Computation Element Protocol.

TE LSP: Traffic Engineered Label Switched Path.

[2](#). Introduction

[RFC4655] describes the motivations and architecture for a Path Computation Element (PCE)-based path computation model for Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineered Label Switched Paths (TE LSPs). The model allows for the separation of the PCE from a Path Computation Client (PCC) (also referred to as a non co-located PCE) and allows for cooperation between PCEs (where one PCE acts as a PCC to make requests of the other PCE). This relies on a communication protocol between PCC and PCE, and also between PCEs. The requirements for such a communication protocol can be found in [[RFC4657](#)], and the communication protocol is defined in [[PCEP](#)].

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

A network may contain a large number of PCEs, each with potentially distinct capabilities. In such a context it is highly desirable to have a mechanism for automatic and dynamic PCE discovery that allows PCCs to automatically discover a set of PCEs along with additional information about each PCE that may be used by a PCC to perform PCE selection. Additionally, it is valuable for a PCC to dynamically

detect new PCEs, failed PCEs, or any modification to the PCE information. Detailed requirements for such a PCE discovery mechanism are provided in [[RFC4674](#)].

Note that the PCE selection algorithm applied by a PCC is out of the scope of this document.

When PCCs are LSRs participating in the IGP (OSPF or IS-IS), and PCEs are either LSRs or servers also participating in the IGP, an effective mechanism for PCE discovery within an IGP routing domain consists of utilizing IGP advertisements.

This document defines extensions to OSPFv2 [[RFC2328](#)] and OSPFv3 [[RFC2740](#)] to allow a PCE in an OSPF routing domain to advertise its location along with some information useful to a PCC for PCE selection so as to satisfy dynamic PCE discovery requirements set forth in [[RFC4674](#)].

Generic capability advertisement mechanisms for OSPF are defined in [[OSPF-CAP](#)]. These allow a router to advertise its capabilities within an OSPF area or an entire OSPF routing domain. This document leverages this generic capability advertisement mechanism to fully satisfy the dynamic PCE discovery requirements.

This document defines a new TLV (named the PCE Discovery (PCED) TLV) to be carried within the OSPF Router Information LSA ([[OSPF-CAP](#)]).

The PCE information advertised is detailed in [Section 3](#). Protocol extensions and procedures are defined in [Sections 4](#) and [5](#).

The OSPF extensions defined in this document allow for PCE discovery within an OSPF routing domain. Solutions for PCE discovery across AS boundaries are beyond the scope of this document, and for further study.

[3](#). Overview

[3.1](#). PCE Discovery Information

The PCE discovery information is composed of:

- The PCE location: an IPv4 and/or IPv6 address that is used to reach the PCE. It is RECOMMENDED to use an address that is always reachable if there is any connectivity to the PCE;
- The PCE path computation scope (i.e., intra-area, inter-area, inter-AS, or inter-layer);
- The set of one or more PCE-Domain(s) into which the PCE has visibility and for which the PCE can compute paths;
- The set of zero, one or more neighbor PCE-Domain(s) toward which the PCE can compute paths;
- A set of communication capabilities (e.g., support for request prioritization) and path computation-specific capabilities (e.g., supported constraints).

PCE discovery information is by nature fairly static and does not change with PCE activity. Changes in PCE discovery information may occur as a result of PCE configuration updates, PCE deployment/activation, PCE deactivation/suppression, or PCE failure. Hence, this information is not expected to change frequently.

[3.2. Flooding Scope](#)

The flooding scope for PCE information advertised through OSPF can be limited to one or more OSPF areas the PCE belongs to, or can be extended across the entire OSPF routing domain.

Note that some PCEs may belong to multiple areas, in which case the flooding scope may comprise these areas. This could be the case for an ABR, for instance, advertising its PCE information within the backbone area and/or a subset of its attached IGP area(s).

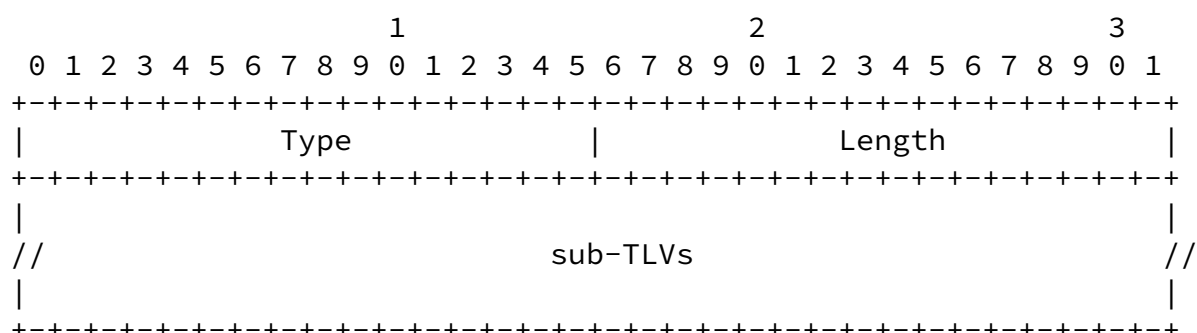
[4. The OSPF PCED TLV](#)

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

The format of the OSPF PCED TLV and its sub-TLVs is identical to the TLV format used by the Traffic Engineering Extensions to OSPF [RFC3630]. 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 four-octet aligned. Unrecognized types are ignored.

The OSPF PCED TLV has the following format:



Type	To be defined by IANA (suggested value=5)
Length	Variable
Value	This comprises one or more sub-TLVs

Five sub-TLVs are defined:

Sub-TLV type	Length	Name
1	variable	PCE-ADDRESS sub-TLV
2	4	PATH-SCOPE sub-TLV
3	4	PCE-DOMAIN sub-TLV
4	4	NEIG-PCE-DOMAIN sub-TLV
5	variable	PCE-CAP-FLAGS sub-TLV

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

The PCE-DOMAIN and NEIG-PCE-DOMAIN sub-TLVs are optional. They MAY be present in the PCED TLV to facilitate selection of inter-domain PCEs.

The PCE-CAP-FLAGS sub-TLV is optional and MAY be present in the PCED TLV to facilitate the PCE selection process.

Malformed PCED TLVs or sub-TLVs not explicitly described in this document MUST cause the LSA to be treated as malformed according to the normal procedures of OSPF.

Any unrecognized sub-TLV MUST be silently ignored.

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

No additional sub-TLVs will be added to the PCED TLV in the future. If a future application requires the advertisement of additional PCE information in OSPF, this will not be carried in the Router Information LSA.

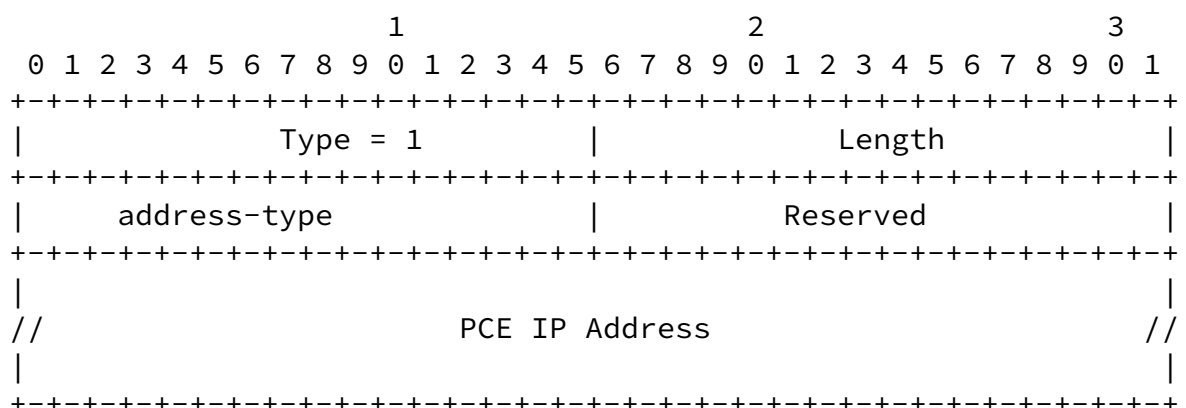
The following sub-sections describe the sub-TLVs which may be carried within the PCED sub-TLV.

[4.1](#). PCE-ADDRESS Sub-TLV

The PCE-ADDRESS sub-TLV specifies an IP address that can be used to reach the PCE. It is RECOMMENDED to make use of an address that is always reachable, provided that the PCE is alive and reachable.

The PCE-ADDRESS sub-TLV is mandatory; it MUST be present within the PCED TLV. It MAY appear twice, when the PCE has both an IPv4 and IPv6 address. It MUST NOT appear more than once for the same address type. If it appears more than once, only the first occurrence is processed and any others MUST be ignored.

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



PCE-ADDRESS sub-TLV format

Type 1
Length 8 (IPv4) or 20 (IPv6)

Address-type:

Internet Draft [draft-ietf-pce-disco-protocol-08.txt](https://datatracker.ietf.org/doc/draft-ietf-pce-disco-protocol-08.txt) October 2007

- 1 IPv4
- 2 IPv6

Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

PCE IP Address: The IP address to be used to reach the PCE.

4.2. PATH-SCOPE Sub-TLV

The PATH-SCOPE sub-TLV indicates the PCE path computation scope, which refers to the PCE's ability to compute or take part in the computation of paths for intra-area, inter-area, inter-AS, or inter-layer_TE LSPs.

The PATH-SCOPE sub-TLV is mandatory; it MUST be present within the PCED TLV. There MUST be exactly one instance of the PATH-SCOPE sub-TLV within each PCED TLV. If it appears more than once, only the first occurrence is processed and any others MUST be ignored.

The PATH-SCOPE sub-TLV contains a set of bit-flags indicating the supported path scopes, and four fields indicating PCE preferences.

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

```

          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 = 2             |               Length             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|1|2|3|4|5|   Reserved             |PrefL|PrefR|PrefS|PrefY| Res   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type 2
Length 4
Value This comprises a 2-octet flag field where each bit represents a supported path scope, as well as four

preference fields used to specify PCE preferences.

The following bits are defined:

Bit	Path Scope
0	L bit: Can compute intra-area paths
1	R bit: Can act as PCE for inter-area TE LSP computation
2	Rd bit: Can act as a default PCE for inter-area TE LSP computation
3	S bit: Can act as PCE for inter-AS TE LSP computation
4	Sd bit: Can act as a default PCE for inter-AS TE LSP

Le Roux, Vasseur et al. OSPF extensions for PCE Discovery [Page 8]

Internet Draft [draft-ietf-pce-disco-proto-ospf-08.txt](#) October 2007

	computation
5	Y bit: Can compute or take part into the computation of paths across layers.

PrefL field: PCE's preference for intra-area TE LSPs computation.

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

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

PrefY field: PCE's preference for inter-layer TE LSPs computation.

Res: Reserved for future use.

The L, R, S, and Y bits are set when the PCE can act as a PCE for intra-area, inter-area, inter-AS, or 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 (that is, the PCE can compute a path toward any neighbor area). Similarly, when set, the Sd bit indicates that the PCE can act as a default PCE for inter-AS TE LSP computation (the PCE can compute a path toward any neighbor AS).

When the Rd and Sd bit are set the PCED TLV MUST NOT contain a NEIG-PCE-DOMAIN sub-TLV (see [Section 4.1.4](#)).

When the R bit is clear, the Rd bit SHOULD be clear on transmission and MUST be ignore on receipt. When the S bit is clear, the Sd bit SHOULD be clear on transmission and MUST be ignored on receipt.

The PrefL, PrefR, PrefS, and PrefY fields are each three bits long and allow the PCE to specify a preference for each computation scope, where 7 reflects the highest preference. Such preferences can be used for weighted load balancing of path computation requests. An operator may decide to configure a preference for each computation scope at each PCE so as to balance the path computation load among them. The algorithms used by a PCC to load balance its path computation requests according to such PCE preferences is out of the scope of this document and is a matter for local or network-wide policy. The same or different preferences may be used for each scope. For instance, an operator that wants a PCE capable of both inter-area and inter-AS computation to be preferred for use for inter-AS computations may configure PrefS higher than PrefR.

When the L, R, S, or Y bits are cleared, the PrefL, PrefR, PrefS, and PrefY fields SHOULD respectively be set to 0 on transmission and MUST be ignored on receipt.

Both reserved fields SHOULD be set to zero on transmission and MUST be ignored on receipt.

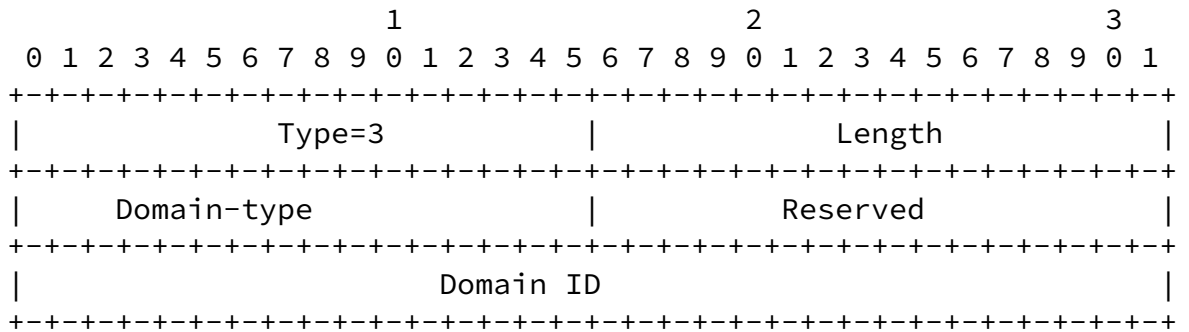
[4.3](#). PCE-DOMAIN Sub-TLV

The PCE-DOMAIN sub-TLV specifies a PCE-Domain (area or AS) where the PCE has topology visibility and through which the PCE can compute paths.

The PCE-DOMAIN sub-TLV SHOULD be present when PCE-Domains for which the PCE can operate 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 (see [Section 5](#) for a discussion of how the flooding scope is set and interpreted).

A PCED TLV may include multiple PCE-DOMAIN sub-TLVs when the PCE has visibility into multiple PCE-Domains.

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



PCE-DOMAIN sub-TLV format

Type 3
Length 8

Two domain-type values are defined:

- 1 OSPF Area ID
- 2 AS Number

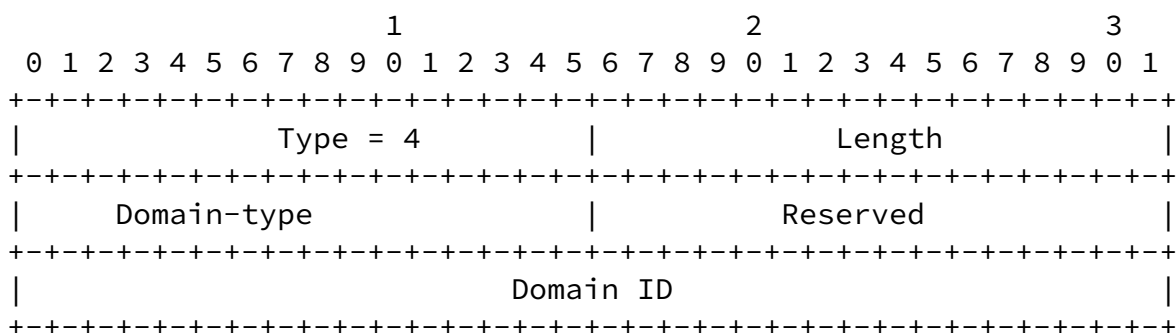
Domain ID: With the domain-type set to 1, this indicates the 32 bit Area ID of an area where the PCE has visibility and can compute paths. With domain-type set to 2, this indicates an AS number of an AS where the PCE has visibility and can compute paths. When the AS number is coded in two octets, the AS Number field MUST have its first two octets set to 0.

[4.4.](#) NEIG-PCE-DOMAIN Sub-TLV

The NEIG-PCE-DOMAIN sub-TLV specifies a neighbor PCE-domain (area or AS) toward which a PCE can compute paths. It means that the PCE can take part in the computation of inter-domain TE LSPs with paths that transit this neighbor PCE-domain.

A PCED sub-TLV may include several NEIG-PCE-DOMAIN sub-TLVs when the PCE can compute paths towards several neighbour PCE-domains.

The NEIG-PCE-DOMAIN sub-TLV has the same format as the PCE-DOMAIN sub-TLV:



NEIG-PCE-DOMAIN sub-TLV format

Type 4
Length 8

Two domain-type values are defined:

- 1 OSPF Area ID
- 2 AS Number

Domain ID: With the domain-type set to 1, this indicates the 32 bit Area ID of a neighbour area toward which the PCE can compute paths. With domain-type set to 2, this indicates the AS number of a neighbor AS toward which the PCE can compute paths. When the AS number is coded in two octets, the AS Number field MUST have its first two octets set to 0.


The NEIG-PCE-DOMAIN sub-TLV MUST be present at least once with domain-type set to 1 if the R bit is set and the Rd bit is cleared, and MUST be present at least once with domain-type set to 2 if the S bit is set and the Sd bit is cleared.

[4.5.](#) PCE-CAP-FLAGS Sub-TLV

The PCE-CAP-FLAGS sub-TLV is an optional sub-TLV used to indicate PCE capabilities. It MAY be present within the PCED TLV. It MUST NOT be present more than once. If it appears more than once, only the first occurrence is processed and any others MUST be ignored.

The value field of the PCE-CAP-FLAGS sub-TLV is made up of an array of units of 32 bit-flags numbered from the most significant bit as bit zero, where each bit represents one PCE capability.

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



Type 5
 Length Multiple of 4 octets
 Value This contains an array of units of 32 bit flags
 numbered from the most significant as bit zero, where
 each bit represents one PCE capability.

IANA is requested to manage the space of the PCE Capability Flags

The following bits are to be assigned by IANA:

Bit	Capabilities
0	Path computation with GMPLS link constraints
1	Bidirectional path computation
2	Diverse path computation
3	Load-balanced path computation
4	Synchronized path computation
5	Support for multiple objective functions
6	Support for additive path constraints (max hop count, etc.)
7	Support for request prioritization
8	Support for multiple requests per message

9-31 Reserved for future assignments by IANA.

These capabilities are defined in [[RFC4657](#)].

Reserved bits SHOULD be set to zero on transmission and MUST be ignored on receipt.

5. Elements of Procedure

The PCED TLV is advertised within OSPFv2 Router Information LSAs (Opaque type of 4 and Opaque ID of 0) or OSPFv3 Router Information LSAs (function code of 12) which are defined in [[OSPF-CAP](#)]. As such, elements of procedure are inherited from those defined in [[OSPF-CAP](#)].

In OSPFv2 the flooding scope is controlled by the opaque LSA type (as defined in [[RFC2370](#)]) and in OSPFv3 by the S1/S2 bits (as defined in [[RFC2740](#)]). If the flooding scope is local to an area then the PCED TLV MUST be carried within an OSPFv2 type 10 router information LSA or an OSPFv3 Router Information LSA with the S1 bit set and the S2 bit clear. If the flooding scope is the entire IGP domain then the PCED TLV MUST be carried within an OSPFv2 type 11 Router Information LSA or OSPFv3 Router Information LSA with the S1 bit clear and the S2 bit set. When only the L bit of the PATH-SCOPE sub-TLV is set, the flooding scope MUST be area local.

When the PCE function is deactivated, the OSPF speaker advertising this PCE MUST originate a new Router Information LSA that no longer includes the corresponding PCED TLV, provided there are other TLVs in the LSA. If there are no other TLVs in the LSA, it MUST either send an empty Router Information LSA or purge it by prematurely aging it.

The PCE address (i.e., the address indicated within the PCE ADDRESS TLV) SHOULD be reachable via some prefixes advertised by OSPF. This allows the detection of a PCE failure to be sped up. When the PCE address is no longer reachable, the PCE node has failed, has been torn down, or there is no longer IP connectivity to the PCE.

A change in information in the PCED TLV MUST NOT trigger any SPF computation at a receiving router.

The way PCEs determine the information they advertise is out of the scope of this document. Some information may be configured on the PCE (e.g., address, preferences, scope) and other information may be automatically determined by the PCE (e.g., areas of visibility).

6. Backward Compatibility

The PCED TLV defined in this document does not introduce any interoperability issues.

A router not supporting the PCED TLV will just silently ignore the TLV as specified in [[OSPF-CAP](#)].

Internet Draft [draft-ietf-pce-disco-protocol-08.txt](https://datatracker.ietf.org/doc/draft-ietf-pce-disco-protocol-08.txt) October 2007

[7.](#) IANA Considerations

[7.1.](#) OSPF TLV

IANA has defined a registry for TLVs carried in the Router Information LSA defined in [[OSPF-CAP](#)]. IANA is requested to assign a new TLV code-point for the PCED TLV carried within the Router Information LSA.

Value	TLV Name	Reference
-----	-----	-----
5	PCED	(this document)

[7.2.](#) PCE Capability Flags registry

This document provides new capability bit flags, which are present in the PCE-CAP-FLAGS TLV referenced in [section 4.1.5](#).

The IANA is requested to create a new top-level OSPF registry, the "PCE Capability Flags" registry, and to manage the space of PCE capability bit flags numbering them in the usual IETF notation starting at zero and continuing at least through 31, with the most significant bit as bit zero.

New bit numbers may be allocated only by an IETF Consensus action.

Each bit should be tracked with the following qualities:

- Bit number
- Capability Description
- Defining RFC

Several bits are defined in this document. Here are the suggested values:

Bit	Capability Description
-----	------------------------

- 0 Path computation with GMPLS link constraints
- 1 Bidirectional path computation
- 2 Diverse path computation
- 3 Load-balanced path computation
- 4 Synchronized paths computation
- 5 Support for multiple objective functions
- 6 Support for additive path constraints
(max hop count, etc.)
- 7 Support for request prioritization
- 8 Support for multiple requests per message

8. Security Considerations

This document defines OSPF extensions for PCE discovery within an administrative domain. Hence the security of the PCE discovery relies on the security of OSPF.

Mechanisms defined to ensure authenticity and integrity of OSPF LSAs [[RFC2154](#)], and their TLVs, can be used to secure the PCE Discovery information as well.

OSPF provides no encryption mechanism for protecting the privacy of LSAs, and in particular the privacy of the PCE discovery information.

9. Manageability Considerations

Manageability considerations for PCE Discovery are addressed in [Section 4.10 of \[RFC4674\]](#).

9.1. Control of Policy and Functions

Requirements for the configuration of PCE discovery parameters on PCCs and PCEs are discussed in [Section 4.10.1 of \[RFC4674\]](#).

In particular, a PCE implementation SHOULD allow the following parameters to be configured on the PCE:

- The PCE IPv4/IPv6 address(es) (see [Section 4.1](#))
- The PCE Scope, including the inter-domain functions (inter-area, inter-AS, inter-layer), the preferences, and whether the

- PCE can act as default PCE (see [Section 4.2](#))
- The PCE domains (see [Section 4.3](#))
- The neighbour PCE domains (see [Section 4.4](#))
- The PCE capabilities (see [Section 4.5](#))

[9.2.](#) Information and Data Model

A MIB module for PCE Discovery is defined in [[PCED-MIB](#)].

[9.3.](#) Liveness Detection and Monitoring

PCE Discovery Protocol liveness detection relies upon OSPF liveness detection. OSPF already includes a liveness detection mechanism (Hello protocol), and PCE discovery does not require additional capabilities.

Procedures defined in [Section 5](#) allow a PCC to detect when a PCE has been deactivated, or is no longer reachable.

[9.4.](#) Verify Correct Operations

The correlation of information advertised against information received can be achieved by comparing the information in the PCED TLV received by the PCC with that stored at the PCE using the PCED MIB [[PCED-MIB](#)]. The number of dropped, corrupt, and rejected information elements are available through the PCED MIB.

[9.5.](#) Requirements on Other Protocols and Functional Components

The OSPF extensions defined in this document do not imply any requirement on other protocols.

[9.6.](#) Impact on Network Operations

Frequent changes in PCE information advertised in the PCED TLV, may have a significant impact on OSPF and might destabilize the operation of the network by causing the PCCs to swap between PCEs.

As discussed in [Section 4.10.4 of \[RFC4674\]](#), it MUST be possible to

apply at least the following controls:

- Configurable limit on the rate of announcement of changed parameters at a PCE.
- Control of the impact on PCCs such as through rate-limiting the processing of PCED TLVs.
- Configurable control of triggers that cause a PCC to swap to another PCE.

10. Acknowledgments

We would like to thank Lucy Wong, Adrian Farrel, Les Ginsberg, Mike Shand, and Lou Berger for their useful comments and suggestions.

We would also like to thank Dave Ward, Lars Eggert, Sam Hartman, and Tim Polk for their comments during the final stages of publication.

11. References

11.1. 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", [RFC 2328](#), April 1998.

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", [RFC 2740](#), December 1999.

[RFC2370] Coltun, R., "The OSPF Opaque LSA Option", [RFC 2370](#), July 1998.

Le Roux, Vasseur et al. OSPF extensions for PCE Discovery [Page 16]

Internet Draft [draft-ietf-pce-disco-proto-ospf-08.txt](#) October 2007

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

[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.

[RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", [RFC 2154](#), June 1997.

11.2. Informative References

[RFC4657] Ash, J., Le Roux, J.L., "PCE Communication Protocol Generic Requirements", [RFC4657](#), September 2006.

[PCEP] Vasseur, Le Roux, et al., "Path Computation Element (PCE) communication Protocol (PCEP) - Version 1", [draft-ietf-pce-pcep](#), work in progress.

[PCED-MIB] Stephan, E., "Definitions of Managed Objects for Path Computation Element Discovery", [draft-ietf-pce-disc-mib](#), work in progress.

[PCED-ISIS] Le Roux, Vasseur, et al. "IS-IS protocol extensions for Path Computation Element (PCE) Discovery", [draft-ietf-pce-disco-proto-isis](#), work in progress.

[RFC4655] Farrel, A., Vasseur, J.P., Ash, J., "Path Computation Element (PCE)-based Architecture", [RFC4655](#), August 2006.

[RFC4674] Le Roux, J.L., et al. "Requirements for PCE discovery", [RFC4674](#), October 2006.

12. Editor's Addresses

Jean-Louis Le Roux (Editor)
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@orange-ftgroup.com

Jean-Philippe Vasseur (Editor)
Cisco Systems, Inc.
1414 Massachusetts avenue
Boxborough , MA - 01719
USA
Email: jpv@cisco.com

13. Contributors' Addresses

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

Raymond Zhang
BT Infonet
2160 E. Grand Ave.
El Segundo, CA 90025
USA
Email: raymond_zhang@bt.infonet.com

14. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary 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

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY

Internet Draft [draft-ietf-pce-disco-proto-ospf-08.txt](#) October 2007

WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The IETF Trust (2007). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

