

Network Working Group
Internet Draft
Category: Standard Track
Expires: November 2007

J.L. Le Roux (Editor)
France Telecom

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

Yuichi Ikejiri
NTT Communications

Raymond Zhang
BT Infonet

May 2007

IS-IS protocol extensions for Path Computation Element (PCE) Discovery

[draft-ietf-pce-disco-proto-isis-04.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>.

This Internet-Draft will expire on November 3, 2007.

Copyright Notice

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

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 (PCE), along with some information that can be used 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 discover PCEs consists of using IGP flooding. For that purpose this document defines extensions to the Intermediate System to Intermediate System (IS-IS) routing protocol for the advertisement of PCE Discovery information within an IS-IS area or within the entire IS-IS 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 Information.....	5
3.1.1.	PCE Discovery Information.....	5
3.1.2.	PCE Congestion Information.....	6
3.2.	Flooding Scope.....	6
4.	IS-IS Extensions.....	7
4.1.	The IS-IS PCED Sub-TLV.....	7
4.1.1.	PCE-ADDRESS Sub-TLV.....	8
4.1.2.	The PATH-SCOPE Sub-TLV.....	8
4.1.3.	PCE-DOMAIN Sub-TLV.....	10
4.1.4.	NEIG-PCE-DOMAIN Sub-TLV.....	11
4.1.5.	PCE-CAP-FLAGS Sub-TLV.....	11
4.1.6.	The CONGESTION Sub-TLV.....	12
5.	Elements of Procedure.....	13
5.1.1.	CONGESTION Sub-TLV Specific Procedures.....	14
6.	Backward Compatibility.....	14

7.	IANA Considerations.....	15
7.1.	IS-IS Sub-TLV.....	15
7.2.	PCED Sub-TLVs registry.....	15
8.	Security Considerations.....	15
9.	Manageability Considerations.....	16
9.1.	Control of Policy and Functions.....	16
9.2.	Information and Data Model.....	16
9.3.	Liveness Detection and Monitoring.....	16
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.....	17
11.	References.....	17
11.1.	Normative References.....	17
11.2.	Informative References.....	18
12.	Editors' Addresses:.....	18
13.	Contributors' Adresses:.....	18
14.	Intellectual Property Statement.....	19

[1.](#) Terminology

Terminology used in this document

AS: Autonomous System.

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

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

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

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.

IS-IS LSP: Link State PDU

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 "domain" in [RFC4655]). Examples of PCE-Domains include IGP areas and Autonomous Systems. This should be distinguished from an IS-IS routing domain as defined by [ISO].

Le Roux, Vasseur et al. IS-IS extensions for PCE discovery [Page 3]

Internet Draft [draft-ietf-pce-disco-proto-isis-04.txt](#)

May 2007

PCEP: Path Computation Element communication 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. This relies on a communication protocol between PCC and PCE, and 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 some PCEs in other domains, e.g. in case of inter-domain TE LSP computation.

A network may contain a large number of PCEs with potentially

distinct capabilities. In such a context it is highly desirable to have a mechanism for automatic and dynamic PCE discovery, which allows PCCs to automatically discover a set of PCEs, along with additional information about each PCE that may be required for the PCC to perform PCE selection. Additionally, it is valuable for a PCC to dynamically detect new PCEs or any modification of the PCE information. Detailed requirements for such a PCE discovery mechanism are provided in [\[RFC4674\]](#).

Moreover, it may also be useful to discover when a PCE experiences processing congestion and when it exits such a state, in order for the PCCs to take some appropriate actions (e.g. redirect their requests to another PCE). 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, 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 IS-IS extensions to allow a PCE in an IS-IS 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 [\[RFC4674\]](#). This document also defines extensions allowing a PCE in an IS-IS routing domain to advertise its processing congestion state.

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

This document defines a new sub-TLV (named PCE Discovery (PCED) to be carried within the IS-IS Router Capability TLV ([\[IS-IS-CAP\]](#)).

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

This document does not define any new IS-IS elements of procedure. The procedures defined in [\[IS-IS-CAP\]](#) MUST be used.

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

This document defines a set of sub-TLVs that are nested within each other. When the degree of nesting TLVs is 2 (a TLV is carried within another TLV) the TLV carried within a TLV is called a sub-TLV. Strictly speaking, when the degree of nesting is 3, a subsub-TLV is carried within a sub-TLV that is itself carried within a TLV. For the sake of terminology simplicity, we refer to sub-TLV, a TLV carried within a TLV regardless of the degree of nesting.

[3. Overview](#)

[3.1. PCE Information](#)

The PCE information advertised via IS-IS falls into two categories: PCE Discovery information and PCE Congestion information.

[3.1.1. PCE Discovery Information](#)

The PCE Discovery information is comprised 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;
- The PCE path computation scope (i.e. inter-area, inter-AS, inter-layer);
- The set of one or more PCE-Domain(s) into which the PCE has visibility and can compute paths;
- The set of one or more neighbor PCE-Domain(s) towards which a PCE can compute paths;

- A set of communication capabilities (e.g. support for request prioritization) and path computation specific capabilities (e.g. supported constraints).

Optional elements to describe more complex capabilities may also be advertised.

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.

[3.1.2.](#) PCE Congestion Information

The PCE Congestion information is optional and can be used to report a PCE's processing congestion state along with an estimated congestion duration. This is dynamic information, which may change with PCE activity.

Procedures for a PCE to move from a processing congestion state to a non-congestion state are beyond the scope of this document, but the rate at which a PCE Status change is advertised MUST NOT impact by any means the IGP scalability. Particular attention should be given on procedures to avoid state oscillations.

[3.2.](#) Flooding Scope

The flooding scope for PCE information advertised through IS-IS can be a single L1 area, a L1 area and the L2 sub-domain, or the entire IS-IS routing domain.

[4.](#) IS-IS Extensions

[4.1.](#) The IS-IS PCED Sub-TLV

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

The format of the IS-IS PCED sub-TLV and its sub-TLVs is identical to the TLV format used by the Traffic Engineering Extensions to IS-IS [[RFC3784](#)]. That is, the TLV is composed of 1 octet for the type, 1 octet specifying the TLV length, and a value field. The Length field defines the length of the value portion in bytes.

The IS-IS PCED sub-TLV has the following format:

TYPE: To be assigned by IANA (suggested value = 5)
LENGTH: Variable
VALUE: set of sub-TLVs

Sub-TLVs types are under IANA control.

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

Sub-TLV type	Length	Name
1	variable	PCE-ADDRESS sub-TLV
2	3	PATH-SCOPE sub-TLV
3	variable	PCE-DOMAIN sub-TLV
4	variable	NEIG-PCE-DOMAIN sub-TLV
5	variable	PCE-CAP-FLAGS sub-TLV
6	1	CONGESTION sub-TLV

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

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

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

The CONGESTION sub-TLV is optional and MAY be present in the PCED sub-TLV, to indicate a PCE's processing congestion state.

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 sub-TLV is carried within an IS-IS CAPABILITY TLV defined in [[IS-IS-CAP](#)].

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

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

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

The PCE-ADDRESS sub-TLV is mandatory; it MUST be present within the PCED sub-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.

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

TYPE: To be assigned by IANA (Suggested value =1)
LENGTH: 5 for IPv4 address and 17 for IPv6 address
VALUE: This comprises one octet indicating the address-type and 4 or 16 bytes encoding the IPv4 or IPv6 address to be used to reach the PCE.

Address-type:

- 1 IPv4
- 2 IPv6

[4.1.2](#). The 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 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 sub-TLV. There MUST be exactly one instance of the PATH-SCOPE sub-TLV within each PCED sub-TLV.

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:

TYPE: To be assigned by IANA (Suggested value =2)

LENGTH: 3

VALUE: This comprises a one-byte flags field where flag represents a supported path scope, followed by a 2-bytes preferences field indicating PCE preferences.

Here is the structure of the bits flag:

```

+-----+
|0|1|2|3|4|5|Res|
+-----+

```

Bit	Path Scope
0	L bit: Can compute intra-area path
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 LSPs computation
5	Y bit: Can compute or take part into the computation of paths across layers
6-7	Reserved for future usage.

Here is the structure of the preferences field

```

+-----+
|PrefL|PrefR|PrefS|PrefY| Res |
+-----+

```

Res: Reserved for future usage.

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

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

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

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

Res: Reserved for future usage.

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 LSP computation (that is the PCE can compute a path towards 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 towards any neighbor AS).

When the Rd and Sd bit are set, the PCED sub-TLV MUST NOT contain any NEIG-PCE-DOMAIN sub-TLV (see 4.1.4).

When the R/S bit is cleared, the Rd/Sd bit SHOULD be cleared and MUST be ignored.

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 preference can be used for weighted load balancing of requests. An operator may decide to configure a preference for each computation scope to each PCE so as to balance the path computation load among them. The algorithms used by a PCC to balance its path computation requests according to such PCE preference are out of the scope of this document and is a matter for local or network wide policy. The same or distinct preferences may be used for each scopes. For instance an operator that wants a PCE capable of both inter-area and inter-AS computation to be used preferably for inter-AS computation may configure a PrefS higher than the PrefR.

When the L bit, R bit, S bit or Y bit are cleared the PrefL, PrefR, PrefS, PrefY fields SHOULD respectively be set to 0 and MUST be

ignored.

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

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

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

The PCE-DOMAIN sub-TLV MAY 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 (see [section 5](#) for a discussion of how the flooding scope is set and interpreted).

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

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

TYPE: To be assigned by IANA (Suggested value =3)

LENGTH: Variable

VALUE: This comprises one octet indicating the domain-type (area ID or AS Number) and a variable length IS-IS area ID or a 32 bits AS number, identifying a PCE-domain where the PCE has visibility.

Two domain types are defined:

- 1 Area ID
- 2 AS Number

The Area ID is the area address as defined in [\[ISO\]](#).

When coded in 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.

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

The NEIG-PCE-DOMAIN sub-TLV specifies a neighbour PCE-domain (area, AS) toward which a PCE can compute paths. It means that the PCE can

take part in the computation of inter-domain TE LSPs whose path transits this neighbour 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:

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

LENGTH: Variable

VALUE: This comprises one octet indicating the domain-type (area ID or AS Number) and a variable length IS-IS area ID or a 32 bits AS number, identifying a PCE-domain towards which the PCE can compute paths.

Two domain types are defined:

- 1 Area ID
- 2 AS Number

The Area ID is the area address as defined in [\[ISO\]](#).

When coded in 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.

The NEIG-PCE-DOMAIN sub-TLV MUST be present 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.

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

The PCE-CAP-FLAGS sub-TLV is an optional sub-TLV used to indicate PCEP related capabilities. It MAY be present within the PCED sub-TLV. It MUST NOT be present more than once.

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 as bit zero, where each bit represents one PCE capability.

The PCE-CAP-FLAGS sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =4)
 LENGTH: Multiple of 4
 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.

The PCE capability registry is managed by IANA, it is common with OSPF and defined in [[PCED-OSPF](#)].

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

[4.1.6](#). The CONGESTION Sub-TLV

The CONGESTION sub-TLV is used to indicate that a PCE is experiencing a processing congestion state and may optionally include expected PCE congestion duration.

The CONGESTION sub-TLV is optional, it MAY be carried within the PCED sub-TLV. It MUST NOT be present more than once.

The format of the CONGESTION sub-TLV is as follows:

TYPE: To be assigned by IANA (Suggested value =6)
 LENGTH: 3
 VALUE: This comprises a one byte of bit flags indicating the congestion status, followed by a 2-bytes field indicating the congestion duration.

Here is the TLV structure

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|C|      Reserved|      Congestion Duration      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Value

-C bit: When set this indicates that the PCE is experiencing congestion and cannot accept any new request. When cleared this indicates that the PCE is not experiencing congestion and can accept new requests.

-Congestion Duration: 2-bytes, the estimated PCE congestion duration in seconds.

When C is set and the Congestion Duration field is equal to 0, this means that the Congestion Duration is unknown.

When C is cleared the Congestion Duration SHOULD be set to 0 and MUST be ignored.

5. Elements of Procedure

The PCED sub-TLV is advertised within an IS-IS Router Capability TLV defined in [[IS-IS-CAP](#)]. As such, elements of procedures are inherited from those defined in [[IS-IS-CAP](#)].

The flooding scope is controlled by the S flag in the IS-IS Router Capability TLV (see [[IS-IS-CAP](#)]). When the scope of the PCED sub-TLV is area local it MUST be carried within an IS-IS Router Capability TLV having the S bit cleared. When the scope of the PCED sub-TLV is the entire IS-IS routing domain, it MUST be carried within an IS-IS Router Capability TLV having the S bit set. Note that when only the L bit of the PATH-SCOPE sub-TLV is set, the flooding scope MUST be area local.

Note that a L1L2 node may include both in its L1 and L2 LSPs a PCED TLV in a Router Capability TLV with the S bit cleared. This allows restricting the flooding scope to the L1 area and the L2 sub-domain.

An IS-IS router MUST originate a new IS-IS LSP whenever there is a change in a PCED TLV associated with a PCE it advertises.

When a PCE is deactivated, the IS-IS Router advertising this PCE MUST originate a new IS-IS LSP that does no longer include the corresponding PCED TLV.

The PCE address(s), i.e. the address(s) indicated within the PCE ADDRESS sub-TLV, must be reachable via some prefix(es) advertised by IS-IS; this allows speeding up the detection of a PCE failure. Note that when the PCE address is no longer reachable, this means that the PCE node has failed or has been torn down, or that there is no longer IP connectivity to the PCE node.

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

A change in PCED information 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 (e.g.,

address, preferences, scope) and other information may be automatically determined by the PCE (e.g. areas of visibility).

5.1.1. CONGESTION Sub-TLV Specific Procedures

When a PCE enters into a processing congestion state, the conditions of which are implementation dependent, a new IS-IS LSP with a CONGESTION sub-TLV with the C bit set and optionally a non-null expected congestion duration MAY be generated.

When a PCE exists from the processing congestion state, the conditions of which are implementation dependent, two cases are considered:

- If the congestion duration in the previously originated CONGESTION sub-TLV was null, a CONGESTION sub-TLV with the C bit cleared SHOULD be generated;
- If the congestion duration in the previously originated CONGESTION sub-TLV was non null, a CONGESTION sub-TLV with the C bit cleared MAY be generated. Note that in some particular cases it may be desired to originate a CONGESTION sub-TLV with the C bit cleared if the congestion duration was over estimated.

The congestion duration allows a reduction in the amount of IS-IS flooding, as only uncongested-to-congested state transitions are advertised.

An IS-IS implementation SHOULD support an appropriate dampening algorithm so as to dampen flooding of PCE Congestion information in order to not impact the IS-IS scalability. It is RECOMMENDED to introduce some hysteresis for congestion state transition, so as to avoid state oscillations that may impact IS-IS performance. For instance two thresholds MAY be configured: a resource congestion upper-threshold and a resource congestion lower-threshold. An LSR enters the congested state when the CPU load reaches the upper threshold and leaves the congested state when the CPU load goes under the lower threshold.

Upon receipt of an updated CONGESTION sub-TLV a PCC should take appropriate actions. In particular, the PCC SHOULD stop sending requests to a congested PCE, and SHOULD gradually start sending again

requests to a PCE that is no longer congested.

[6. Backward Compatibility](#)

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

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

[7. IANA Considerations](#)

[7.1. IS-IS Sub-TLV](#)

Once a registry for the IS-IS Router Capability sub-TLVs, defined in [[IS-IS-CAP](#)] has been assigned, IANA will assign a new sub-TLV code-point for the PCED sub-TLV carried within the Router Capability TLV.

Value	Sub-TLV	References
-----	-----	-----
5	PCED sub-TLV	(this document)

[7.2. PCED Sub-TLVs Registry](#)

The PCED sub-TLV referenced above is constructed from sub-TLVs. Each sub-TLV includes a 8-bit type identifier.

The IANA is requested to create a new sub-registry of the IS-IS Router Capability sub-TLVs registry, named the "PCED sub-TLVs" registry, and manage sub-TLV type identifiers as follows:

- sub-TLV Type
- sub-TLV Name
- Reference

This document defines five sub-TLVs as follows (suggested values):

Value	TLV name	References
-----	-----	-----
1	PCE-ADDRESS	This document
2	PATH-SCOPE	This document
3	PCE-DOMAIN	This document
4	NEIG-PCE-DOMAIN	This document
5	PCE-CAP-FLAGS	This document
6	CONGESTION	This document

New sub-TLV type values may be allocated only by an IETF Consensus action.

8. Security Considerations

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

Mechanisms defined to ensure authenticity and integrity of IS-IS LSPs [[RFC3567](#)], and their TLVs, can be used to secure the PCED sub-TLV as well.

IS-IS provides no mechanism for protecting the privacy of LSPs, and in particular the privacy of the PCE discovery information.

Le Roux, Vasseur et al. IS-IS extensions for PCE discovery [Page 15]

Internet Draft [draft-ietf-pce-disco-proto-isis-04.txt](#)

May 2007

9. Manageability Considerations

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

9.1. Control of Policy and Functions

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

Particularly, a PCE implementation SHOULD allow configuring the following parameters on the PCE:

- The PCE IPv4/IPv6 address(es) (see [section 4.1.1](#))
- 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 4.1.2](#))
- The PCE domains (see [section 4.1.3](#))
- The neighbour PCE domains (see [section 4.1.4](#))
- The PCE capabilities (see [section 4.1.5](#))

[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 IS-IS liveness detection. IS-IS already includes a liveness detection mechanism (Hello PDUs), and PCE discovery does not require additional capabilities.

Procedures defined in [section 5.1](#) allow a PCC detecting 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 PCED information in the PCC and in the PCE, which is stored in the PCED MIB [[PCED-MIB](#)]. The number of dropped, corrupt, and rejected information elements are stored in the PCED MIB.

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

The IS-IS extensions defined in this documents do not imply any requirement on other protocols.

[9.6.](#) Impact on Network Operations

Frequent changes in PCE information, and particularly in PCE congestion information, may have a significant impact on IS-IS and

might destabilize the operation of the network by causing the PCCs to swap between PCEs.

As discussed in [section 5.1](#), a PCE implementation SHOULD support an appropriate dampening algorithm so as to dampen IS-IS flooding in order to not impact the IS-IS scalability.

Also, 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 discovery messages rate-limiting.
- 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.

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.

[ISO] "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO DP 10589, February 1990.

[RFC3784] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", [RFC 3784](#), June 2004.

[IS-IS-CAP] Vasseur, J.P. et al., "IS-IS extensions for advertising router information", [draft-ietf-isis-caps](#), 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.

[RFC3567] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", [RFC 3567](#), July 2003.

Path Computation Element (PCE) Discovery", [draft-ietf-pce-disco-proto-ospf](#), work in progress.

[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.

[12](#). Editors' 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' Adresses:

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 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.

