

Network Working Group
Internet Draft
Category: Standards Track
Expires: June 2007

E. Oki
NTT
A. Farrel
Old Dog Consulting

January 2007

Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions

[draft-oki-pce-pcep-xro-00.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/1id-abstracts.txt>.

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

Abstract

The Path Computation Element (PCE) provides functions of path computation in support of traffic engineering in Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

When a Path Computation Client (PCC) requests a PCE for a route, it may be useful for the PCC to specify as constraints to the path computation abstract nodes, resources, and Shared Risk Link Groups (SRLGs) that are to be explicitly excluded from routes. Such constraints are termed route exclusions.

The PCE Communication Protocol (PCEP) is designed as a communication protocol between PCCs and PCEs. This document presents PCEP extensions for route exclusions.

Conventions used in this document

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

Table of Contents

1. Introduction.....	2
2. Protocol Procedures and Extensions.....	3
2.1. Exclude Route Object (XRO).....	4
2.2.1. Processing Rules.....	7
2.2. Explicit Route Exclusion.....	8
2.2.1. Processing Rules.....	9
3. IANA Considerations.....	10
3.1. PCEP Objects.....	10
3.2. Error Object Field Values.....	10
4. Manageability considerations.....	10
5. Security Considerations.....	11
6. References.....	11
6.1. Normative Reference.....	11
6.2. Informative Reference.....	11
7. Authors EAddresses.....	11

[1. Introduction](#)

The Path Computation Element (PCE) defined in [[RFC4655](#)] is an entity that is capable of computing a network path or route based on a network graph, and applying computational constraints. A Path Computation Client (PCC) may make requests to a PCE for paths to be computed.

When a PCC requests a PCE for a route, it may be useful for the PCC to specify abstract nodes, resources, and Shared Risk Link Groups (SRLGs) that are to be explicitly excluded from the route.

For example, disjoint paths for inter-domain LSPs may be computed by cooperation between PCEs, each of which computes segments of the paths across one domain. In order to achieve path computation for a secondary (backup) path, a PCE may act as a PCC to request another PCE for a route that must be a node/link/SRLG disjoint from the primary (working) path. Another example is where a network operator wants path to avoid specified nodes for administrative reasons perhaps because the specified nodes will be out-of-services in near future.

[RFC4657] specifies generic requirements for a communication protocol between PCCs and PCEs. Generic constraints described in [RFC4657] include route exclusions for links, nodes, and SRLGs. That is, the requirement for support of route exclusions within the PCC-PCE communication protocol is already established.

The PCE communication protocol (PCEP) is designed as a communication protocol between PCCs and PCEs and is defined in [PCEP]. This document presents PCEP extensions to satisfy the requirements for route exclusions as described in Sections 5.1.4 and 5.1.16 of [RFC4657].

Note that MPLS-TE and GMPLS signaling extensions for communicating route exclusions between network nodes for specific Label Switched Paths (LSPs) are described in [XRO]. Route exclusions may be specified during provisioning requests for specific LSPs using the `mplsTunnelHopInclude` object of MPLS-TE-STD-MIB defined in [RFC3812].

2. Protocol Procedures and Extensions

This section describes the procedures adopted by a PCE handling a request for path computation with route exclusions received from a PCC, and defines how those exclusions are encoded.

There are two types of route exclusion described in [XRO].

1. Exclusion of certain abstract nodes or resources on the whole path. This set of abstract nodes is referred to as the Exclude Route List.
2. Exclusion of certain abstract nodes or resources between a specific pair of abstract nodes present in an explicit path. Such specific exclusions are referred to as an Explicit Route Exclusion.

This document defines protocol extensions to allow a PCC to specify both types of route exclusions to a PCE on a path computation request.

A new PCEP object is defined as the Exclude Route Object (XRO) to convey the Exclude Route List. The existing Include Route Object (IRO) in PCEP [PCEP] is modified by introducing a new IRO subobject, the Explicit Exclusion Route Subobject (EXRS), to convey Explicit Route Exclusions.

2.1. Exclude Route Object (XRO)

The XRO is OPTIONAL and MAY be carried within PCReq and PCRep messages.

When present in a PCReq message, the XRO provides a list of network resources that the PCE is requested to exclude from the path that it computes. Flags associated with each list member instruct the PCE as to whether the network resources must be excluded from the computed path or whether the PCE should make best efforts to exclude the resources from the computed path.

The XRO MAY be used on PCRep message with the NO-PATH object to indicate the set of elements of the original XRO that prevented the PCE from finding a path. The XRO MAY also be used on a PCRep message for a successful path computation when the PCE wishes to provide a set of exclusions to be signaled during LSP setup.

XRO Object-Class is to be assigned by IANA (recommended value=17)

XRO Object-Type is to be assigned by IANA (recommended value=1)

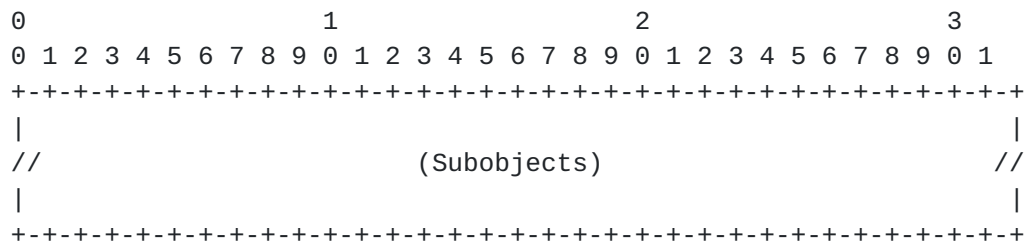


Figure 1: XRO body format

Subobjects. The XRO is up made of one or more subobject(s). An XRO with no subobjects MUST not be sent and SHOULD be ignored on receipt.

In the following subobject definitions a set of fields have consistent meaning as follows:

X

The X-bit indicates whether the exclusion is mandatory or desired. 0 indicates that the resource specified MUST be excluded from the path computed by the PCE 1 indicates that the resource specified SHOULD be excluded from the path computed by the PCE, but MAY be included subject to PCE policy and the absence of a viable path that meets the other other constraints and excludes the resource.

Type

The type of the subobject. The following subobject types are defined.

Type	Subobject
-----+-----	
1	IPv4 prefix
2	IPv6 prefix
3	Unnumbered Interface ID
4	Autonomous system number
5	SRLG

Length

The length of the subobject including the Type and Length fields.

Prefix Length

Where present, this field can be used to indicate a set of addresses matching a prefix. If the subobject indicates a single address, the prefix length **MUST** be set to the full length of the address.

Attribute

The Attribute field indicates how the exclusion subobject is to be interpreted.

0 Interface

The subobject is to be interpreted as an interface or set of interfaces. All interfaces identified by the subobject are to be excluded from the computed path according to the setting of the X-bit. This value is valid only for subobject types 1, 2, and 3.

1 Node

The subobject is to be interpreted as a node or set of nodes. All nodes identified by the subobject are to be excluded from the computed path according to the setting of the X-bit. This value is valid only for subobject types 1, 2, 3, and 4.

2 SRLG

The subobject identifies an SRLG explicitly or indicates all of the SRLGs associated with the resource or resources identified by the subobject. Resources that share any SRLG with those identified are to be excluded from the computed path according to the setting of the X-bit. This value is valid for all subobjects.

Reserved

Reserved fields **MUST** be transmitted as zero and **SHOULD** be

ignored on receipt.

Oki and Farrel

Expires July 2007

[Page 5]

The subobjects are encoded as follows:

IPv4 prefix Subobject

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|X|  Type = 1   |      Length      | IPv4 address (4 bytes)          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv4 address (continued)          | Prefix Length |  Attribute  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

IPv6 prefix Subobject

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|X|  Type = 2   |      Length      | IPv6 address (16 bytes)          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv6 address (continued)          |                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv6 address (continued)          |                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv6 address (continued)          |                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| IPv6 address (continued)          | Prefix Length |  Attribute  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Unnumbered Interface ID Subobject

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|X|  Type = 3   |      Length      |  Reserved  |  Attribute  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               TE Router ID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Interface ID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The TE Router ID and Interface ID fields are as defined in [\[RFC3477\]](#).

Autonomous System Number Subobject

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

If a two-octet AS number is used, the optional AS Number High Octets MUST be set to zero.

SRLG Subobject

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									
X Type = 5										Length										SRLG Id (4 bytes)																			
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									
SRLG Id (continued)										Reserved										Attribute																			
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									

The Attribute SHOULD be set to two (2) and SHOULD be ignored on receipt.

2.2.1. Processing Rules

A PCC builds an XRO to encode all of the resources that it wishes the PCE to exclude from the path that it is requested to compute. For each exclusion, the PCC clears the X-bit to indicate that the PCE is required to exclude the resources, or sets the X-bit to indicate that the PCC simply desires that the resources are excluded. For each exclusion, the PCC also sets the Attribute field to indicate how the PCE should interpret the contents of the exclusion subobject.

When a PCE receives a PCReq message it looks for an XRO to see if exclusions are required. If the PCE finds more than one XRO it MUST use the first one in the message and MUST ignore subsequent instances.

If the PCE does not recognize the XRO it MUST return a PCErr message with Error-Type "Unknown Object" as described in [PCEP].

If the PCE is unwilling or unable to process the XRO it MUST return a PCErr message with the Error-Type "Not supported object" and follow the relevant procedures described in [PCEP].

If the PCE processes the XRO and attempts to compute a path, it MUST adhere to the requested exclusions as expressed in the XRO. That is, the returned path MUST NOT include any resources encoded with the X-bit clear, and SHOULD NOT include any with the X-bit set unless alternate paths that match the other constraints expressed in the PCReq are unavailable.

When a PCE returns a path in a PCRep it MAY also supply an XRO. In this case, the PCC SHOULD apply the contents using the same rules as in [XRO] and SHOULD signal the an RSVP-TE XRO to indicate the exclusions that downstream LSRs should apply. This may be particularly useful in per-domain path computation scenarios. [Note that this does not match the behavior for an explicit path where an IRO is used to force inclusions and an ERO is used to report a computed path. We could consider using a separate object to report the XRO that should be signaled.]

In the event that no suitable path can be computed and the PCE returns a PCRep message containing a NO-PATH object, the PCE MAY also include an XRO that lists one or more subobjects from the original XRO that have contributed to the PCE's inability to select a path.

2.2. Explicit Route Exclusion

Explicit Route Exclusion defines network elements that must not or should not be used on the path between two abstract nodes or resources explicitly indicated in the Include Route Object (IRO) [[PCEP](#)]. This information is encoded by defining a new subobject for the IRO .

The new IRO subobject, the Explicit Exclusion Route Subobject (EXRS), has type defined by IANA (see [Section 3](#)). The EXRS contains one or more subobjects in its own right. An EXRS MUST NOT be sent with no subobjects, and if received with no subobjects MUST be ignored.

The format of the EXRS is as follows:

[illegible]

L

MUST be set to zero on transmission and MUST be ignored on receipt.

Reserved

MUST be set to zero on transmission and MUST be ignored on receipt.

The EXRS subobject may carry any of the subobjects defined for inclusion in the XRO by this document or by future documents. The meanings of the fields of the XRO subobjects are unchanged when the subobjects are included in an EXRS, except that scope of the exclusion is limited to the single hop between the previous and subsequent elements in the IRO.

2.2.1. Processing Rules

A PCC that supplies a partial explicit route to a PCE in an IRO MAY also specify explicit exclusions by including one or more EXRSes in the IRO.

If a PCE parses an IRO in a received PCReq message and encounters an EXRS and does not recognize the subobject it MUST respond with a PCErr message using the Error-Type "Unrecognized IRO subobject" and set the Error-Value to the subobject type code of the EXRS (see [Section 3](#)).

If a PCE parses an IRO and encounters an EXRS that it recognizes, but detects an EXRS subobject that it does not recognize it MUST act according to the setting of the X-bit in the subobject. If the X-bit is clear, the PCE MUST respond with a PCErr with Error-Type "Unrecognized EXRS subobject" and set the Error-Value to the EXRS subobject type code (see [Section 3](#)). If the X-bit is set, the PCE MAY respond with a PCErr as already stated or MAY ignore the EXRS subobject: this choice is a local policy decision.

If a PCE parses an IRO and encounters an EXRS subobject that it recognizes, it MUST act according to the requirements expressed in the subobject. That is, if the X-bit is clear, the PCE MUST NOT produce a path that includes any resource identified by the EXRS subobject in the path between the previous abstract node in the IRO and the next abstract node in the IRO. If the X-bit is set, the PCE SHOULD NOT produce a path that includes any resource identified by the EXRS subobject in the path between the previous abstract node in the IRO and the next abstract node in the IRO unless it is not possible to construct a path that avoids that resource while still complying with the other constraints expressed in the PCReq message.

A successful path computation reported in a PCRep message MUST

include an ERO to specify the path that has been computed. That ERO

Oki and Farrel

Expires July 2007

[Page 9]

MAY contain specific route exclusions using the EXRS as specified in [[XRO](#)].

If the path computation fails and a PCErr is returned with a NO-PATH object, the PCE MAY include an IRO to report the hops that could not be complied with, and that IRO MAY include EXRSes.

[3. IANA Considerations](#)

[3.1. PCEP Objects](#)

The "PCEP Parameters" registry contains a subregistry "PCEP Objects". IANA is requested to make the following allocations from this registry.

Object Class	Name	Object Type	Name
17	XRO	1	Route exclusion

[3.2. Error Object Field Values.](#)

The "PCEP Parameters" registry contains a subregistry "PCEP Errors". IANA is requested to make the following allocations from this registry.

Values in this section are recommended and to be confirmed by IANA.

Error Type	Meaning and Error-Values
11	Unrecognized IRO subobject

Note that this Error-Type has been omitted from [[PCEP](#)] where it is required. It is expected that it will be added to a later version of [[PCEP](#)] and removed from this document.

12	Unrecognized EXRS subobject
----	-----------------------------

[4. Manageability Considerations](#)

A MIB module for management of the PCEP is specified in a separate document. This MIB module allows examination of individual PCEP messages, in particular requests, responses and errors.

The MIB module MUST be extended to include the ability to view the route exclusion extensions defined in this document.

5. Security Considerations

The new exclude route mechanisms defined in this document allow finer and more specific control of the path computed by a PCE. Such control increases the risk if a PCEP message is intercepted, modified, or spoofed. Therefore, the security techniques described in [[PCEP](#)] are considered more important.

6. References

6.1. Normative Reference

[RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", [RFC 2119](#), March 1997.

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

6.2. Informative Reference

[RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", [RFC 3812](#), June 2004.

[RFC4655] A. Farrel, JP. Vasseur and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), September 2006.

[RFC4657] J. Ash and J.L. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", [RFC 4657](#), September 2006.

[XR0] Lee et al, "Exclude Routes - Extension to RSVP-TE", [draft-ietf-ccamp-rsvp-te-exclude-route](#), (work in progress).

7. Authors' Addresses

Eiji Oki
NTT
3-9-11 Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
Email: oki.eiji@lab.ntt.co.jp

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Full 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.

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.

Intellectual Property

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.

