

Network Working Group
Request for Comments: 5521
Category: Standards Track

E. Oki
University of Electro-Communications
T. Takeda
NTT
A. Farrel
Old Dog Consulting
April 2009

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

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

The Path Computation Element (PCE) provides functions of path computation in support of traffic engineering (TE) 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 the computed route. 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.

Table of Contents

| | | |
|--------|--|----|
| 1. | Introduction | 3 |
| 1.1. | Conventions Used in This Document | 3 |
| 2. | Protocol Procedures and Extensions | 4 |
| 2.1. | Exclude Route Object (XRO) | 4 |
| 2.1.1. | Definition | 4 |
| 2.1.2. | Processing Rules | 8 |
| 2.2. | Explicit Route Exclusion | 9 |
| 2.2.1. | Definition | 9 |
| 2.2.2. | Processing Rules | 10 |
| 3. | Exclude Route with Confidentiality | 11 |
| 3.1. | Exclude Route Object (XRO) Carrying Path-Key | 11 |
| 3.1.1. | Definition | 11 |
| 3.1.2. | Processing Rules | 12 |
| 4. | IANA Considerations | 13 |
| 4.1. | PCEP Objects | 13 |
| 4.2. | New Subobject for the Include Route Object | 13 |
| 4.3. | Error Object Field Values | 13 |
| 4.4. | Exclude Route Flags | 14 |
| 5. | Manageability Considerations | 14 |
| 6. | Security Considerations | 14 |
| 7. | References | 15 |
| 7.1. | Normative References | 15 |
| 7.2. | Informative References | 15 |
| | Acknowledgements | 16 |

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 Label Switched Paths (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 node/link/SRLG disjoint from the primary (working) path. Another example is where a network operator wants a path to avoid specified nodes for administrative reasons, perhaps because the specified nodes will be out-of-service in the 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 [RFC5440]. 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 [RFC4874]. Route exclusions may be specified during provisioning requests for specific LSPs by setting the `mplsTunnelHopInclude` object of MPLS-TE-STD-MIB defined in [RFC3812] to false (2).

1.1. 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].

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 [RFC4874].

1. Exclusion of certain abstract nodes or resources from 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, the Exclude Route Object (XRO), is defined to convey the Exclude Route List. The existing Include Route Object (IRO) in PCEP [RFC5440] 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)

2.1.1. Definition

The XRO is OPTIONAL and MAY be carried within Path Computation Request (PCReq) and Path Computation Reply (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 a PCRep message that carries the NO-PATH object (i.e., one that reports a path computation failure) 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 using the extensions to Resource Reservation Protocol (RSVP)-TE [RFC4874].

The XRO Object-Class is 17.

The XRO Object-Type is 1.

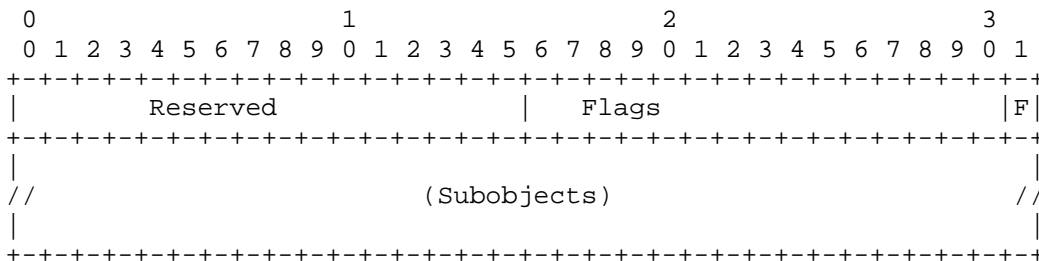


Figure 1: XRO Body Format

Reserved: 16 bits - MUST be set to zero on transmission and SHOULD be ignored on receipt.

Flags: 16 bits - The following flags are currently defined:

F (Fail - 1 bit): when set, the requesting PCC requires the computation of a new path for an existing TE LSP that has failed. If the F bit is set, the path of the existing TE LSP MUST be provided in the PCReq message by means of a Record Route Object (RRO) defined in [RFC5440]. This allows the path computation to take into account the previous path and reserved resources to avoid double bandwidth booking should the Traffic Engineering Database (TED) have not yet been updated or the corresponding resources not be yet been released. This will usually be used in conjunction with the exclusion from the path computation of the failed resource that caused the LSP to fail.

Subobjects: The XRO is made up 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 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 |
| 4 | Unnumbered Interface ID |
| 32 | Autonomous system number |
| 34 | 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

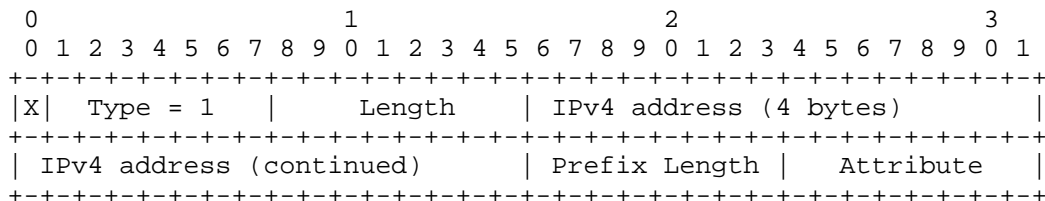
The subobject identifies an SRLG explicitly or indicates all of the SRLGs associated with the resource or resources identified by the subobject. Resources that share any SRLG with those identified are to be excluded from the computed path according to the setting of the X-bit. This value is valid for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

Reserved

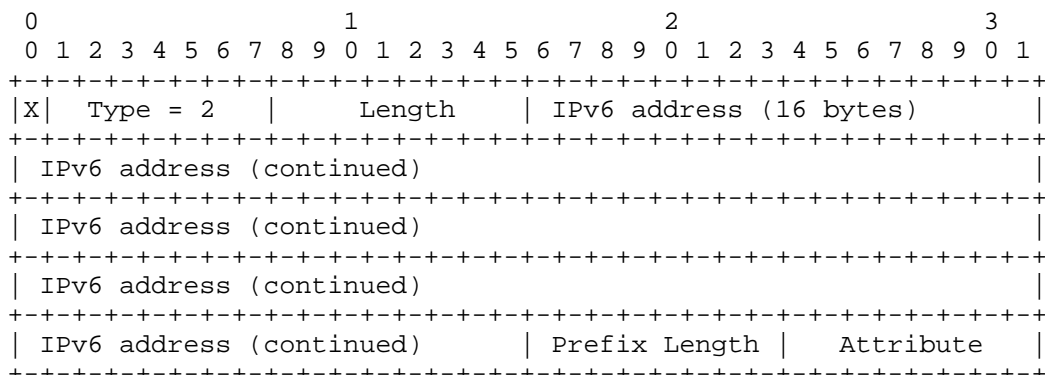
Reserved fields within subobjects MUST be transmitted as zero and SHOULD be ignored on receipt.

The subobjects are encoded as follows:

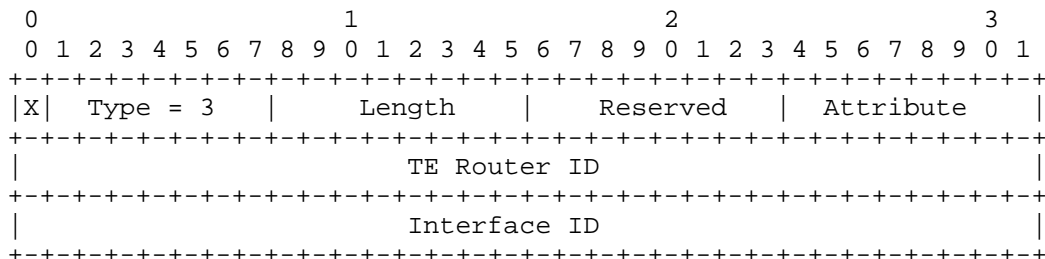
IPv4 prefix Subobject



IPv6 prefix Subobject

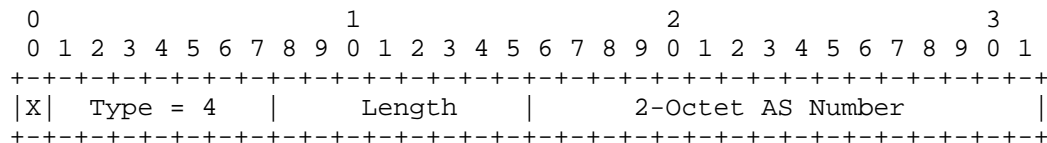


Unnumbered Interface ID Subobject




The TE Router ID and Interface ID fields are as defined in [RFC3477].

Autonomous System Number Subobject



Note that as in other PCEP objects [RFC5440] and RSVP-TE objects [RFC3209], no support for 4-octet Autonomous System (AS) Numbers is provided. It is anticipated that, as 4-octet AS Numbers become more common, both PCEP and RSVP-TE will be updated in a consistent way to add this support.

SRLG Subobject



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

2.1.2. 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 [RFC5440].

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 [RFC5440].

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. An XRO in a PCRep message with the NO-PATH object indicates that the set of elements of the original XRO prevented the PCE from finding a path. On the other hand, if an XRO is present in a PCRep message without a NO-PATH object, the PCC SHOULD apply the contents using the same rules as in [RFC4874] and the PCC or a corresponding LSR SHOULD signal an RSVP-TE XRO to indicate the exclusions that downstream LSRs should apply. This may be particularly useful in per-domain path computation scenarios [RFC5152].

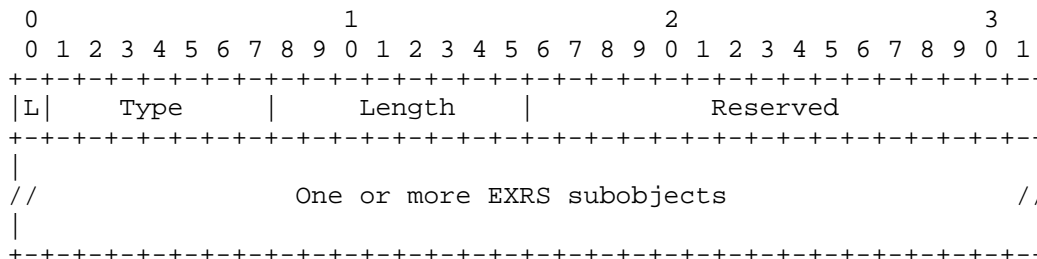
2.2. Explicit Route Exclusion

2.2.1. Definition

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) [RFC5440]. This information is encoded by defining a new subobject for the IRO.

The new IRO subobject, the Explicit Exclusion Route subobject (EXRS), has type 33 (see Section 4). 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:



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

Reserved

MUST be set to zero on transmission and SHOULD 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.

2.2.2. 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 EXRSs in the IRO.

If a PCE that does not support the use of EXRS receives an IRO in a PCReq message that contains an EXRS, it will respond according to the rules for a malformed object as described in [RFC5440]. The PCE MAY also include the IRO in the PCErr to indicate in which case the IRO SHOULD be terminated immediately after the unrecognized EXRS.

If a PCE that supports the EXRS in an IRO parses an IRO and encounters an EXRS that contains a subobject that it does not support or 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 4). 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 as specified in [RFC5440]. That ERO MAY contain specific route exclusions using the EXRS as specified in [RFC4874].

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 compiled with as described in [RFC5440], and that IRO MAY include EXRSs.

3. Exclude Route with Confidentiality

3.1. Exclude Route Object (XRO) Carrying Path-Key

3.1.1. Definition

In PCE-based inter-domain diverse path computation, an XRO may be used to find a backup (secondary) path. A sequential path computation approach may be applied for this purpose, where a working (primary) path route is computed first and a backup path route that must be a node/link/SRLG disjoint route from the working path is then computed [RFC5298]. Backward Recursive Path Computation (BRPC) may be used for inter-domain path computation [RFC5441].

In some cases of inter-domain computation (e.g., where domains are administered by different service providers), confidentiality must be kept. For primary path computation, to preserve confidentiality, instead of explicitly expressing the computed route, Path-Key Subobjects (PKSs) [RFC5520] are carried in the Explicit Route Object (ERO) in the PCRep Message.

Therefore, during inter-domain diverse path computation, it may be necessary to request diversity from a path that is not fully known and where a segment of the path is represented by a PKS. This means that a PKS may be present as a subobject of the XRO on a PCReq message.

The format and definition of PKS when it appears as an XRO subobject are as defined in [RFC5520], except for the definition of the L bit. The L bit of the PKS subobject in the XRO MUST be ignored.

3.1.2. Processing Rules

Consider that BRPC is applied for both working and backup path computation in a sequential manner. First, PCC requests PCE for the computation of a working path. After BRPC processing has completed, the PCC receives the results of the working-path computation expressed in an ERO in a PCRep message. The ERO may include PKSs if certain segments of the path are to be kept confidential.

For backup path computation, when the PCC constructs a PCReq Message, it includes the entire working-path in the XRO so that the computed path is node/link disjoint from the working path. The XRO may also include SRLGs to ensure SRLG diversity from the working path. If the working path ERO includes PKS subobjects, these are also included in the XRO to allow the PCE to ensure diversity.

A set of PCEs for backup path computation may be the same as ones for working path computation, or they may be different.

- Identical PCEs

In the case where the same PCEs are used for both path computations, the processing is as follows. During the process of BRPC for backup path computation, a PCE may encounter a PKS as it processes the XRO when it creates a virtual path tree (VPT) in its own domain. The PCE retrieves the PCE-ID from the PKS, recognizes itself, and converts the PKS into a set of XRO subobjects that it uses for the local calculation to create the VPT. The XRO subobjects created in this way MUST NOT be shared with other PCEs. Other operations are the same as BRPC.

- Different PCEs

In the case where a set of PCEs for backup path computation is different from the ones used for working path computation, the processing is as follows. If a PCE encounters a PKS in an XRO when it is creating a virtual path tree in its own domain, the PCE retrieves the PCE-ID from the PKS and sends a PCReq message to the identified PCE to expand the PKS. The PCE computing the VPT treats the path segment in the response as a set of XRO subobjects in performing its path computation. The XRO subobjects determined in this way MUST NOT be shared with other PCEs.

4. IANA Considerations

4.1. PCEP Objects

The "PCEP Parameters" registry contains a subregistry "PCEP Objects". IANA has made the following allocations from this registry.

| Object Class | Name | Reference |
|--------------|--|-----------|
| 17 | XRO Object-Type 1: Route exclusion | [RFC5521] |

This object should be registered as being allowed to carry the following subobjects:

| Subobject Type | Reference |
|---------------------------------|-----------|
| 1 IPv4 prefix | [RFC3209] |
| 2 IPv6 prefix | [RFC3209] |
| 4 Unnumbered Interface ID | [RFC3477] |
| 32 Autonomous system number | [RFC3209] |
| 34 SRLG | [RFC4874] |
| 64 Path-Key with 32-bit PCE ID | [RFC5520] |
| 65 Path-Key with 128-bit PCE ID | [RFC5520] |

4.2. New Subobject for the Include Route Object

The "PCEP Parameters" registry contains a subregistry "PCEP Objects" with an entry for the Include Route Object (IRO).

IANA added a further subobject that can be carried in the IRO as follows:

| Subobject Type | Reference |
|--|-----------|
| 33 Explicit Exclusion Route subobject (EXRS) | [RFC4874] |

4.3. Error Object Field Values

The "PCEP Parameters" registry contains a subregistry "Error Types and Values". IANA made the following allocations from this subregistry.

| Error Type | Meaning | Reference |
|------------|-----------------------------|-----------|
| 11 | Unrecognized EXRS subobject | [RFC5521] |

4.4. Exclude Route Flags

IANA created a subregistry of the "PCEP Parameters" for the bits carried in the Flags field of the Exclude Route Object (XRO). The subregistry is called "XRO Flag Field".

New bits may be allocated only by an IETF Consensus action.

The field contains 16 bits numbered from bit 0 as the most significant bit.

| Bit | Name | Description | Reference |
|-----|-------|-------------|-----------|
| 15 | F-bit | Fail | [RFC5221] |

5. Manageability Considerations

A MIB module for management of the PCEP is being specified in a separate document [PCEP-MIB]. That 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.

Several local policy decisions should be made at the PCE. Firstly, the exact behavior with regard to desired exclusions must be available for examination by an operator and may be configurable. Second, the behavior on receipt of an unrecognized XRO or EXRS subobject with the X-bit set should be configurable and must be available for inspection. The inspection and control of these local policy choices may be part of the PCEP MIB module.

6. 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 because it allows the attacker to exert control over the path that the PCE will compute or to make the path computation impossible. Therefore, the security techniques described in [RFC5440] are considered more important.

Note, however, that the route exclusion mechanisms also provide the operator with the ability to route around vulnerable parts of the network and may be used to increase overall network security.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.
- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.
- [RFC5441] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, April 2009.
- [RFC5520] Bradford, R., Ed., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.

7.2. Informative References

- [PCEP-MIB] Koushik, A. S. K., and E. Stephan, "PCE Communication Protocol(PCEP) Management Information Base", Work in Progress, November 2008.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004.

- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4657] Ash, J., Ed., and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.
- [RFC5298] Takeda, T., Ed., Farrel, A., Ed., Ikejiri, Y., and JP. Vasseur, "Analysis of Inter-Domain Label Switched Path (LSP) Recovery", RFC 5298, August 2008.

Acknowledgements

The authors would like to thank Fabien Verhaeghe for valuable comments on subobject formats. Thanks to Magnus Westerlund, Dan Romascanu, Tim Polk, and Dave Ward for comments during IESG review.

Authors' Addresses

Eiji Oki
University of Electro-Communications
1-5-1 Chofugaoka
Chofu, Tokyo 182-8585
JAPAN

E-Mail: oki@ice.uec.ac.jp

Tomonori Takeda
NTT
3-9-11 Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
E-Mail: takeda.tomonori@lab.ntt.co.jp

Adrian Farrel
Old Dog Consulting
E-Mail: adrian@olddog.co.uk