

Network Working Group
Request for Comments: 5520
Category: Standards Track

R. Bradford, Ed.
JP. Vasseur
Cisco Systems, Inc.
A. Farrel
Old Dog Consulting
April 2009

Preserving Topology Confidentiality in
Inter-Domain Path Computation Using a Path-Key-Based Mechanism

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.

Abstract

Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) Label Switched Paths (LSPs) may be computed by Path Computation Elements (PCEs). Where the TE LSP crosses multiple domains, such as Autonomous Systems (ASes), the path may be computed by multiple PCEs that cooperate, with each responsible for computing a segment of the path. However, in some cases (e.g., when ASes are administered by separate Service Providers), it would break confidentiality rules for a PCE to supply a path segment to a PCE in another domain, thus disclosing AS-internal topology information. This issue may be circumvented by returning a loose hop and by invoking a new path computation from the domain boundary Label Switching Router (LSR) during TE LSP setup as the signaling message enters the second domain, but this technique has several issues including the problem of maintaining path diversity.

This document defines a mechanism to hide the contents of a segment of a path, called the Confidential Path Segment (CPS). The CPS may be replaced by a path-key that can be conveyed in the PCE Communication Protocol (PCEP) and signaled within in a Resource Reservation Protocol TE (RSVP-TE) explicit route object.

Table of contents

1. Introduction	3
1.1. Terminology	4
2. Path-Key Solution	5
2.1. Mode of Operation	5
2.2. Example	6
3. PCEP Protocol Extensions	7
3.1. Path-Keys in PCRep Messages	7
3.1.1. PKS with 32-Bit PCE ID	8
3.1.2. PKS with 128-Bit PCE ID	9
3.2. Unlocking Path-Keys	10
3.2.1. Path-Key Bit	10
3.2.2. PATH-KEY Object	10
3.2.3. Path Computation Request (PCReq) Message with Path-Key	11
4. PCEP Mode of Operation for Path-Key Expansion	12
5. Security Considerations	12
6. Manageability Considerations	13
6.1. Control of Function through Configuration and Policy	13
6.2. Information and Data Models	14
6.3. Liveness Detection and Monitoring	15
6.4. Verifying Correct Operation	15
6.5. Requirements on Other Protocols and Functional Components	15
6.6. Impact on Network Operation	16
7. IANA Considerations	16
7.1. New Subobjects for the ERO Object	16
7.2. New PCEP Object	17
7.3. New RP Object Bit Flag	17
7.4. New NO-PATH-VECTOR TLV Bit Flag	17
8. References	17
8.1. Normative References	17
8.2. Informative References	18
Acknowledgements	19

1. Introduction

Path computation techniques using the Path Computation Element (PCE) are described in [RFC4655] and allow for path computation of inter-domain Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) and Generalized MPLS (GMPLS) Label Switched Paths (LSPs).

An important element of inter-domain TE is that TE information is not shared between domains for scalability and confidentiality reasons ([RFC4105] and [RFC4216]). Therefore, a single PCE is unlikely to be able to compute a full inter-domain path.

Two path computation scenarios can be used for inter-domain TE LSPs: one using per-domain path computation (defined in [RFC5152]), and the other using a PCE-based path computation technique with cooperation between PCEs (as described in [RFC4655]). In this second case, paths for inter-domain LSPs can be computed by cooperation between PCEs each of which computes a segment of the path across one domain. Such a path computation procedure is described in [RFC5441].

If confidentiality is required between domains (such as would very likely be the case between Autonomous Systems (ASes) belonging to different Service Providers), then cooperating PCEs cannot exchange path segments or else the receiving PCE and the Path Computation Client (PCC) will be able to see the individual hops through another domain thus breaking the confidentiality requirement stated in [RFC4105] and [RFC4216]. We define the part of the path that we wish to keep confidential as the Confidential Path Segment (CPS).

One mechanism for preserving the confidentiality of the CPS is for the PCE to return a path containing a loose hop in place of the segment that must be kept confidential. The concept of loose and strict hops for the route of a TE LSP is described in [RFC3209]. The Path Computation Element Communication Protocol (PCEP) defined in [RFC5440] supports the use of paths with loose hops, and it is a local policy decision at a PCE whether it returns a full explicit path with strict hops or uses loose hops. Note that a path computation request may request an explicit path with strict hops or may allow loose hops as detailed in [RFC5440].

The option of returning a loose hop in place of the CPS can be achieved without further extensions to PCEP or the signaling protocol. If loose hops are used, the TE LSPs are signaled as normal ([RFC3209]), and when a loose hop is encountered in the explicit route, it is resolved by performing a secondary path computation to reach the resource or set of resources identified by the loose hop. Given the nature of the cooperation between PCEs in computing the original path, this secondary computation occurs at or on behalf of a

Label Switching Router (LSR) at a domain boundary (i.e., an Area Border Router (ABR) or an AS Border Router (ASBR)) and the path is expanded as described in [RFC5152].

The PCE-based computation model is particularly useful for determining mutually disjoint inter-domain paths such as might be required for service protection [RFC5298]. A single path computation request is used. However, if loose hops are returned, the path of each TE LSP must be recomputed at the domain boundaries as the TE LSPs are signaled, and since the TE LSP signaling proceeds independently for each TE LSP, disjoint paths cannot be guaranteed since the LSRs in charge of expanding the explicit route objects (EROs) are not synchronized. Therefore, if the loose hop technique is used without further extensions, path segment confidentiality and path diversity are mutually incompatible requirements.

This document defines the notion of a Path-Key that is a token that replaces a path segment in an explicit route. The Path-Key is encoded as a Path-Key Subobject (PKS) returned in the PCEP Path Computation Reply message (PCRep) ([RFC5440]). Upon receiving the computed path, the PKS will be carried in an RSVP-TE Path message (RSVP-TE [RFC3209] and [RSVP-PKS]) during signaling.

The BNF in this document follows the format described in [RBNF].

Please note that the term "path-key" used in this document refers to an identifier allocated by a PCE to represent a segment of a computed path. This term has no relation to the term "cryptographic key" used in some documents that describe security mechanisms.

1.1. Terminology

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].

This document makes use of the following terminology and acronyms.

AS: Autonomous System.

ASBR: Autonomous System Border Routers used to connect to another AS of a different or the same Service Provider via one or more links inter-connecting between ASes.

CPS: Confidential Path Segment. A segment of a path that contains nodes and links that the AS policy requires to not be disclosed outside the AS.

Inter-AS TE LSP: A TE LSP that crosses an AS boundary.

LSR: Label Switching Router.

LSP: Label Switched Path.

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.

TE LSP: Traffic Engineering Label Switched Path.

2. Path-Key Solution

The Path-Key solution may be applied in the PCE-based path computation context as follows. A PCE computes a path segment related to a particular domain and replaces any CPS in the path reported to the requesting PCC (or another PCE) by one or more subobjects referred to as PKSes. The entry boundary LSR of each CPS SHOULD be specified using its TE Router Id as a hop in the returned path immediately preceding the CPS, and other subobjects MAY be included in the path immediately before the hop identifying the boundary LSR to indicate link and label choices. Where two PKSes are supplied in sequence with no intervening nodes, the entry node to the second CPS MAY be part of the first CPS and does not need to be explicitly present in the returned path. The exit node of a CPS MAY be present as a strict hop immediately following the PKS.

2.1. Mode of Operation

During path computation, when local policy dictates that confidentiality must be preserved for all or part of the path segment being computed or if explicitly requested by the path computation request, the PCE associates a path-key with the computed path for the CPS, places its own identifier (its PCE ID as defined in Section 3.1) along with the path-key in a PKS, and inserts the PKS object in the path returned to the requesting PCC or PCE immediately after the subobject that identifies (using the TE Router Id) the LSR that will expand the PKS into explicit path hops. This will usually be the LSR that is the starting point of the CPS. The PCE that generates a PKS SHOULD store the computed path segment and the path-key for later retrieval. A local policy SHOULD be used to determine for how long to retain such stored

information, and whether to discard the information after it has been queried using the procedures described below. It is RECOMMENDED for a PCE to store the PKS for a period of 10 minutes.

A path-key value is scoped to the PCE that computed it as identified by the PCE-ID carried in the PKS. A PCE MUST NOT re-use a path-key value to represent a new CPS for at least 30 minutes after discarding the previous use of the same path-key. A PCE that is unable to retain information about previously used path-key values over a restart SHOULD use some other mechanism to guarantee uniqueness of path-key values such as embedding a timestamp or version number in the path-key.

A head-end LSR that is a PCC converts the path returned by a PCE into an explicit route object (ERO) that it includes in the Resource Reservation Protocol (RSVP) Path message. If the path returned by the PCE contains a PKS, this is included in the ERO. Like any other subobjects, the PKS is passed transparently from hop to hop, until it becomes the first subobject in the ERO. This will occur at the start of the CPS, which will usually be the domain boundary. The PKS MUST be preceded by an ERO subobject that identifies the LSR that must expand the PKS. This means that (following the rules for ERO processing set out in [RFC3209]) the PKS will not be encountered in ERO processing until the ERO is being processed by the LSR that is capable of correctly handling the PKS.

An LSR that encounters a PKS when trying to identify the next hop retrieves the PCE-ID from the PKS and sends a Path Computation Request (PCReq) message as defined in [RFC5440] to the PCE identified by the PCE-ID that contains the path-key object .

Upon receiving the PCReq message, the PCE identifies the computed path segment using the supplied path-key, and returns the previously computed path segment in the form of explicit hops using an ERO object contained in the Path Computation Reply (PCRep) to the requesting node as defined in [RFC5440]. The requesting node inserts the explicit hops into the ERO and continues to process the TE LSP setup as per [RFC3209].

2.2. Example

Figure 1 shows a simple two-AS topology with a PCE responsible for the path computations in each AS. An LSP is requested from the ingress LSR in one AS to the egress LSR in the other AS. The ingress, acting as the PCC, sends a path computation request to PCE-1. PCE-1 is unable to compute an end-to-end path and invokes PCE-2 (possibly using the techniques described in [RFC5441]). PCE-2 computes a path segment from ASBR-2 to the egress as {ASBR-2, C, D,

Egress}. It could pass this path segment back to PCE-1 in full, or it could send back the path {ASBR-2, Egress} where the second hop is a loose hop.

However, in order to protect the confidentiality of the topology in the second AS while still specifying the path in full, PCE-2 may send PCE-1 a path segment expressed as {ASBR-2, PKS, Egress} where the PKS is a Path-Key Subobject as defined in this document. In this case, PCE-2 has identified the segment {ASBR-2, C, D, Egress} as a Confidential Path Segment (CPS). PCE-1 will compute the path segment that it is responsible for, and will supply the full path to the PCC as {Ingress, A, B, ASBR-1, ASBR-2, PKS, Egress}.

Signaling proceeds in the first AS as normal, but when the Path message reaches ASBR-2, the next hop is the PKS, and this must be expanded before signaling can progress further. ASBR-2 uses the information in the PKS to request PCE-2 for a path segment, and PCE-2 will return the segment {ASBR-2, C, D, Egress} allowing signaling to continue to set up the LSP.




Figure 1 : A Simple network to demonstrate the use of the PKS

3. PCEP Protocol Extensions

3.1. Path-Keys in PCRep Messages

Path-Keys are carried in PCReq and PCRep messages as part of the various objects that carry path definitions. In particular, a Path-Key is carried in the Explicit Route Object (ERO) on PCRep messages.

In all cases, the Path-Key is carried in a Path-Key Subobject (PKS).


The PKS is a fixed-length subobject containing a Path-Key and a PCE-ID. The Path-Key is an identifier, or token used to represent the CPS within the context of the PCE identified by the PCE-ID. The PCE-ID identifies the PCE that can decode the Path-Key using an identifier that is unique within the domain that the PCE serves. The PCE-ID has to be mapped to a reachable IPv4 or IPv6 address of the PCE by the first node of the CPS (usually a domain border router) and a PCE MAY use one of its reachable IP addresses as its PCE-ID. Alternatively and to provide greater security (see Section 5) or increased confidentiality, according to domain-local policy, the PCE MAY use some other identifier that is scoped only within the domain.

To allow IPv4 and IPv6 addresses to be carried, two subobjects are defined in the following subsections.

The Path-Key Subobject may be present in the PCEP ERO or the PCEP PATH-KEY object (see Section 3.2).

3.1.1.1. PKS with 32-Bit PCE ID

The Subobject Type for the PKS with 32-bit PCE ID is 64. The format of this subobject is as follows:



L

The L bit SHOULD NOT be set, so that the subobject represents a strict hop in the explicit route.

Type

Subobject Type for a Path-Key with 32-bit PCE ID (64).

Length

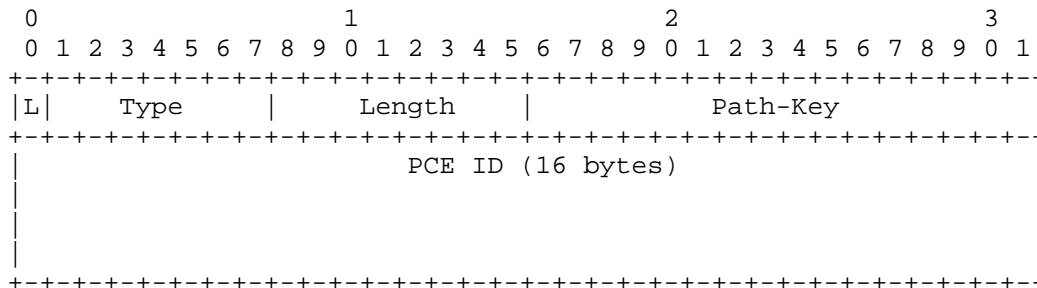
The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is always 8.

PCE ID

A 32-bit identifier of the PCE that can decode this path-key. The identifier MUST be unique within the scope of the domain that the CPS crosses, and MUST be understood by the LSR that will act as PCC for the expansion of the PKS. The interpretation of the PCE-ID is subject to domain-local policy. It MAY be an IPv4 address of the PCE that is always reachable, and MAY be an address that is restricted to the domain in which the LSR that is called upon to expand the CPS lies. Other values that have no meaning outside the domain (for example, the Router ID of the PCE) MAY be used to increase security or confidentiality (see Section 5).

3.1.2. PKS with 128-Bit PCE ID

The Subobject Type for the PKS with 128-bit PCE ID is 65. The format of the subobject is as follows.



L

As above.

Type

Subobject Type for a Path-Key with 128-bit PCE ID (65).

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is always 20.

PCE ID

A 128-bit identifier of the PCE that can decode this path-key. The identifier MUST be unique within the scope of the domain that the CPS crosses, and MUST be understood by the LSR that

will act as PCC for the expansion of the PKS. The interpretation of the PCE-ID is subject to domain-local policy. It MAY be an IPv6 address of the PCE that is always reachable, but MAY be an address that is restricted to the domain in which the LSR that is called upon to expand the CPS lies. Other values that have no meaning outside the domain (for example, the IPv6 TE Router ID) MAY be used to increase security (see Section 5).

3.2. Unlocking Path-Keys

When a network node needs to decode a Path-Key so that it can continue signaling for an LSP, it must send a PCReq to the designated PCE. The PCReq defined in [RFC5440] needs to be modified to support this usage, which differs from the normal path computation request. To that end, a new flag is defined to show that the PCReq relates to the expansion of a PKS, and a new object is defined to carry the PKS in the PCReq. These result in an update to the BNF for the message. The BNF used in this document is as described in [RBNF].

3.2.1. Path-Key Bit

[RFC5440] defines the Request Parameters (RP) object that is used to specify various characteristics of the Path Computation Request (PCReq).

In this document, we define a new bit named the Path-Key bit as follows. See Section 7.3 for the IANA assignment of the appropriate bit number.

Path-Key bit: When set, the requesting PCC requires the retrieval of a Confidential Path Segment that corresponds to the PKS carried in a PATH-KEY object in the path computation request. The Path-Key bit MUST be cleared when the path computation request is not related to a CPS retrieval.

3.2.2. PATH-KEY Object

When a PCC needs to expand a path-key in order to expand a CPS, it issues a Path Computation Request (PCReq) to the PCE identified in the PKS in the RSVP-TE ERO that it is processing. The PCC supplies the PKS to be expanded in a PATH-KEY Object in the PCReq message.

The PATH-KEY Object is defined as follows:

PATH-KEY Object-Class is 16.

Path-Key Object-Type is 1.

The PATH-KEY Object MUST contain at least one Path-Key Subobject (see Section 3.1). The first PKS MUST be processed by the PCE. Subsequent subobjects SHOULD be ignored.

3.2.3. Path Computation Request (PCReq) Message with Path-Key

The format of a PCReq message including a PATH-KEY object is unchanged as follows:

```
<PCReq Message> ::= <Common Header>
                    [<SVEC-list>]
                    <request-list>
```

where:

```
<svec-list> ::= <SVEC> [<svec-list>]
<request-list> ::= <request> [<request-list>]
```

To support the use of the message to expand a PKS, the definition of <request> is modified as follows :

```
<request> ::= <RP>
              <segment-computation> | <path-key-expansion>
```

where:

```
<segment-computation> ::= <END-POINTS>
                          [<LSPA>]
                          [<BANDWIDTH>]
                          [<BANDWIDTH>]
                          [<metric-list>]
                          [<RRO>]
                          [<IRO>]
                          [<LOAD-BALANCING>]
<path-key-expansion> ::= <PATH-KEY>
```

Thus, the format of the message for use in normal path computation is unmodified.

4. PCEP Mode of Operation for Path-Key Expansion

The retrieval of the explicit path (the CPS) associated with a PKS by a PCC is no different than any other path computation request with the exception that the PCReq message MUST contain a PATH-KEY object and the Path-Key bit of the RP object MUST be set. On receipt of a PCRep containing a CPS, the requesting PCC SHOULD insert the CPS into the ERO that it will signal, in accordance with local policy.

If the receiving PCE does not recognize itself as identified by the PCE ID carried in the PKS, it MAY forward the PCReq message to another PCE according to local policy. If the PCE does not forward such a PCReq, it MUST respond with a PCRep message containing a NO-PATH object.

If the receiving PCE recognizes itself, but cannot find the related CPS, or if the retrieval of the CPS is not allowed by policy, the PCE MUST send a PCRep message that contains a NO-PATH object. The NO-PATH-VECTOR TLV SHOULD be used as described in [RFC5440] and a new bit number (see Section 7.4) is assigned to indicate "Cannot expand PKS".

Upon receipt of a negative reply, the requesting LSR MUST fail the LSP setup and SHOULD use the procedures associated with loose hop expansion failure [RFC3209].

5. Security Considerations

This document describes tunneling confidential path information across an untrusted domain (such as an AS). There are many security considerations that apply to PCEP and RSVP-TE.

Issues include:

- Confidentiality of the CPS (can other network elements probe for expansion of path-keys, possibly at random?).
- Authenticity of the path-key (resilience to alteration by intermediaries, resilience to fake expansion of path-keys).
- Resilience from Denial-of-Service (DoS) attacks (insertion of spurious path-keys; flooding of bogus path-key expansion requests).

Most of the interactions required by this extension are point to point, can be authenticated and made secure as described in [RFC5440] and [RFC3209]. These interactions include the:

- PCC->PCE request
- PCE->PCE request(s)
- PCE->PCE response(s)
- PCE->PCC response
- LSR->LSR request and response. Note that a rogue LSR could modify the ERO and insert or modify Path-Keys. This would result in an LSR (which is downstream in the ERO) sending decode requests to a PCE. This is actually a larger problem with RSVP. The rogue LSR is an existing issue with RSVP and will not be addressed here.
- LSR->PCE request. Note that the PCE can check that the LSR requesting the decode is the LSR at the head of the Path-Key. This largely contains the previous problem of DoS rather than a security issue. A rogue LSR can issue random decode requests, but these will amount only to DoS.
- PCE->LSR response

Thus, the major security issues can be dealt with using standard techniques for securing and authenticating point-to-point communications. In addition, it is recommended that the PCE providing a decode response should check that the LSR that issued the decode request is the head end of the decoded ERO segment.

Further protection can be provided by using a PCE ID to identify the decoding PCE that is only meaningful within the domain that contains the LSR at the head of the CPS. This may be an IP address that is only reachable from within the domain, or some not-address value. The former requires configuration of policy on the PCEs, the latter requires domain-wide policy.

6. Manageability Considerations

6.1. Control of Function through Configuration and Policy

The treatment of a path segment as a CPS, and its substitution in a PCRep ERO with a PKS, is a function that MUST be under operator and policy control where a PCE supports the function. The operator MUST be given the ability to specify which path segments are to be replaced and under what circumstances. For example, an operator might set a policy that states that every path segment for the operator's domain will be replaced by a PKS when the PCReq has been issued from outside the domain.

The operation of the PKS extensions require that path-keys are retained by the issuing PCE to be available for retrieval by an LSR (acting as a PCC) at a later date. But it is possible that the retrieval request will never be made, so good housekeeping requires that a timer is run to discard unwanted path-keys. A default value for this timer is suggested in Section 2.1. Implementations SHOULD provide the ability for this value to be overridden through operator configuration or policy.

After a PKS has been expanded in response to a retrieval request, it may be valuable to retain the path-key and CPS for debugging purposes. Such retention SHOULD NOT be the default behavior of an implementation, but MAY be available in response to operator request.

Once a path-key has been discarded, the path-key value SHOULD NOT be immediately available for re-use for a new CPS since this might lead to accidental misuse. A default timer value is suggested in Section 2.1. Implementations SHOULD provide the ability for this value to be overridden through operator configuration or policy.

A PCE must set a PCE-ID value in each PKS it creates so that PCCs can correctly identify it and send PCReq messages to expand the PKS to a path segment. A PCE implementation SHOULD allow operator or policy control of the value to be used as the PCE-ID. If the PCE allows PCE-ID values that are not routable addresses to be used, the PCCs MUST be configurable (by the operator or through policy) to allow the PCCs to map from the PCE-ID to a routable address of the PCE. Such mapping may be algorithmic, procedural (for example, mapping a PCE-ID equal to the IGP Router ID into a routable address), or configured through a local or remote mapping table.

6.2. Information and Data Models

A MIB module for PCEP is already defined in [PCEP-MIB]. The configurable items listed in Section 6.1 MUST be added as readable objects in the module and SHOULD be added as writable objects.

A new MIB module MUST be created to allow inspection of path-keys. For a given PCE, this MIB module MUST provide a mapping from path-key to path segment (that is, a list of hops), and MUST supply other information including:

- The identity of the PCC that issued the original request that led to the creation of the path-key.
- The request ID of the original PCReq.

- Whether the path-key has been retrieved yet, and if so, by which PCC.
- How long until the path segment associated with the path-key will be discarded.
- How long until the path-key will be available for re-use.

6.3. Liveness Detection and Monitoring

The procedures in this document extend PCEP, but do not introduce new interactions between network entities. Thus, no new liveness detection or monitoring is required.

It is possible that a head-end LSR that has been given a path including PKSs replacing specific CPSs will want to know whether the path-keys are still valid (or have timed out). However, rather than introduce a mechanism to poll the PCE that is responsible for the PKS, it is considered pragmatic to simply signal the associated LSP.

6.4. Verifying Correct Operation

The procedures in this document extend PCEP, but do not introduce new interactions between network entities. Thus, no new tools for verifying correct operation are required.

A PCE SHOULD maintain counters and logs of the following events that might indicate incorrect operation (or might indicate security issues).

- Attempts to expand an unknown path-key.
- Attempts to expand an expired path-key.
- Duplicate attempts to expand the same path-key.
- Expiry of path-key without attempt to expand it.

6.5. Requirements on Other Protocols and Functional Components

The procedures described in this document require that the LSRs signal PKSs as defined in [RSVP-PKS]. Note that the only changes to LSRs are at the PCCs. Specifically, changes are only needed at the head-end LSRs that build RSVP-TE Path messages containing Path-Key Subobjects in their EROs, and the LSRs that discover such subobjects as next hops and must expand them. Other LSRs in the network, even if they are on the path of the LSP, will not be called upon to process the PKS.

6.6. Impact on Network Operation

As well as the security and confidentiality aspects addressed by the use of the PKS, there may be some scaling benefits associated with the procedures described in this document. For example, a single PKS in an explicit route may substitute for many subobjects and can reduce the overall message size correspondingly. In some circumstances, such as when the explicit route contains multiple subobjects for each hop (including node IDs, TE link IDs, component link IDs for each direction of a bidirectional LSP, and label IDs for each direction of a bidirectional LSP) or when the LSP is a point-to-multipoint LSP, this scaling improvement may be very significant.

Note that a PCE will not supply a PKS unless it knows that the LSR that will receive the PKS through signaling will be able to handle it. Furthermore, as noted in Section 6.5, only those LSRs specifically called upon to expand the PKS will be required to process the subobjects during signaling. Thus, the only backward compatibility issues associated with the procedures introduced in this document arise when a head-end LSR receives a PCRep with an ERO containing a PKS, and it does not know how to encode this into signaling.

Since the PCE that inserted the PKS is required to keep the CPS confidential, the legacy head-end LSR cannot be protected. It must either fail the LSP setup, or request a new path computation avoiding the domain that has supplied it with unknown subobjects.

7. IANA Considerations

IANA assigns values to PCEP parameters in registries defined in [RFC5440]. IANA has made the following additional assignments.

7.1. New Subobjects for the ERO Object

IANA has previously assigned an Object-Class and Object-Type to the ERO carried in PCEP messages [RFC5440]. IANA also maintains a list of subobject types valid for inclusion in the ERO.

IANA assigned two new subobject types for inclusion in the ERO as follows:

Subobject Type	Type	Reference
64	Path-Key with 32-bit PCE ID	[RFC5520]
65	Path-Key with 128-bit PCE ID	[RFC5520]

7.2. New PCEP Object

IANA assigned a new object class in the registry of PCEP Objects as follows.

Object Class	Name	Object Type	Name	Reference
16	PATH-KEY	1	Path-Key	[RFC5520]

Subobjects

This object may carry the following subobjects as defined for the ERO object.

64	Path-Key with 32-bit PCE ID	[RFC5520]
65	Path-Key with 128-bit PCE ID	[RFC5520]

7.3. New RP Object Bit Flag

IANA maintains a registry of bit flags carried in the PCEP RP object as defined in [RFC5440]. IANA assigned a new bit flag as follows:

Bit Number	Hex	Name	Reference
23	0x000017	Path-Key (P-bit)	[RFC5520]

7.4. New NO-PATH-VECTOR TLV Bit Flag

IANA maintains a registry of bit flags carried in the PCEP NO-PATH-VECTOR TLV in the PCEP NO-PATH object as defined in [RFC5440]. IANA assigned a new bit flag as follows:

Bit Number	Name Flag	Reference
27	PKS expansion failure	[RFC5520]

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

8.2. Informative References

- [PCEP-MIB] Koushik, K., and E. Stephan, "PCE Communication Protocol (PCEP) Management Information Base", Work in Progress, November 2008.
- [RBNF] Farrel, A., "Reduced Backus-Naur Form (RBNF) A Syntax Used in Various Protocol Specifications", Work in Progress, November 2008.
- [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.
- [RFC4105] Le Roux, J.-L., Ed., Vasseur, J.-P., Ed., and J. Boyle, Ed., "Requirements for Inter-Area MPLS Traffic Engineering", RFC 4105, June 2005.
- [RFC4216] Zhang, R., Ed., and J.-P. Vasseur, Ed., "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [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.
- [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.
- [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.
- [RSVP-PKS] Bradford, R., Vasseur, JP., and A. Farrel, "RSVP Extensions for Path Key Support", Work in Progress, February 2008.

Acknowledgements

The authors would like to thank Eiji Oki, Ben Campbell, and Ross Callon for their comments on this document.

Authors' Addresses

Rich Bradford (Editor)
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA
EMail: rbradfor@cisco.com

JP. Vasseur
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA
EMail: jpv@cisco.com

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