Networking Working Group
IETF Internet Draft

Rich Bradford (Ed)
JP Vasseur
Cisco Systems, Inc.
Adrian Farrel
Old Dog Consulting

Proposed Status: Standard
Expires: July 3, 2007

January 3, 2007

draft-bradford-pce-path-key-02.txt

Preserving Topology Confidentiality in Inter-Domain Path Computation using a key based mechanism

Status of this Memo

Abstract

Multiprotocol Label Switching (MPLS) 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 (ASs), 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 ASs 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 internal topology information. This issue may be circumvented by returning a loose hop and by invoking a new path computation from the domain boundary LSR during TE LSP setup as the LSP 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 (RSVP) explicit route object.

Table of contents

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].

1. Introduction

Path computation techniques using the Path Computation Element (PCE) have been described in [PCE-ARCH] and allow for path computation of inter-domain Multiprotocol Label Switching (MPLS) traffic engineering (TE) 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 [PD-PATH-COMP]), and the other using a PCE-based path computation technique with cooperation between PCEs (as described in [PCE-ARCH]). 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 [BRPC].

If confidentiality is required between domains (such as would very likely be the case between ASs 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 non conforming to the confidentiality requirement stated in [RFC4105] and [RFC4216]. We define the part of the path which 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 for the segment internal to a domain that must be kept confidential. The concept of loose hops for the route of a TE LSP is described in [RFC3209].

The Path Computation Element Communication Protocol (PCEP) defined in [PCEP] 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 or uses loose hops. Note that a Path computation Request may request a loose or explicit path as detailed in [PCEP].

One option may consist of returning loose hop without further extensions: 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 next loose hop. Given the nature of the cooperation between PCEs in computing the original path, this secondary computation occurs at a Label Switching Router (LSR) at a domain boundary (i.e. an ABR or ASBR) and the path is expanded as described in [PD-PATH-COMP].

The PCE-based computation model is particularly useful for determining mutually disjoint inter-domain paths such as might be required for service protection. 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 EROs are not synchronized. Therefore, using the loose hop technique 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 Sub-object (PKS) returned in the PCEP Path Computation Reply message (PCReq) ([PCEP]). Upon receiving the computed path, the PKS sub-object will be carried out in an RSVP-TE Path message (RSVP-TE [RFC3209]) during signaling. The PKS may also, optionally, be used in recorded routes in RSVP-TE.

2. Terminology

ASBR: border routers used to connect to another AS of a different or the same Service Provider via one or more links inter-connecting between ASs.

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

3. 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 it in the path reported to the requesting PCC (or another PCE) by one or more sub-objects referred to as the PKS. The entry and boundary LSR of each CPS SHOULD be specified as hops in the returned path immediately preceding the PKS, but where two PKSs are supplied in sequence the entry node to the second MAY be encoded within the first. The exit node of a CPS MAY be present as a strict hop immediately following the PKS, but MAY also be hidden as part of the PKS.

3.1. 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 [PCE-MONITORING]) 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 IPv4 sub-object defined in [RFC3209]sub-object that identifies the LSR that will expand the PKS into a explicit path hops. This will usually be the LSR that is the start point of the CPS. The PCE that generates a PKS MUST 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 strore the PKS for a period of 10 minutes.

TBD: Need to define the scope of the PKS and spell out the restrictions on Path Key re-use.

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 PKSs these are included in the ERO. Like any other sub-objects, the PKS is passed transparently from hop to hop, until it becomes the first sub-object 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 sub-object that identifies the LSR that must expand the PKS, so the PKS will not be encountered in ERO processing until the LSR that can process it.

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 [PCEP] 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 (PCReqp) as define in [PCEP] to the requesting node. The requesting node inserts the explicit hops into the ERO and continues to process the TE LSP setup as per [RFC3209].

4. PCEP Protocol Extensions

4.1. PKS sub-object

The PKS object format is identical as in [RSVP-PKS] but redefined in the context of this document since a PCEP codepoint is required.

The PKS is a fixed-length sub-object 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 a reachable IPv4 or IPv6 address of the PCE.

Because of the IPv4 and IPv6 variants, two sub-objects are defined as follows.

PKS IPv4 sub-object

PKS Object-Class is to be assigned by IANA (recommended value=16)

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

```
  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
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |L|    Type     |     Length    |           Path Key            |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                      IPv4 address (4 bytes)                   |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L

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

Type

TBD Path Key with IPv4 address

Length

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

IPv4 address

An IPv4 address of the PCE that can decode this key. The address used SHOULD be an 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.

PKS IPv6 sub-object

PKS Object-Class is to be assigned by IANA (recommended value=16)

PKS Object-Type is to be assigned by IANA (recommended value=2)

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    |           Path Key            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    IPv6 address (16 bytes)                    |
|                                                               |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L

As above.

Type

TBD Path Key with IPv6 address

Length

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

IPv6 address

An IPv6 address of the PCE that can decode this key. The address used SHOULD be an 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.

4.2. PKS bit

[PCEP] specifies the RP object that is used to specify various characteristics of the path computation request.

In this document we define a new bit named the PKS bit defined as follow:

PKS (PKS - 1 bit - Value=0x60): when set, the requesting PCC requires the retrieval of a strict path segment that corresponds to a PKS carried within the path computation request. The PKS bit MUST be cleared when the path computation request is not related to a CPS retrieval.

5. PCEP Mode of operation

The retrieval of the explicit path 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 PKS object and the PKS bit of the RP object MUST the set.

If the receiving PCE cannot find any related strict path or the retrieval of such strict path is not allowed by policy, the PCE MUST send a PCRep message that contains a NO-PATH object.

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

6. Security Considerations

This document proposes tunneling secure topology information across an untrusted AS, so the security considerations are many and apply to PCEP and RSVP-TE.
Issues include:

- Security 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 DNS 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. 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 to 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 pt-pt links. 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.

7. Manageability Considerations

To be detailed in a further revision of this document.

8. IANA considerations

IANA assigns value to PCEP parameters. Each PCEP object has an Object-Class and an Object-Type.

Two new PCEP objects are defined in this document: the IPv4 PKS and the IPv6 PKS objects.

```
 Object-Class      Name

      16           PKS IPv4
                 Object-Type
                     1

      16           PKS IPv6
                 Object-Type
```



```
                     2
```

9. Intellectual Property Considerations

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.

10. References

10.1. 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.

[PCEP] Vasseur, J.P., Le Roux, J.L., Ayyangar, A., Oki, E., Ikejiri, A., Atlas, A., Dolganow, A., "Path Computation Element (PCE) communication Protocol (PCEP)", draft-vasseur-pce-pcep, work in progress.

[RSVP-PKS] Bradford, R., Vasseur, J.P., Farrel, A., "RSVP Extensions for Path Key Support", draft-bradford-ccamp-path-key-ero, work in progress.

[PCE-MONITORING] Vasseur, J.P, "A set of monitoring tools for Path Computation Element based Architecture", draft-vasseur-pce-monitoring, work in progress.

10.2. Informational References

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

[PD-PATH-COMP] Vasseur, J., et al "A Per-domain path computation method for establishing Inter-domain Traffic Engineering (TE) Label Switched Paths (LSPs)", draft-ietf-ccamp-inter-domain-pd-path-comp, work in progress.

[BRPC] Vasseur, J., et al "A Backward Recursive PCE-based Computation (BRPC) procedure to compute shortest inter-domain Traffic Engineering Label Switched Path", draft-ietf-pce-brpc, work in progress.

[RFC4105] Le Roux, J., Vasseur, JP, Boyle, J., "Requirements for Support of Inter-Area and Inter-AS MPLS Traffic Engineering", RFC 4105, June 2005.

[RFC4216] Zhang, R., Vasseur, JP., et. al., "MPLS Inter-AS Traffic Engineering requirements", RFC 4216, November 2005.

11. Authors' Addresses

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

J.-P 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

Intellectual Property