Network Working Group
Request for Comments: 5553
Category: Standards Track

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

May 2009

Resource Reservation Protocol (RSVP) Extensions for Path Key Support

Status of This Memo

Copyright Notice

Abstract

The paths taken by 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.

To preserve confidentiality of topology within each AS, the PCEs support a mechanism to hide the contents of a segment of a path (such as the segment of the path that traverses an AS), called the Confidential Path Segment (CPS), by encoding the contents as a Path Key Subobject (PKS) and embedding this subobject within the result of its path computation.

This document describes how to carry Path Key Subobjects in the Resource Reservation Protocol (RSVP) Explicit Route Objects (EROs) and Record Route Objects (RROs) so as to facilitate confidentiality in the signaling of inter-domain TE LSPs.

1. Introduction

Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) Label Switched Paths (LSPs) are signaled using the TE extensions to the Resource Reservation Protocol (RSVP-TE) [RFC3209], [RFC3473]. The routes followed by MPLS and GMPLS TE LSPs may be computed by Path Computation Elements (PCEs) [RFC4655].

Where the TE LSP crosses multiple domains [RFC4726], 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. To preserve confidentiality of topology with each AS, the PCE Communications Protocol (PCEP) [RFC5440] supports a mechanism to hide the contents of a segment of a path, called the Confidential Path Segment (CPS), by encoding the contents as a Path Key Subobject (PKS) [RFC5520].

This document defines RSVP-TE protocol extensions necessary to support the use of Path Key Subobjects in MPLS and GMPLS signaling by including them in Explicit Route Objects (EROs) and Record Route Object (RROs) so as to facilitate confidentiality in the signaling of inter-domain TE LSPs.

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].

1.2. Usage Scenario

Figure 1 shows a simple network constructed of two ASes. An LSP is desired from the ingress in AS-1 to the egress in AS-2. As described in [RFC4655], the ingress Label Switching Router (LSR) acts as a Path Computation Client (PCC) and sends a request to its PCE (PCE-1). PCE-1 can compute the path within AS-1 but has no visibility into AS-2. So PCE-1 cooperates with PCE-2 to complete the path computation.

However, PCE-2 does not want to share the information about the path across AS-2 with nodes outside the AS. So, as described in [RFC5520], PCE-2 reports the AS-2 path segment using a PKS rather than the explicit details of the path.

PCE-1 can now return the path to be signaled to the ingress LSR in a path computation response with the AS-2 segment still hidden as a PKS.

In order to set up the LSP, the ingress LSR signals using RSVP-TE and encodes the path reported by PCE-1 in the Explicit Route Object (ERO). This process is as normal for RSVP-TE but requires that the PKS is also included in the ERO, using the mechanisms defined in this document.

When the signaling message (the RSVP-TE Path message) reaches ASBR-2 (Autonomous System Border Router), it consults PCE-2 to ’decode’ the PKS and return the expanded explicit path segment to ASBR-2. (The information that PCE-2 uses to decode the PKS is encoded within the PKS itself.) The PKS is replaced in the ERO with the expanded information about the path.

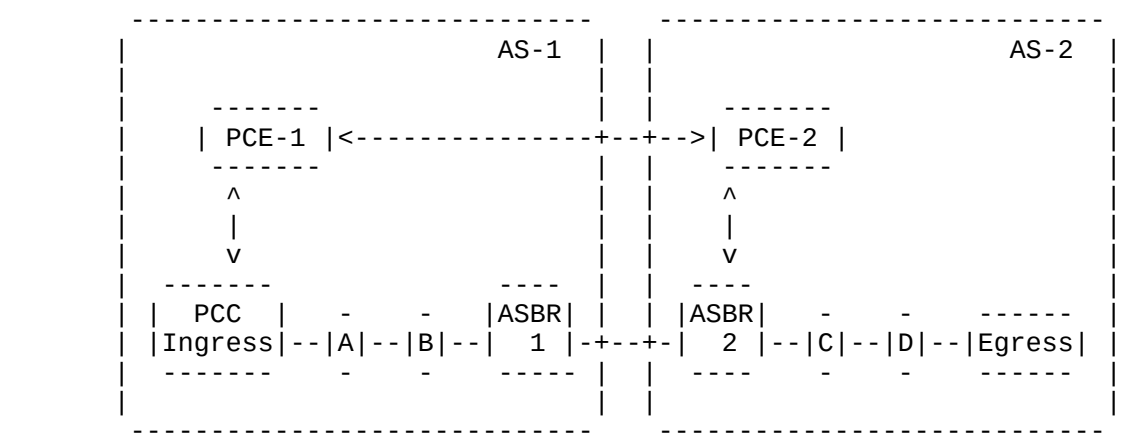

Figure 1: A Simple Network to Demonstrate the Use of the PKS

Note that PCE-2 may in some case be co-located with ASBR-2.

2. Terminology

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.

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.

PKS: Path Key Subobject. A subobject of an Explicit Route Object that encodes a CPS so as to preserve confidentiality.

3. RSVP-TE Path Key Subobject

The Path Key Subobject (PKS) may be carried in the Explicit Route Object (ERO) of an RSVP-TE Path message [RFC3209]. 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 a reachable IPv4 or IPv6 address of the PCE. In most cases, the decoding PCE is also the PCE that computed the Path Key and the associated path. Because of the IPv4 and IPv6 variants, two subobjects are defined as follows.

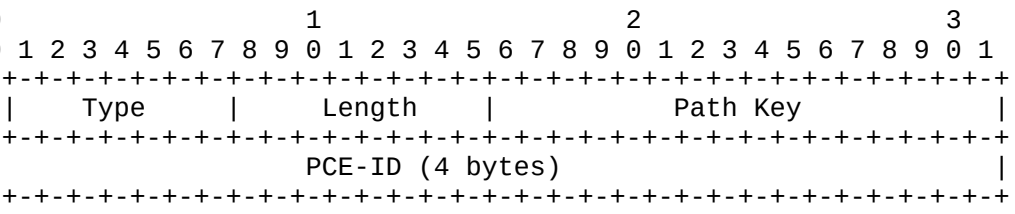

Figure 2: RSVP-TE Path Key Subobject using an IPv4 address for the PCE-ID

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 a 32-bit PCE-ID as assigned by IANA.

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 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.

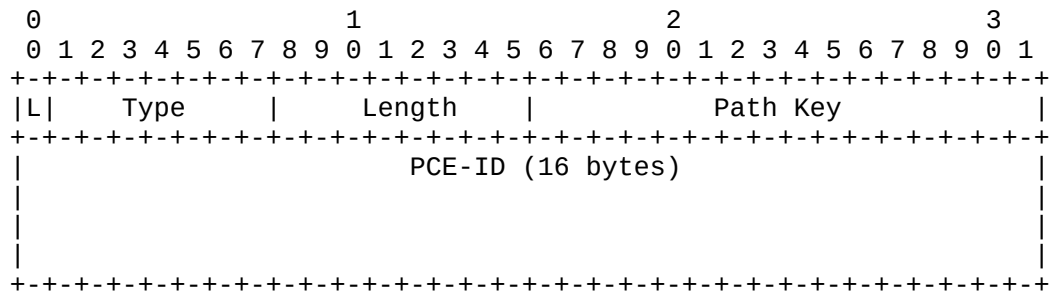

Figure 3: RSVP-TE Path Key Subobject using an IPv6 address for the PCE-ID

L

As above.

Type

Subobject Type for a Path Key with a 128-bit PCE-ID as assigned by IANA.

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 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, 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 IPv6 TE Router ID) MAY be used to increase security (see Section 4).

Note: The twins of these subobjects are carried in PCEP messages as defined in [RFC5520].

3.1. Explicit Route Object Processing Rules

The basic processing rules of an ERO are not altered. Refer to [RFC3209] for details. In particular, an LSR is not required to "look ahead" in the ERO beyond the first subobject that is non-local.

[RFC5520] requires that any path fragment generated by a PCE that contains a PKS be such that the PKS is immediately preceded by a subobject that identifies the head end of the PKS (for example, an incoming interface or a node ID). This rule is extended to the PKS in the ERO so that the following rules are defined.

- If an LSR receives a Path message where the first subobject of the ERO is a PKS, it MUST respond with a PathErr message carrying the error code/value combination "Routing Problem" / "Bad initial subobject".

- If an LSR strips all local subobjects from an ERO carried in a Path message (according to the procedures in [RFC3209]) and finds that the next subobject is a PKS, it MUST attempt to resolve the PKS to a CPS.

  Resolution of the PKS MAY take any of the following forms or use some other technique subject to local policy and network implementation.

  o The LSR can use the PCE-ID contained in the PKS to contact the identified PCE using PCEP [RFC5440] and request that the PKS be expanded.

  o The LSR can contact any PCE using PCEP [RFC5440] to request that the PKS be expanded, relying on cooperation between the PCEs.

  o The LSR can use the information in the PKS to index a CPS previously supplied to it by the PCE that originated the PKS.

  If a CPS is derived, the path fragment SHOULD be inserted into the ERO of the Path message as a direct replacement for the PKS. Other processing of the CPS and ERO are permitted as described in [RFC3209].

  This processing can give rise to the following error cases:

  o PCE-ID cannot be matched to a PCE to decode the PKS.

    The LSR sends a PathErr message with the error code "Routing Problem" and the new error value "Unknown PCE-ID for PKS expansion" (see Section 6.3).

- o PCE identified by the PCE-ID cannot be reached.

  The LSR sends a PathErr message with the error code "Routing Problem" and the new error value "Unreachable PCE for PKS expansion" (see Section 6.3).

- o The PCE is unable to decode the PKS, perhaps because the Path Key has expired.

  The LSR sends a PathErr message with the error code "Routing Problem" and the new error value "Unknown Path Key for PKS expansion" (see Section 6.3).

- o PKS cannot be decoded for policy reasons.

  The LSR sends a PathErr message with the error code "Policy Control Failure" and the error value "Inter-domain policy failure".

- o Addition of CPS to ERO causes Path message to become too large.

  The LSR MAY replace part of the ERO with loose hops [RFC3209] or with a further PKS, according to local policy, if the loss of specifics within the explicit path is acceptable. If the LSR is unable to take steps to reduce the size of the ERO, it MUST send a PathErr message with the error code "Routing Problem" and the new error value "ERO too large for MTU" (see Section 6.3).

- An LSR that is called on to process a PKS within an ERO but that does not recognize the subobject, will react according to [RFC3209] and send a PathErr message with the error code/value combination "Routing Problem" / "Bad Explicit Route Object".

3.2. Reporting Path Key Segments in Record Route Objects

The Record Route Object (RRO) is used in RSVP-TE to record the route traversed by an LSP. The RRO may be present on a Path message and on a Resv message. The intention of [RFC3209] is that an RRO on a Resv message that is received by an ingress LSR is suitable for use as an ERO on a Path message sent by that LSR to achieve an identical LSP.

The PKS offers an alternative that can be more useful to diagnostics. When the signaling message crosses a domain boundary, the path segment that needs to be hidden (that is, a CPS) MAY be replaced in the RRO with a PKS. In the case of an RRO on a Resv message, the PKS used SHOULD be the one originally signaled in the ERO of the Path message. On a Path message, the PKS SHOULD identify the LSR replacing the CPS and provide a Path Key that can be used to expand

the path segment. In the latter case, the Path Key and its expansion SHOULD be retained by the LSR that performs the substitution for at least the lifetime of the LSP. In both cases, the expansion of the PKS SHOULD be made available to diagnostic tools under the control of local policy.

4. Security Considerations

The protocol interactions required by the mechanisms described in this document are point-to-point and can be authenticated and made secure as described in [RFC5440] and [RFC3209]. The protocol interactions for PCEP are listed in [RFC5520], while general considerations for securing RSVP-TE in MPLS-TE and GMPLS networks can be found in [MPLS-SEC].

Thus, 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 PKS expansion check that the LSR that issued the request for PKS expansion is the head end of the resulting CPS.

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 either an IP address that is only reachable from within the domain or some non-address value. The former requires configuration of policy on the PCEs; the latter requires domain-wide policy.

The following specific security issues need to be considered.

- Confidentiality of the CPS. The question to be answered is whether other network elements can probe a PCE for the expansion of PKSs, possibly generating Path Keys at random. This can be protected against by only allowing PKS expansion to be successfully completed if requested by the LSR that is at the head end of the resulting CPS. Under specific circumstances, PKS expansion might also be allowed by configured management stations.

  The CPS itself may be kept confidential as it is exchanged in the PCEP and RSVP-TE protocols using standard security mechanisms defined for those protocols.

- Determination of information by probing. In addition to the probing described above, a node might deduce information from the error responses that are generated when PKS expansion fails as described in Section 3.1. Any LSR that determines that supplying one of the detailed error codes described in Section 3.1 might

provide too much information that could be used as part of a systematic attack MAY simply use the error code/value "Policy Control Failure" / "Inter-domain policy failure" in all cases.

- Authenticity of the Path Key. A concern is that the Path Key in the PKS will be altered or faked, leading to erroneous Path Key expansion and use of the wrong CPS. The consequence would be a bad ERO in a Path message, causing the LSP to be set up incorrectly and resulting in incorrect network resource usage, diversion of traffic to where it can be intercepted, or failure to set up the LSP. These problems can be prevented by protecting the protocol exchanges in PCEP and RSVP-TE using the security techniques described in [RFC5440], [RFC3209], and [MPLS-SEC].

- Resilience to denial-of-service (DoS) attacks. A PCE can be attacked through a flood of Path Key expansion requests -- this issue is addressed in [RFC5520] and is out of scope for this document. A further attack might consist of sending a flood of RSVP-TE Path messages with deliberately spurious PKSs. This attack is prevented by ensuring the integrity of the Path messages using standard RSVP-TE security mechanisms and by enforcing the RSVP-TE chain-of-trust security model.

5. Manageability Considerations

5.1. Control of Function through Configuration and Policy

Policy forms an important part of the use of PKSs in EROs and RROs. There are local and domain-wide policies that SHOULD be available for configuration in an implementation.

- Handling of an ERO containing a PKS. As described in Section 3.1, an LSR that receives a Path message containing a PKS can be configured to reject the Path message according to policy.

- Handling of PKS requests at a PCE. As described in Section 3.1, in [RFC5520], and in [RFC5394], a PCE can be configured with policy regarding how it should handle requests for PKS expansion.

- PKS expansion. Section 3.1 explains that the PKS can be expanded by the local LSR, the specific PCE identified in the PKS, any PCE acting as a proxy, or by some other method. The behavior of the LSR needs to be locally configurable but is subject to the domain-wide policy.

- Interpretation of PCE-ID. The interpretation of the PCE-ID component of PKSs is subject to domain-local policy and needs to be configurable as such. See Section 3 and Section 4 for the options.

- ERO too large. The behavior of an LSR when it finds that adding a CPS to the ERO causes the Path message to be too large is an implementation choice. However, implementations may choose to provide configuration of behavior as described in Section 3.1.

- Masking of RRO. As described in Section 3.2, a border router can choose to mask segments of the path by replacing them with PKSs. This behavior needs to be configurable, with the default being to not hide any part of the RRO.

- Inspection / decoding of PKS by diagnostic tools. A PCE can allow access from management or diagnostic tools to request the expansion of a PKS. Note that this must be regulated with the security and confidentiality behavior described in Section 4.

- Hiding of reason codes. An LSR can support the configuration of local policy to hide reason codes associated with the failure to expand a PKS and, as described in Section 4, report all errors as policy failures.

The treatment of a path segment as a CPS, and its substitution in a PCRep ERO with a PKS, is a PCE function and is described in [RFC5520].

6. IANA Considerations

6.1. Explicit Route Object Subobjects

IANA maintains a registry called "Resource Reservation Protocol (RSVP) Parameters" with a subregistry called "Class Names, Class Numbers, and Class Types".

Within this subregistry, there is a definition of the EXPLICIT_ROUTE object with Class Number 20. The object definition lists a number of acceptable subobjects for the Class Type 1.

IANA has allocated two further subobjects as described in Section 3. The resulting entry in the registry is as follows.

20	EXPLICIT_ROUTE	[RFC3209]
	Class Types or C-Types:	
1	Type 1 Explicit Route	[RFC3209]
	Subobject type	
64	Path Key with 32-bit PCE-ID	[RFC5553]
65	Path Key with 128-bit PCE-ID	[RFC5553]

Note well: [RFC5520] defines the PKS for use in PCEP. IANA has assigned the same subobject numbers for use in RSVP-TE as are assigned for the PKS in PCEP. The numbers above are the same as in [RFC5520].

6.2. Record Route Objects Subobjects

IANA maintains a registry called "Resource Reservation Protocol (RSVP) Parameters" with a subregistry called "Class Names, Class Numbers, and Class Types".

Within this subregistry, there is a definition of the ROUTE_RECORD object (also known as the RECORD_ROUTE object) with Class Number 21. The object definition lists a number of acceptable subobjects for the Class Type 1.

IANA has allocated two further subobjects as described in Section 3. The resulting entry in the registry is as follows.

```
 21  ROUTE_RECORD                                    [RFC3209]
     (also known as RECORD_ROUTE)
     Class Types or C-Types:
       1   Type 1 Route Record                       [RFC3209]
           Subobject type
            64   Path Key with 32-bit PCE-ID        [RFC5553]
            65   Path Key with 128-bit PCE-ID       [RFC5553]
```

Note well: IANA is requested to use the same subobject numbers as are defined for the EXPLICIT_ROUTE object in Section 6.1.

6.3. Error Codes and Error Values

IANA maintains a registry called "Resource Reservation Protocol (RSVP) Parameters" with a subregistry called "Error Codes and Globally-Defined Error Value Sub-Codes".

Within this subregistry, there is a definition of the "Routing Problem" error code with error code value 24. The definition lists a number of error values that may be used with this error code.

IANA has allocated further error values for use with this error code as described in Section 3.1. The resulting entry in the registry is as follows.

24 Routing Problem [RFC3209]

This Error Code has the following globally defined Error Value sub-codes:

31 = Unknown PCE-ID for PKS expansion [RFC5553]
32 = Unreachable PCE for PKS expansion [RFC5553]
33 = Unknown Path Key for PKS expansion [RFC5553]
34 = ERO too large for MTU [RFC5553]

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.

[RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

7.2. Informative References

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

[RFC4726] Farrel, A., Vasseur, J.-P., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, December 2008.

[RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 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.

[MPLS-SEC] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, March 2009.

Authors' Addresses

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

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

Jean-Philippe Vasseur
Cisco Systems, Inc
11, Rue Camille Desmoulins
L'Atlantis
92782 Issy Les Moulineaux
France
EMail: jpv@cisco.com