

CCAMP Working Group  
Internet Draft  
Category: Standards track

Xian Zhang  
Fatai Zhang  
Huawei  
O. Gonzalez de Dios  
Telefonica I+D  
Igor Bryskin  
ADVA Optical Networking  
Dhruv Dhody  
Huawei

Expires: August 13, 2014

February 14, 2014

Extensions to Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) to Support Route Exclusion Using Path Key Subobject

draft-zhang-ccamp-route-exclusion-pathkey-01.txt

#### Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on August 13, 2014.

Copyright Notice

Copyright (c) 2014 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 (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Requirements Language

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 [RFC2119].

Abstract

This document extends the Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) eXclude Route Object (XRO) and Explicit eXclusion Route Subobject (EXRS) within Explicit Route Object (ERO) to support specifying route exclusion requirement using Path Key Subobject (PKS).

Table of Contents

- 1. Introduction ..... 3
  - 1.1. Example Use ..... 3
- 2. RSVP-TE Extensions ..... 4
  - 2.1. Path Key Subobject (PKS)..... 4
  - 2.2. PKS Processing Rules..... 5
- 3. Other considerations..... 6
  - 3.1. Path-Key Retention and Reuse..... 6
  - 3.2. Path-Key Uniqueness..... 7
  - 3.3. PKS Update ..... 7
- 4. Manageability Considerations .....7
  - 4.1. Control of Function through Configuration and Policy.....7
- 5. Security Considerations..... 8
- 6. IANA Considerations ..... 8

6.1. New Subobject Type.....	8
6.2. New Error Code .....	8
7. Acknowledgments .....	9
8. References .....	9
8.1. Normative References.....	9
8.2. Informative References.....	9
9. Contributors .....	9
10. Authors' Addresses .....	9

## 1. Introduction

[RFC5520] defines the concept of a Path Key for confidentiality in a multi-domain environment. This can be used by a Path Computation Element (PCE) in place of a segment of a path that it wishes to keep confidential. The Path Key can be signaled in Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) protocol by placing it in an Explicit Route Object (ERO) as described in [RFC5553].

When establishing a set of LSPs to provide protection services [RFC4427], it is often desirable that the LSPs should take different paths through the network. This can be achieved by path computation entities that have full end-to-end visibility, but it is more complicated in multi-domain environments when segments of the path may be hidden so that they are not visible outside the domain they traverse.

This document describes how the Path Key object can be used in the RSVP-TE eXclude Route Object (XRO), and the Explicit eXclusion Route subobject (EXRS) of the ERO in order to facilitate path hiding, but allow diverse end-to-end paths to be established in multi-domain environments.

### 1.1. Example Use

Figure 1 shows a simple network with two domains. It is desired to set up a pair of path-disjoint LSPs from the source in Domain 1 to the destination in Domain 2, but the domains keep strict confidentiality about all path and topology information.

The first LSP will be signaled by the source with ERO {A, B, loose Dst} and will be set up with the path {Src, A, B, U, V, W, Dst}. But when sending the RRO out of Domain 2, node U would normally strip the path and replace it with a loose hop to the destination. With this limited information, the source is unable to include enough detail in the ERO of the second LSP to avoid it taking, for example, the path {Src, C, D, X, V, W, Dst} which is not path-disjoint.

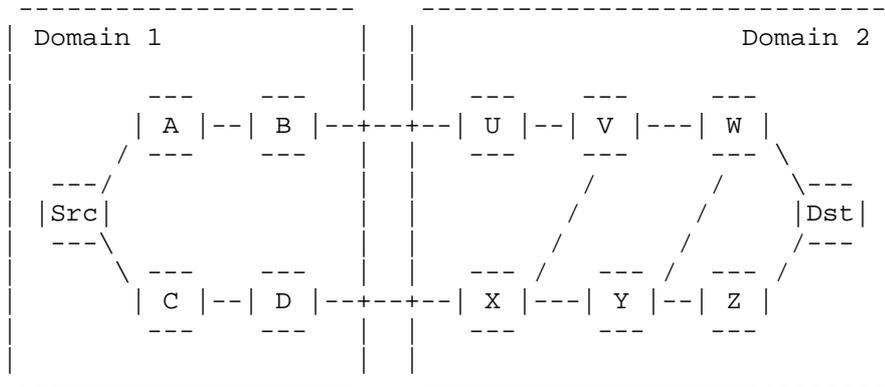


Figure 1: A Simple Multi-Domain Network

In order to improve the outcome, node U can replace the path segment {U, V, W} in the RRO with a Path Key Subobject. The Path Key Subobject assigns an identifier to the key and also indicates that it was node U that made the replacement.

With this additional information, the source is able to signal the second LSP with ERO set to {C, D, exclude Path Key(EXRS), loose Dst}. When the signaling message reaches node X, it can consult node U to expand the Path Key and so know to avoid the path of the first LSP. Alternatively, the source could use an ERO of {C, D, loose Dst} and include an XRO containing the Path Key.

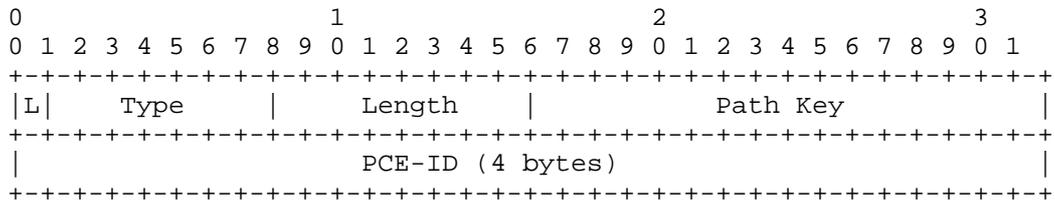
This example uses a PCE deployed in each border router, having at least the capability to expand PKS. Other deployment scenarios (Domain PCE, PCE being part of the Management system) may be used.

## 2. RSVP-TE Extensions

This section defines the Path Key Subobject that can be either in the XRO object or Explicit eXclusion Route subobject (EXRS) as defined in [RFC4874].

### 2.1. Path Key Subobject (PKS)

The IPv4 PKS has the same format as defined in [RFC5553] and is detailed as below:



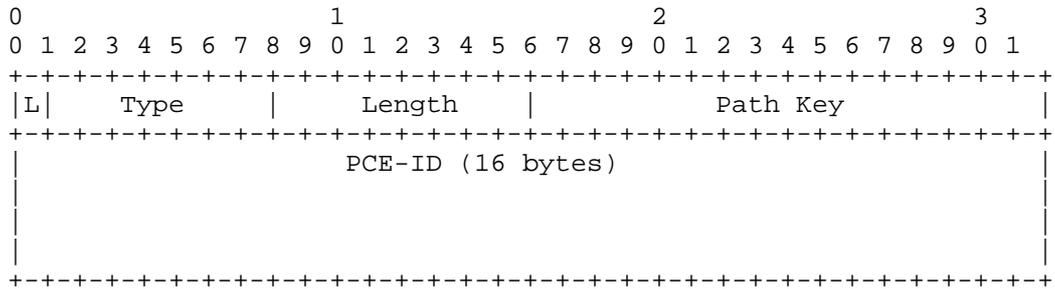
The meaning of the field Length and Path Key is defined in [RFC5553].

L: 0 indicates that the path or path segment hidden with the Path Key specified MUST be excluded. 1 indicates that the path or path segment hidden with the Path Key specified SHOULD be avoided.

Type: sub-object type for XRO Path Key; TBD.

PCE-ID: The IPv4 address of a node that assigned the Path Key identifier and that can return an expansion of the Path Key or use the Path Key as an exclusion in a path computation. Note this draft does not confine whether it is the network element or a dedicated server for path key generation and decoding.

Similarly, the format of IPv6 PKS is as follows:



### 2.2. PKS Processing Rules

The exclude route list is encoded as a series of subobjects contained in an EXCLUDE\_ROUTE object or an EXRS of the ERO. Multiple Path-Keys may be included in XRO or EXRO of ERO if more than segment of a path are kept hidden during diverse path establishment. The procedure defined in [RFC4874] for processing XRO and EXRS is not changed by this document.

Irrespective of the L flag, if the node, receiving the PKS, cannot recognize the subobject, it will react according to [RFC4874] and SHOULD ignore the constraint.

Otherwise, if it decodes the PKS but cannot find a route/route segment meeting the constraint:

- if L flag is set to 0, it will react according to [RFC4874] and SHOULD send a PathErr message with the error code/value combination "Routing Problem" / "Route Blocked by Exclude Route".

- if L flag is set to 1, which means the node SHOULD try to be as much diversified as possible with the specified resource. If it cannot fully support the constraint, it SHOULD send a PathErr message with the error code/value combination "Notify Error" / "Fail to find diversified path" (TBD).

If it cannot decode the PKS, the error handling procedure defined in Section 3.1 of [RFC5553] is not changed by this document.

This mechanism can work with all the PKS resolution mechanism, as detailed in [RFC5553] section 3.1. A PCE, co-located or not, may be used to resolve the PKS, but the node (i.e., a Label Switcher Router(LSR)) can also use the PKS information to index a Path Segment previously supplied to it by the entity that originated the PKS, for example the LSR that inserted the PKS in the RRO or a management system.

### 3. Other considerations

#### 3.1. Path-Key Retention and Reuse

The use of the path key relies on the availability of a PCE function supporting [RFC5520] functionality.

Following [RFC4655] a simple deployment option is when the PCE function is collocated with each border domain node generating the RRO. This collocated PCE functionality can be restricted to only serve the PKS resolution. This PCE function is only required to be accessible to the nodes excluding this PKS, so this can be restricted to one domain. This option can very easily tie the lifetime of the PKS to the lifetime of the LSP.

Alternatively, if a dedicated server, such as a PCE, is in charge of this, it may need to be explicitly informed of the LSP tear-down in order to recycle the path key allocated already. This can be easily supported by a stateful PCE [Stateful-PCE]. Note this draft does not confine the methods for path key generation and decoding.

Last, options including allowing a LSR can use the PKS information to index a Path Segment previously supplied to it by the entity that

originated the PKS, for example the LSR that inserted the PKS in the RRO or a management system, can also be used.

### 3.2. Path-Key Uniqueness

In the CCAMP mailing list, there is concern about whether 16-bit Path key is still enough and future proof. This can be easily solved by confining the scope of a path key. If an ingress node is responsible for managing the Path Key, it should not be an issue since the LSP across domains do not expected to be larger than 65535. On the other hand, if a dedicated entity, such as a PCE server, is used to allocate and recycle the Path Key, it is advised to allocate the Path Key per ingress node basis to avoid the limitation of Path Key numbers facing a domain-based allocation space. These are only illustrative examples and other methods that can guarantee the uniqueness of Path-Key are not precluded.

### 3.3. PKS Update

When the information of a path is changed, the LSPs using that path and corresponding PKS should be aware of the changes. The procedures defined in Section 4.4.3 of RFC 3209 [RFC3209] MUST be used to refresh the PKS information if the PKS change is to be communicated to other nodes according to the local node's policy. If local policy is that the PKS change should be suppressed or would result in no change to the PKS expansion, the node does not need to send an update. This procedure allows for ingress node to react on path change.

## 4. Manageability Considerations

### 4.1. Control of Function through Configuration and Policy

In addition to the set of policies described in [RFC5553] the following policies (are local and domain-wide) SHOULD be available for configuration in an implementation:

- Handling a XRO or EXRS containing a PKS. As described in Section 2.2, an LSR that receives a Path message containing a PKS exclusion can be configured to reject the Path message according to policy.
- Hiding of reason codes. The policy described in [RFC5553] section 5.1 is also applicable to policies for PKS in XRO or EXRS.

This document makes no other new management consideration to RSVP and PCE, the existing consideration applies.

## 5. Security Considerations

The use of path keys proposed in this draft allows nodes to hide parts of the path as it is signaled. This can be used to improve the confidentiality of the LSP setup. Moreover, it may serve to improve security of the control plane for the LSP as well as data plane traffic carried on this LSP. However, the benefits of using path key are lost unless there is an appropriate access control of any tool that allows expansion of the path key.

## 6. IANA Considerations

### 6.1. New Subobject Type

IANA registry: RSVP PARAMETERS  
 Subsection: Class Names, Class Numbers, and Class Types

This document introduces two new subobjects for the EXCLUDE\_ROUTE object [RFC4874], C-Type 1.

Subobject Type	Subobject Description
-----	-----
64(TBD by IANA)	IPv4 Path Key Subobject
65(TBD By IANA)	IPv6 Path Key Subobject

Note: [RFC5520] defines the PKS for use in PCEP. The above number suggestions for use in RSVP-TE follow that assigned for the PKS in PCEP [RFC5520].

### 6.2. New Error Code

IANA registry: RSVP PARAMETERS

Subsection: Error Codes and Globally-Defined Error Value Sub-Codes

New Error Values sub-codes have been registered for the Error Code 'Notify Error' (25).

TBD = "Fail to find diversified path"

## 7. Acknowledgments

The authors would like to thank John Drake, Daniele Ceccarelli and Zafar Ali for their comments and discussions.

## 8. References

### 8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.
- [RFC3209] D. Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, December 2001.
- [RFC4874] CY. Lee, A. Farrel, S. De Cnodder, "Exclude Routes - Extension to Resource Reservation Protocol-Traffic Engineering (RSVP-TE)", RFC4874, April 2007.
- [RFC5553] A. Farrel, Ed., "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC5553, May 2009.

### 8.2. Informative References

- [RFC5520] R. Bradford, Ed., "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC5520, April 2009.
- [RFC4427] E. Mannie, Ed., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC4427, March 2006.
- [Stateful-PCE] Crabbe, E., Medved, J., Minei, I., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-06 (work in progress), August 2013.

## 9. Contributors

Cyril  
cyril.margaria@gmail.com

## 10. Authors' Addresses

Xian Zhang  
Huawei Technologies  
  
Email: zhang.xian@huawei.com

Fatai Zhang  
Huawei Technologies  
  
Email: zhangfatai@huawei.com

Oscar Gonzalez de Dios  
Telefonica I+D  
Don Ramon de la Cruz  
Madrid, 28006  
Spain

Phone: +34 913328832  
Email: ogondio@tid.es

Igor Bryskin  
ADVA Optical Networking  
  
Email: ibryskin@advaoptical.com

Dhruv Dhody  
Huawei Technologies  
Leela Palace  
Bangalore, Karnataka 560008  
INDIA  
  
EMail: dhruv.ietf@gmail.com

