

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: April 26, 2015

Zafar Ali, Ed.
George Swallow, Ed.
Cisco Systems
F. Zhang, Ed.
Huawei
D. Beller, Ed.
Alcatel-Lucent
October 27, 2014

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Path
Diversity using Exclude Route

draft-ietf-ccamp-lsp-diversity-05.txt

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 26, 2015.

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.

Abstract

RFC 4874 specifies methods by which path exclusions can be communicated during RSVP-TE signaling in networks where precise explicit paths are not computed by the LSP source node. This document specifies procedures for additional route exclusion subobject based on Paths currently existing or expected to exist within the network.

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].

Table of Contents

1. Introduction	2
1.1. Client-Initiated Identifier	5
1.2. PCE-allocated Identifier	6
1.3. Network-Assigned Identifier	7
2. RSVP-TE signaling extensions	9
2.1. Diversity XRO Subobject	9
2.1.1. IPv4 Diversity XRO Subobject	9
2.1.2. IPv6 Diversity XRO Subobject	14
2.2. Processing rules for the Diversity XRO subobject	17
2.3. Diversity EXRS Subobject	20
3. Security Considerations	22
4. IANA Considerations	22
4.1. New XRO subobject types	22
4.2. New EXRS subobject types	23
4.3. New RSVP error sub-codes	23
5. Acknowledgements	23
6. References	24
6.1. Normative References	24
6.2. Informative References	24

1. Introduction

Path diversity for multiple connections is a well-known Service Provider requirement. Diversity constraints ensure that Label-Switched Paths (LSPs) can be established without sharing resources, thus greatly reducing the probability of simultaneous connection failures.

When a source node has full topological knowledge and is permitted to signal an Explicit Route Object, diverse paths for LSPs can be computed by this source node. However, there are scenarios when

path computations are performed by different nodes, and there is therefore a need for relevant diversity constraints to be communicated to those nodes. These include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), e.g. inter-domain LSPs;
- . Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI), where path computation may be performed by the core node [RFC4208].

[RFC4874] introduced a means of specifying nodes and resources to be excluded from a route, using the eXclude Route Object (XRO) and Explicit Exclusion Route Subobject (EXRS). It facilitates the calculation of diverse paths for LSPs based on known properties of those paths including addresses of links and nodes traversed, and Shared Risk Link Groups (SRLGs) of traversed links. Employing these mechanisms requires that the source node that initiates signaling knows the relevant properties of the path(s) from which diversity is desired. However, there are circumstances under which this may not be possible or desirable, including (but not limited to):

- . Exclusion of a path which does not originate, terminate or traverse the source node of the diverse LSP, in which case the addresses of links and SRLGs of the path from which diversity is required are unknown to the source node.
- . Exclusion of a path which is known to the source node of the diverse LSP for which the node has incomplete or no path information, e.g. due to operator policy. In this case, the existence of the reference path is known to the source node but the information required to construct an XRO object to guarantee diversity from the reference path is not fully known. Inter-domain and GMPLS overlay networks can present such restrictions.

This is exemplified in the Figure 1, where overlay reference model from [RFC4208] is shown.

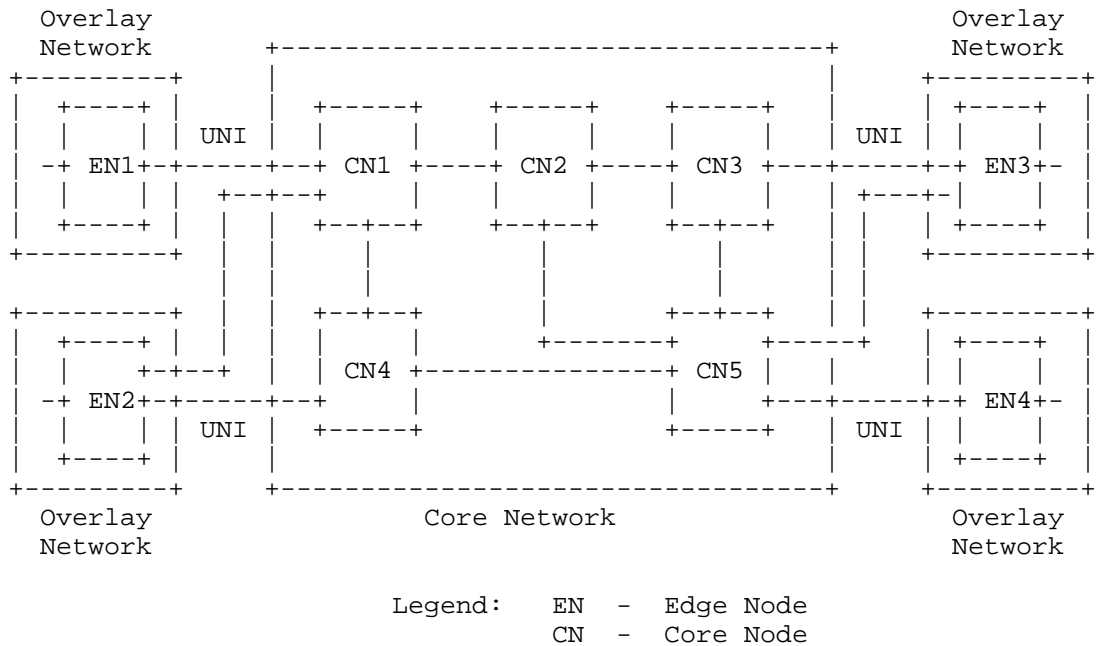


Figure 1: Overlay Reference Model [RFC4208]

Figure 1 depicts two types of UNI connectivity: single-homed and dual-homed ENs (which also applies to higher order multi-homed connectivity.). Single-homed EN devices are connected to a single CN device via a single UNI link. This single UNI link may constitute a single point of failure. UNI connection between EN1 and CN1 is an example of single-homed UNI connectivity.

A single point of failure caused by a single-homed UNI can be avoided when the EN device is connected to two different CN devices, as depicted for EN2 in Figure 1. For the dual-homing case, it is possible to establish two different UNI connections from the same source EN device to the same destination EN device. For example, two connections from EN2 to EN3 may use the two UNI links EN2-CN1 and EN2-CN4. To avoid single points of failure within the provider network, it is necessary to also ensure path (LSP) diversity within the core network.

In a UNI network such as that shown in Figure 1, the CNs typically perform path computation. Information sharing across

the UNI boundary is restricted based on the policy rules imposed by the core network. Typically, the core network topology information is not exposed to the ENs. In the network shown in Figure 1, consider a use case where an LSP from EN2 to EN4 needs to be SRLG diverse from an LSP from EN1 to EN3. In this case, EN2 may not know SRLG attributes of the EN1- EN3 LSP and hence cannot construct an XRO to exclude these SRLGs. In this example EN2 cannot use the procedures described in [RFC4874]. Similarly, an LSP from EN2 to EN3 traversing CN1 needs to be diverse from an LSP from EN2 to EN3 going via CN4. Again in this case, exclusions based on [RFC4874] cannot be used.

This document addresses these diversity requirements by introducing the notion of excluding the path taken by particular LSP(s). The reference LSP(s) or route(s) from which diversity is required is/are identified by an "identifier". The type of identifier to use is highly dependent on the networking deployment scenario; it could be client-initiated, allocated by the (core) network or managed by a PCE. This document defines three different types of identifiers corresponding to these three cases: a client initiated identifier, a PCE allocated Identifier and CN ingress node (UNI-N) allocated Identifier.

1.1. Client-Initiated Identifier

There are scenarios in which the ENs have the following requirements for the diversity identifier:

- The identifier is controlled by the client side and is specified as part of the service request.
- Both client and server understand the identifier.
- It is necessary to be able to reference the identifier even if the LSP referenced by it is not yet signaled.
- The identifier is to be stable for a long period of time.
- The identifier is to be stable even when the referenced tunnel is rerouted.
- The identifier is to be human-readable.

These requirements are met by using the Resource ReserVation Protocol (RSVP) tunnel/ LSP Forwarding Equivalence Class (FEC) as the identifier.

The usage of the client-initiated identifier is illustrated by using Figure 1. Suppose a tunnel from EN2 to EN4 needs to be diverse with respect to a tunnel from EN1 to EN3. The tunnel FEC of the EN1-EN3 tunnel is FEC1, where FEC1 is defined by the tuple (tunnel-id = T1, source address = EN1.ROUTE Identifier (RID), destination address = EN3.RID, extended tunnel-id = EN1.RID). Similarly, tunnel FEC of the EN2-EN3 tunnel is FEC2, where FEC2 is defined by the tuple (tunnel-id = T2, source address = EN2.RID, destination address = EN4.RID, extended tunnel-id = EN2.RID). The EN1-EN3 tunnel is signaled with an exclusion requirement from FEC2, and the EN2-EN3 tunnel is signaled with an exclusion requirement from FEC1. In order to maintain diversity between these two connections within the core network, it is assumed that the core network implements Crankback Signaling [RFC4920]. Note that crankback signaling is known to lead to slower setup times and sub-optimal paths under some circumstances as described by [RFC4920].

1.2. PCE-allocated Identifier

In scenarios where a PCE is deployed and used to perform path computation, the core edge node (e.g., node CN1 in Figure 1) could consult a PCE to allocate identifiers, which are used to signal path diversity constraints. In other scenarios a PCE is deployed in each border node or a PCE is part of a Network Management System (NMS). In all these cases, the Path Key as defined in [RFC5520] can be used in RSVP signaling as the identifier to ensure diversity.

An example of specifying LSP diversity using a Path Key is shown in Figure 2, where a simple network with two domains is shown. 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 is signaled by the source with ERO {A, B, loose Dst} and is set up with the path {Src, A, B, U, V, W, Dst}. However, 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} for path-disjointness.




Figure 1: A Simple Multi-Domain Network

In order to improve the situation, node U performs the PCE function and replaces the path segment {U, V, W} in the RRO with a Path Key Subobject. The Path Key Subobject assigns an "identifier" to the key. The PCE ID in the message indicates that it was node U that made the replacement.

With this additional information, the source is able to signal the subsequent LSPs with the 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 know how 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 mechanism can work with all the Path-Key resolution mechanisms, as detailed in [RFC5553] section 3.1. A PCE, co-located or not, may be used to resolve the Path-Key, but the node (i.e., a Label Switching Router (LSR)) can also use the Path Key information to index a Path Segment previously supplied to it by the entity that originated the Path-Key, for example the LSR that inserted the Path-Key in the RRO or a management system.

1.3. Network-Assigned Identifier

There are scenarios in which the network provides diversity-related information for a service that allows the client device to include this information in the signaling message. If the Shared Resource Link Group (SRLG) identifier information is both available and shareable (by policy) with the ENs, the procedure

defined in [DRAFT-SRLG-RECORDING] can be used to collect SRLG identifiers associated with an LSP (LSP1). When a second LSP (LSP2) needs to be diverse with respect to LSP1, the EN constructing the RSVP signaling message for setting up LSP2 can insert the SRLG identifiers associated with LSP1 as diversity constraints into the XRO using the procedure described in [RFC4874]. However, if the core network SRLG identifiers are either not available or not shareable with the ENs based on policies enforced by core network, existing mechanisms cannot be used.

In this draft, a signaling mechanism is defined where information signaled to the CN via the UNI does not require shared knowledge of core network SRLG information. For this purpose, the concept of a Path Affinity Set (PAS) is used for abstracting SRLG information. The motive behind the introduction of the PAS is to minimize the exchange of diversity information between the core network (CNs) and the client devices (ENs). The PAS contains an abstract SRLG identifier associated with a given path rather than a detailed SRLG list. The PAS is a single identifier that can be used to request diversity and associate diversity. The means by which the processing node determines the path corresponding to the PAS is beyond the scope of this document.

A CN on the core network boundary interprets the specific PAS identifier (e.g. "123") as meaning to exclude the core network SRLG information (or equivalent) that has been allocated by LSPs associated with this PAS identifier value. For example, if a Path exists for the LSP with the identifier "123", the CN would use local knowledge of the core network SRLGs associated with the "123" LSPs and use those SRLGs as constraints for path computation. If a PAS identifier is included for exclusion in the connection request, the CN (UNI-N) in the core network is assumed to be able to determine the existing core network SRLG information and calculate a path that meets the determined diversity constraints.

When a CN satisfies a connection setup for a (SRLG) diverse signaled path, the CN may optionally record the core network SRLG information for that connection in terms of CN based parameters and associates that with the EN addresses in the Path message. Specifically for Layer-1 Virtual Private Networks (L1VPNs), Port Information Tables (PIT) [RFC5251] can be leveraged to translate between client (EN) addresses and core network addresses.

The PAS and the associated SRLG information can be distributed within the core network by an Interior Gateway Protocol (IGP) or

by other means such as configuration. They can then be utilized by other CNs when other ENs are requesting paths to be setup that would require path/connection diversity. In the VPN case, this information is distributed on a VPN basis and contains a PAS identifier, CN addresses and SRLG information. In this way, on a VPN basis, the core network can have additional opaque records for the PAS values for various Paths along with the SRLG list associated with the Path. This information is internal to the core network and is known only to the core network.

2. RSVP-TE signaling extensions

This section describes the signaling extensions required to address the aforementioned requirements and use cases.

2.1. Diversity XRO Subobject

New Diversity XRO subobjects are defined by this document as follows.

2.1.1. IPv4 Diversity XRO Subobject

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	2	3	4	5	6	7	8	9
L XRO Type										Length										DI Type A-Flags E-Flags										Resvd									
										IPv4 Diversity Identifier source address																													
										Diversity Identifier Value																													
//										...																				//									

L:

The L-flag is used as for the XRO subobjects defined in [RFC4874], i.e.,

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

XRO Type

Type for IPv4 diversity XRO subobject (to be assigned by IANA; suggested value: 37).

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is variable, depending on the diversity identifier value.

Diversity Identifier Type (DI Type)

Diversity Identifier Type (DI Type) indicates the way the reference LSP(s) or route(s) with which diversity is required is identified. Three values are defined in this document:

IPv4 Client Initiated Identifier	1 (to be assigned by IANA)
IPv4 PCE Allocated Identifier	2 (to be assigned by IANA)
IPv4 Network Assigned Identifier	3 (to be assigned by IANA)

Attribute Flags (A-Flags):

The Attribute Flags (A-Flags) are used to communicate desirable attributes of the LSP being signaled. The following flags are defined. Each flag acts independently. Any combination of flags is permitted.

0x01 = Destination node exception

Indicates that the exclusion does not apply to the destination node of the LSP being signaled.

0x02 = Processing node exception

Indicates that the exclusion does not apply to the border node(s) performing ERO expansion for the LSP being signaled. An ingress UNI-N node is an example of such a node.

0x04 = Penultimate node exception

Indicates that the penultimate node of the LSP being signaled MAY be shared with the excluded path even when this violates the exclusion flags.

0x08 = LSP ID to be ignored

This flag is only applicable when the diversity is specified using the client-initiated identifier, the flag indicates tunnel level exclusion, as detailed in section 2.2.

Exclusion Flags (E-Flags):

The Exclusion-Flags are used to communicate the desired type(s) of exclusion. The following flags are defined. Any combination of these flags is permitted.

0x01 = SRLG exclusion

Indicates that the path of the LSP being signaled is requested to be SRLG-diverse from the excluded path specified by the Diversity XRO subobject.

0x02 = Node exclusion

Indicates that the path of the LSP being signaled is requested to be node-diverse from the excluded path specified by the Diversity XRO subobject.

(Note: the meaning of this flag may be modified by the value of the Attribute-flags.)

0x04 = Link exclusion

Indicates that the path of the LSP being signaled is requested to be link-diverse from the path specified by the Diversity XRO subobject.

Resvd

This field is reserved. It SHOULD be set to zero on transmission, and MUST be ignored on receipt.

IPv4 Diversity Identifier source address:

This field is set to the IPv4 address of the node that assigns the diversity identifier. Depending on the diversity identifier type, the diversity identifier source may be a client node, PCE entity or network node. Specifically:

- o When the diversity identifier type is set to "IPv4 Client Initiated Identifier", the value is set to IPv4 tunnel sender address of the reference LSP against which diversity is desired. IPv4 tunnel sender address is as defined in [RFC3209].
- o When the diversity identifier type is set to "IPv4 PCE Allocated Identifier", the value indicates the IPv4 address of the node that assigned the Path Key identifier and that can return an expansion of the Path Key or use the Path Key as exclusion in a path computation. The Path Key is defined in [RFC5553].
- o When the diversity identifier type is set to "IPv4 Network Assigned Identifier", the value indicates the IPv4 address of the node publishing the Path Affinity Set (PAS).

Diversity Identifier Value:

Encoding for this field depends on the diversity identifier type, as defined in the following.

When the diversity identifier type is set to "IPv4 Client Initiated Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv4 tunnel end point address                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Must Be Zero           |           Tunnel ID           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Extended Tunnel ID                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Must Be Zero           |           LSP ID           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The IPv4 tunnel end point address, Tunnel ID, Extended Tunnel ID and LSP ID are as defined in [RFC3209].

When the diversity identifier type is set to "IPv4 PCE Allocated Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Must Be Zero           |           Path Key           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The Path Key is defined in [RFC5553].

When the diversity identifier type is set to "IPv4 Network Assigned Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Path Affinity Set (PAS) identifier                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The Path affinity Set (PAS) identifier is a single number that represents a summarized SRLG for the reference path against which diversity is desired. The node identified by the "IPv4 Diversity Identifier source address" field of the diversity XRO subobject assigns the PAS value.

2.1.2. IPv6 Diversity XRO Subobject

```

      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|  XRO Type  |      Length      |DI Type|A-Flags|E-Flags| Resvd |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 Diversity Identifier source address
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 Diversity Identifier source address (cont.)
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 Diversity Identifier source address (cont.)
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 Diversity Identifier source address (cont.)
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Diversity Identifier Value
//                                     ...                                     //
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

L:

The L-flag is used as for the XRO subobjects defined in [RFC4874], i.e.,

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

XRO Type

Type for IPv6 diversity XRO subobject (to be assigned by IANA; suggested value: 38).

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is variable, depending on the diversity identifier value.

Attribute Flags (A-Flags):

As defined in Section 2.1.1 for the IPv4 counterpart.

Exclusion Flags (E-Flags):

As defined in Section 2.1.1 for the IPv4 counterpart.

Resvd

This field is reserved. It SHOULD be set to zero on transmission, and MUST be ignored on receipt.

Diversity Identifier Type (DI Type)

This field is defined in the same fashion as its IPv4 counterpart described in Section 2.1.1.
The DI Types associated with IPv6 addresses are defined, as follows:

IPv6 Client Initiated Identifier	4 (to be assigned by IANA)
IPv6 PCE Allocated Identifier	5 (to be assigned by IANA)
IPv6 Network Assigned Identifier	6 (to be assigned by IANA)

These identifier are assigned and used as defined in Section 2.1.1.

IPv4 Diversity Identifier source address:

This field is set to IPv6 address of the node that assigns the diversity identifier. How identity of node for various diversity types is determined is as described in Section 2.1.1 for the IPv4 counterpart.

Diversity Identifier Value:

Encoding for this field depends on the diversity identifier type, as defined in the following.

When the diversity identifier type is set to "IPv6 Client Initiated Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 tunnel end point address                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 tunnel end point address (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 tunnel end point address (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     IPv6 tunnel end point address (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Must Be Zero          |          Tunnel ID          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Extended Tunnel ID                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Extended Tunnel ID (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Extended Tunnel ID (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Extended Tunnel ID (cont.)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Must Be Zero          |          LSP ID          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The IPv6 tunnel end point address, Tunnel ID, IPv6 Extended Tunnel ID and LSP ID are as defined in [RFC3209].

When the diversity identifier type is set to "IPv6 PCE Allocated Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Must Be Zero          |          Path Key          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The Path Key is defined in [RFC5553].

When the diversity identifier type is set to "IPv6 Network Assigned Identifier", the diversity identifier value is encoded as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Path Affinity Set (PAS) identifier                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The Path affinity Set (PAS) identifier is as defined in Section 2.1.1.

2.2. Processing rules for the Diversity XRO subobject

The procedure defined in [RFC4874] for processing XRO and EXRS is not changed by this document. If the processing node cannot recognize the IPv4/ IPv6 Diversity XRO subobject, the node is expected to follow the procedure defined in [RFC4874].

An XRO object MAY contain multiple Diversity subobjects. E.g., In order to exclude multiple Path Keys, an EN may include multiple Diversity XRO subobjects each with a different Path Key. Similarly, in order to exclude multiple PAS identifiers, an EN may include multiple Diversity XRO subobjects each with a different PAS identifier. However, all Diversity subobjects in an XRO SHOULD contain the same Diversity Identifier Type. If a Path message contains an XRO with Diversity subobjects with multiple Diversity Identifier Types, the processing node SHOULD return a PathErr with the error code "Routing Problem" (24) and error sub-code "XRO Too Complex" (68).

The attribute-flags affect the processing of the Diversity XRO subobject as follows:

- o When the "destination node exception" flag is set, the exclusion SHOULD be ignored for the destination node.
- o When the "processing node exception" flag is set, the exclusion SHOULD be ignored for the processing node. The processing node is the node performing path calculation.

- o When the "penultimate node exception" flag is set, the exclusion SHOULD be ignored for the penultimate node on the path of the LSP being established.
- o The "LSP ID to be ignored" flag is only defined for the "IPv4/ IPv6 Client Initiated Identifier" diversity types. When the Diversity Identifier Type is set to any other value, this flag SHOULD NOT be set on transmission and MUST be ignored in processing. When this flag is not set, the lsp-id is not ignored and the exclusion applies only to the specified LSP (i.e., LSP level exclusion).

If the L-flag of the diversity XRO subobject is not set, the processing node proceeds as follows.

- "IPv4/ IPv6 Client Initiated Identifiers" Diversity Type: the processing node MUST ensure that any path calculated for the signaled LSP is diverse from the RSVP TE FEC identified by the client in the XRO subobject.
- "IPv4/ IPv6 PCE Allocated Identifiers" Diversity Type: the processing node MUST ensure that any path calculated for the signaled LSP is diverse from the route identified by the Path-Key. The processing node MAY use the PCE identified by the IPv4 Diversity Identifier source address in the subobject for route computation. The processing node MAY use the Path-Key resolution mechanisms described in [RFC5553].
- "IPv4/ IPv6 Network Assigned Identifiers" Diversity Type: the processing node MUST ensure that the path calculated for the signaled LSP respects the requested PAS exclusion. .
- Regardless of whether the path computation is performed locally or at a remote node (e.g., PCE), the processing node MUST ensure that any path calculated for the signaled LSP respects the requested exclusion flags with respect to the excluded path referenced by the subobject, including local resources.
- If the excluded path referenced in the XRO subobject is unknown to the processing node, the processing node SHOULD ignore the diversity XRO subobject and SHOULD proceed with the signaling request. After sending the ResvErr for the signaled LSP, the processing node SHOULD return a PathErr with the error code "Notify Error" (25) and error sub-code "Route reference in diversity XRO identifier unknown" (value to be assigned by IANA, suggested value: 13) for the signaled LSP.

- If the processing node fails to find a path that meets the requested constraint, the processing node MUST return a PathErr with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67).

If the L-flag of the diversity XRO subobject is set, the processing node proceeds as follows:

- "IPv4/ IPv6 Client Initiated Identifiers" Diversity Type: the processing node SHOULD ensure that the path calculated for the signaled LSP is diverse from the RSVP TE FEC identified by the client in the XRO subobject.
- "IPv4/ IPv6 PCE Allocated Identifiers" Diversity Type: the processing node SHOULD ensure that the path calculated for the signaled LSP is diverse from the route identified by the Path-Key.
- "IPv4/ IPv6 Network Assigned Identifiers" Diversity Type: the processing node SHOULD ensure that the path calculated for the signaled LSP respects the requested PAS exclusion. The means by which the processing node determines the path corresponding to the PAS is beyond the scope of this document.
- The processing node SHOULD respect the requested exclusion flags with respect to the excluded path to the extent possible.
- If the processing node fails to find a path that meets the requested constraint, it SHOULD proceed with signaling using a suitable path that meets the constraint as far as possible. After sending the Resv for the signaled LSP, it SHOULD return a PathErr message with error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14) to the source node.

If, subsequent to the initial signaling of a diverse LSP:

- An excluded path referenced in the XRO subobject becomes known to the processing node, or a change in the excluded path becomes known to the processing node, the processing node SHOULD re-evaluate the exclusion and diversity constraints requested by the diverse LSP to determine whether they are still satisfied.
- If the requested exclusion constraints for the diverse LSP are no longer satisfied and an alternative path for the diverse LSP that can satisfy those constraints exists, then:

- o If the L-flag was not set in the original exclusion, the processing node MUST send a PathErr message for the diverse LSP with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67). The PSR flag SHOULD NOT be set. A source node receiving a PathErr message with this error code and sub-code combination SHOULD take appropriate actions to migrate the compliant path.
- o If the L-flag was set in the original exclusion, the processing node SHOULD send a PathErr message for the diverse LSP with the error code "Notify Error" (25) and a new error sub-code "compliant path exists" (value: to be assigned by IANA, suggest value: 15). The PSR flag SHOULD NOT be set. A source node receiving a PathErr message with this error code and sub-code combination MAY signal a new LSP to migrate the compliant path.
- If the requested exclusion constraints for the diverse LSP are no longer satisfied and no alternative path for the diverse LSP that can satisfy those constraints exists, then:
 - o If the L-flag was not set in the original exclusion, the processing node MUST send a PathErr message for the diverse LSP with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67). The PSR flag SHOULD be set.
 - o If the L-flag was set in the original exclusion, the processing node SHOULD send a PathErr message for the diverse LSP with the error code error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14). The PSR flag SHOULD NOT be set.

The following rules apply whether or not the L-flag is set:

- A source node receiving a PathErr message with the error code "Notify Error" (25) and error sub-codes "Route of XRO tunnel identifier unknown" or "Failed to respect Exclude Route" MAY take no action.

2.3. Diversity EXRS Subobject

[RFC4874] defines the EXRS ERO subobject. An EXRS is used to identify abstract nodes or resources that must not or should not be used on the path between two inclusive abstract nodes or

resources in the explicit route. An EXRS contains one or more subobjects of its own, called EXRS subobjects [RFC4874].

An EXRS MAY include Diversity subobject as specified in this document. In this case, the IPv4 EXRS format is as follows:

```

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      |      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|L|  XRO Type  |      Length  |DI Type|A-Flags|E-Flags| Resvd |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IPv4 Diversity Identifier source address      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Diversity Identifier Value      |
//                                     ...                                     //
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Similarly, the IPv6 EXRS format is as follows:

```

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      |      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|L|  XRO Type  |      Length  |DI Type|A-Flags|E-Flags| Resvd |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IPv6 Diversity Identifier source address      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IPv6 Diversity Identifier source address (cont.)      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IPv6 Diversity Identifier source address (cont.)      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IPv6 Diversity Identifier source address (cont.)      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Diversity Identifier Value      |
//                                     ...                                     //
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The meanings of respective fields in EXRS header are as defined in [RFC4874]. The meanings of respective fields in the Diversity subobject are as defined earlier in this document for the XRO subobject.

The processing rules for the EXRS object are unchanged from [RFC4874]. When the EXRS contains one or more Diversity subobject(s), the processing rules specified in Section 2.2 apply to the node processing the ERO with the EXRS subobject.

If a loose-hop expansion results in the creation of another loose-hop in the outgoing ERO, the processing node MAY include the EXRS in the newly created loose hop for further processing by downstream nodes.

The processing node exception for the EXRS subobject applies to the node processing the ERO.

The destination node exception for the EXRS subobject applies to the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

The penultimate node exception for the EXRS subobject applies to the node before the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

3. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473] and [RFC4874].

4. IANA Considerations

4.1. New XRO subobject types

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 Description -----	Subobject Type -----
IPv4 Diversity subobject	To be assigned by IANA (suggested value: 37)
IPv6 Diversity subobject	To be assigned by IANA (suggested value: 38)

4.2. New EXRS subobject types

The diversity XRO subobjects are also defined as new EXRS subobjects.

4.3. New RSVP error sub-codes

IANA registry: RSVP PARAMETERS
Subsection: Error Codes and Globally Defined Error Value Sub-Codes

For Error Code "Notify Error" (25) (see [RFC3209]) the following sub-codes are defined.

Sub-code -----	Value -----
Route of XRO tunnel identifier unknown	To be assigned by IANA. Suggested Value: 13.
Failed to respect Exclude Route	To be assigned by IANA. Suggested Value: 14.
Compliant path exists	To be assigned by IANA. Suggested Value: 15.

5. Acknowledgements

The authors would like to thank Luyuan Fang and Walid Wakim for their review comments.

6. References

6.1. 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., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.
- [RFC5553] Farrel, A., Ed., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, May 2009.

6.2. Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC4920] Farrel, A., Ed., Satyanarayana, A., Iwata, A., Fujita, N., and G. Ash, "Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE", RFC 4920, July 2007.
- [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.
- [DRAFT-SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria, "RSVP-TE Extensions for Collecting SRLG Information", draft-ietf-ccamp-rsvp-te-srlg-collect.txt, work in progress.

- [RFC2205] Braden, R. (Ed.), Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReserVation Protocol -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
- [RFC5253] Takeda, T., Ed., "Applicability Statement for Layer 1 Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July 2008.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Contributors' Addresses

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

Daniele Ceccarelli
Ericsson
Email: Daniele.Ceccarelli@ericsson.com

Dhruv Dhody
Huawei Technologies
EMail: dhruv.ietf@gmail.com

Oscar Gonzalez de Dios
Telefonica I+D
Email: ogondio@tid.es

Don Fedyk
Hewlett-Packard
Email: don.fedyk@hp.com

Clarence Filsfils
Cisco Systems, Inc.
Email: cfilsfil@cisco.com

Xihua Fu
ZTE

Email: fu.xihua@zte.com.cn

Gabriele Maria Galimberti
Cisco Systems
Email: ggalimbe@cisco.com

Ori Gerstel
SDN Solutions Ltd.
Email: origerstel@gmail.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Email: Ruediger.Kunze@telekom.de

Lieven Levrau
Alcatel-Lucent
Email: Lieven.Levrau@alcatel-lucent.com

Cyril Margaria
cyril.margaria@gmail.com

Julien Meuric
France Telecom Orange
Email: julien.meuric@orange.com

Yuji Tochio
Fujitsu
Email: tochio@jp.fujitsu.com

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

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

Dieter Beller
Alcatel-Lucent
Email: Dieter.Beller@alcatel-lucent.com

George Swallow
Cisco Systems
Email: swallow@cisco.com

Fatai Zhang
Huawei Technologies
Email: zhangfatai@huawei.com