

Internet Engineering Task Force (IETF)
Request for Comments: 8390
Updates: 4874
Category: Standards Track
ISSN: 2070-1721

Z. Ali, Ed.
Cisco Systems
G. Swallow, Ed.
SETC
F. Zhang, Ed.
Huawei
D. Beller, Ed.
Nokia
July 2018

RSVP-TE Path Diversity Using Exclude Route

Abstract

RSVP-TE provides support for the communication of exclusion information during Label Switched Path (LSP) setup. A typical LSP diversity use case is for protection, where two LSPs should follow different paths through the network in order to avoid single points of failure, thus greatly improving service availability. This document specifies an approach that can be used for network scenarios where the full path(s) is not necessarily known by use of an abstract identifier for the path. Three types of abstract identifiers are specified: client based, Path Computation Element (PCE) based, and network based. This document specifies two new diversity subobjects for the RSVP eXclude Route Object (XRO) and the Explicit Exclusion Route Subobject (EXRS).

For the protection use case, LSPs are typically created at a slow rate and exist for a long time so that it is reasonable to assume that a given (reference) path currently existing (with a well-known identifier) will continue to exist and can be used as a reference when creating the new diverse path. Re-routing of the existing (reference) LSP, before the new path is established, is not considered.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8390>.

Copyright Notice

Copyright (c) 2018 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 (<https://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.

Table of Contents

1.	Introduction	3
1.1.	Conventions Used in This Document	6
1.2.	Terms and Abbreviations	6
1.3.	Client-Initiated Identifier	7
1.4.	PCE-Allocated Identifier	7
1.5.	Network-Assigned Identifier	9
2.	RSVP-TE Signaling Extensions	10
2.1.	Diversity XRO Subobject	10
2.2.	Diversity EXRS Subobject	16
2.3.	Processing Rules for the Diversity XRO and EXRS Subobjects	16
3.	Security Considerations	20
4.	IANA Considerations	21
4.1.	New XRO Subobject Types	21
4.2.	New EXRS Subobject Types	21
4.3.	New RSVP Error Sub-codes	22
5.	References	22
5.1.	Normative References	22
5.2.	Informative References	23
	Acknowledgements	24
	Contributors	24
	Authors' Addresses	26

1. Introduction

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

The source node can compute diverse paths for LSPs when it has full knowledge of the network topology and is permitted to signal an Explicit Route Object (ERO). However, there are scenarios where different nodes perform path computations, and therefore there is a need for relevant diversity constraints to be signaled to those nodes. These include (but are not limited to):

- o LSPs with loose hops in the Explicit Route Object, e.g., inter-domain LSPs; and
- o Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI), where the core node may perform path computation [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):

- o Exclusion of a path that 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.
- o Exclusion of a path that 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 source node is aware of the existence of the reference path, 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 impose such restrictions.

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

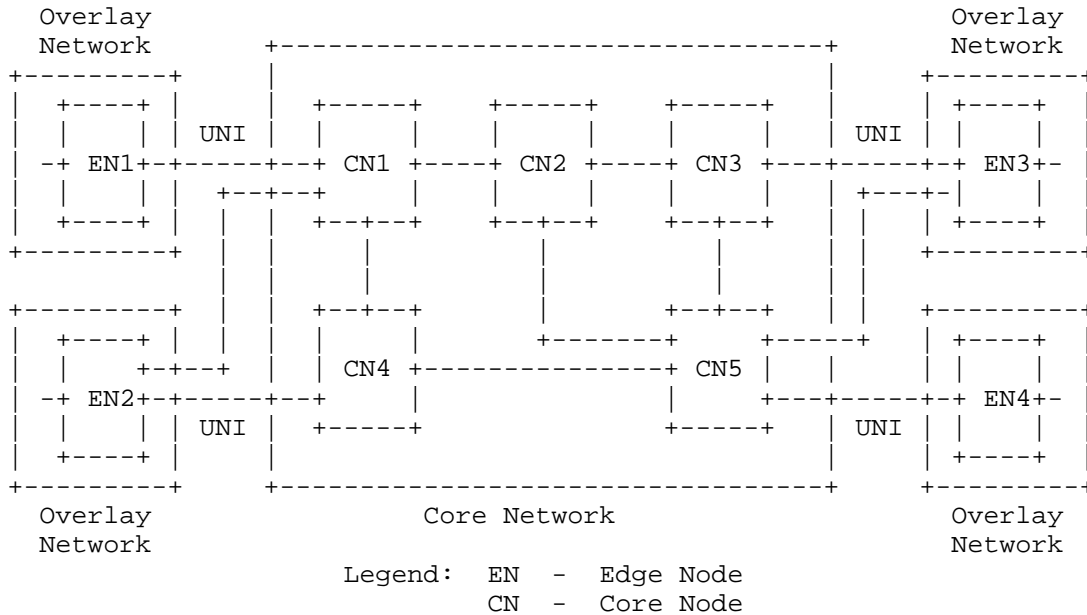


Figure 1: Overlay Reference Model [RFC4208]

Figure 1 depicts two types of UNI connectivity: single-homed and dual-homed ENs (which also applies to higher-order multihomed 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.

Such a single point of failure 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 network providing a set of UNI interfaces between ENs and CNs 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 as well as LSP path 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 an approach of excluding the path taken by these particular LSP(s). Each reference LSP or route from which diversity is required is identified by an abstract "identifier". The type of identifier to use is highly dependent on the core network operator's networking deployment scenario; it could be client initiated (provided by the EN), provided by a PCE, or allocated by the (core) network. This document defines three different types of identifiers corresponding to these three cases: a client-initiated identifier, a PCE-allocated identifier, and an identifier allocated by the CN ingress node (UNI-N), i.e., a network-assigned identifier.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terms and Abbreviations

Diverse LSP: A diverse Label Switched Path (LSP) is an LSP that has a path that does not have any link or SRLG in common with the path of a given LSP. Diverse LSPs are meaningful in the context of protection or restoration.

ERO: Explicit Route Object as defined in [RFC3209].

EXRS: Explicit Exclusion Route Subobject as defined in [RFC4874].

SRLG: Shared Risk Link Group as defined in [RFC4202].

Reference Path: The reference path is the path of an existing LSP to which the path of a diverse LSP shall be diverse.

XRO: eXclude Route Object as defined in [RFC4874].

1.3. Client-Initiated Identifier

The following fields MUST be used to represent the client-initiated identifier: IPv4/IPv6 tunnel sender address, IPv4/IPv6 tunnel endpoint address, Tunnel ID, and Extended Tunnel ID. Based on local policy, the client MAY also include the LSP ID to identify a specific LSP within the tunnel. These fields are defined in Sections 4.6.1.1 and 4.6.2.1 of [RFC3209].

The usage of the client-initiated identifier is illustrated by Figure 1. Suppose an LSP from EN2 to EN4 needs to be diverse with respect to an LSP from EN1 to EN3.

The LSP identifier of the EN1-EN3 LSP is LSP-IDENTIFIER1, where LSP-IDENTIFIER1 is defined by the tuple

```
(tunnel-id = T1,  
LSP ID = L1,  
source address = EN1.RID (Route Identifier),  
destination address = EN3.RID,  
extended tunnel-id = EN1.RID).
```

Similarly, the LSP identifier of the EN2-EN4 LSP is LSP-IDENTIFIER2, where LSP-IDENTIFIER2 is defined by the tuple

```
(tunnel-id = T2,  
LSP ID = L2,  
source address = EN2.RID,  
destination address = EN4.RID,  
extended tunnel-id = EN2.RID).
```

The EN1-EN3 LSP is signaled with an exclusion requirement from LSP-IDENTIFIER2, and the EN2-EN4 LSP is signaled with an exclusion requirement from LSP-IDENTIFIER1. In order to maintain diversity between these two connections within the core network, the core network SHOULD implement crankback signaling extensions as defined in [RFC4920]. Note that crankback signaling is known to lead to slower setup times and suboptimal paths under some circumstances as described by [RFC4920].

1.4. PCE-Allocated Identifier

In scenarios where a PCE is deployed and used to perform path computation, typically the ingress node of the core network (e.g., node CN1 in Figure 1) could consult a PCE to allocate identifiers, which are used to signal path diversity constraints. In other deployment scenarios, a PCE is deployed at a network node(s) or it is

part of a Network Management System (NMS). In all these cases, the PCE is consulted and 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 Record Route Object (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.

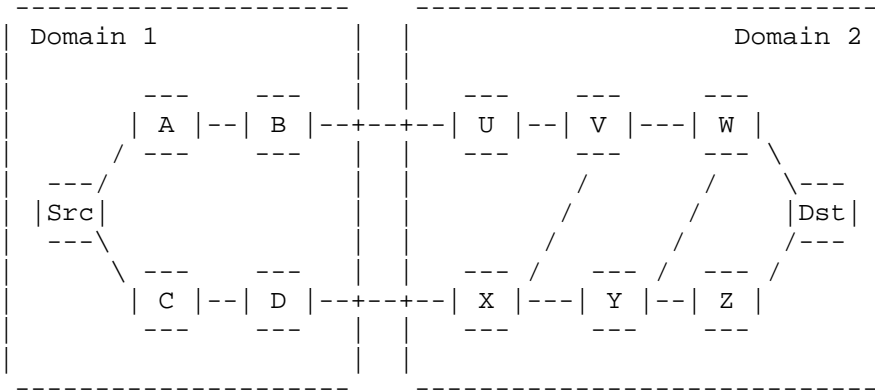


Figure 2: A Simple Multi-domain Network

In order to support LSP diversity, node U consults the PCE and replaces the path segment {U, V, W} in the RRO with a Path Key subobject. The PCE function assigns an "identifier" and puts it into the Path Key field of the Path Key subobject. The PCE ID in the message indicates that this replacement operation was performed by node U.

With this additional information, the source node is able to signal the subsequent LSPs with the ERO set to {C, D, exclude Path Key (signaled in the EXRS RSVP subobject), loose Dst}. When the signaling message reaches node X, it can consult the PCE function associated with node U to expand the Path Key in order to calculate a

path that is diverse with respect to the first LSP. Alternatively, the source node 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 Section 3.1 of [RFC5553]. 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).

1.5. 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 Risk Link Group (SRLG) identifier information is both available and shareable (by policy) with the ENs, the procedure defined in [RFC8001] 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 the core network, existing mechanisms cannot be used.

In this document, 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 defined 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 PAS identifier "123", the CN would use local knowledge of the core network SRLGs associated with the LSPs tagged with PAS attribute "123" and use those SRLGs as constraints

for path computation. If a PAS identifier is used as an exclusion identifier 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 an SRLG-diverse signaled path, the CN may optionally record the core network SRLG information for that connection in terms of CN-based parameters and associate that with the EN addresses in the Path message. Specifically, for Layer 1 Virtual Private Networks (L1VPNs), Port Information Tables (PITs) [RFC5251] can be leveraged to translate between client (EN) addresses and core network addresses.

The means to distribute the PAS information within the core network is beyond the scope of this document. For example, 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. Regardless of means used to distribute the PAS information, the information is kept inside the core network and is not shared with the overlay network (see Figure 1).

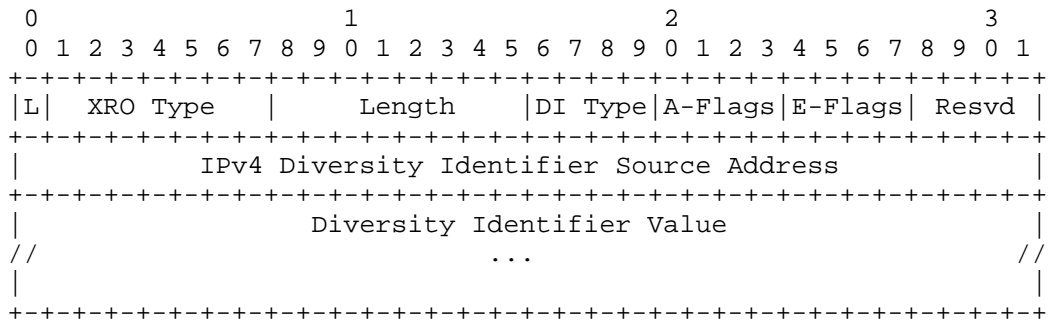
2. RSVP-TE Signaling Extensions

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 below for the IPv4 and IPv6 address families. Most of the fields in the IPv4 and IPv6 Diversity XRO subobjects are common and are described following the definition of the two subobjects.

The IPv4 Diversity XRO subobject is defined as follows:



Similarly, the IPv6 Diversity XRO subobject is defined as follows:



L:
 The L flag is used in the same way as for the XRO subobjects defined in [RFC4874], that is:

- 0 indicates that the diversity constraints MUST be satisfied, and
- 1 indicates that the diversity constraints SHOULD be satisfied.

XRO Type:
 The value is set to 38 for the IPv4 Diversity XRO subobject. The value is set to 39 for the IPv6 Diversity XRO subobject.

Length:
 Per [RFC4874], the Length contains the total length of the IPv4/IPv6 subobject in bytes, including the XRO 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 in the IPv4/IPv6 Diversity subobjects. The following three DI Type values are defined in this document:

DI Type value	Definition
1	Client-Initiated Identifier
2	PCE-Allocated Identifier
3	Network-Assigned Identifier

Attribute Flags (A-Flags):

The Attribute Flags (A-Flags) are used to communicate desirable attributes of the LSP being signaled in the IPv4/IPv6 Diversity subobjects. 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 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. This flag is useful, for example, when an EN is not dual homed (like EN4 in Figure 1, where all LSPs have to go through CN5).

The "Penultimate node exception" flag is typically set when the destination node is single homed (e.g., EN1 or EN4 in Figure 2). In such a case, LSP diversity can only be accomplished inside the core network up to the egress node and the penultimate hop must be the same for the LSPs.

0x08 = LSP ID to be ignored

This flag is used to indicate tunnel-level exclusion. Specifically, this flag is used to indicate that if the diversity identifier contains an LSP ID field, then the LSP ID is to be ignored, and the exclusion applies to any LSP matching the rest of the diversity identifier.

Exclusion Flags (E-Flags):

The Exclusion Flags are used to communicate the desired type(s) of exclusion requested in the IPv4/IPv6 Diversity subobjects. The following flags are defined. Any combination of these flags is permitted. Please note that the exclusion specified by these flags may be modified by the value of the A-Flags. For example, the node exclusion flag is ignored for the penultimate node if the "Penultimate node exception" flag of the A-Flags is set.

0x01 = SRLG exclusion

Indicates that the path of the LSP being signaled is requested to be SRLG disjoint with respect to the excluded path specified by the IPv4/IPv6 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 IPv4/IPv6 Diversity XRO subobject.

0x04 = Link exclusion

Indicates that the path of the LSP being signaled is requested to be "link diverse" from the path specified by the IPv4/IPv6 Diversity XRO subobject.

0x08 = Reserved

This flag is reserved. It MUST be set to zero on transmission and MUST be ignored on receipt for both IPv4/IPv6 Diversity XRO subobjects.

Resvd:

This field is reserved. It MUST be set to zero on transmission and MUST be ignored on receipt for both IPv4/IPv6 Diversity XRO subobjects.

IPv4/IPv6 Diversity Identifier Source Address:


This field MUST be set to the IPv4/IPv6 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:

- * When the Diversity Identifier Type is set to the "Client-Initiated Identifier", the value MUST be set to IPv4/IPv6 tunnel sender address of the reference LSP against which diversity is desired. The IPv4/IPv6 tunnel sender address is as defined in [RFC3209].

- * When the Diversity Identifier Type is set to "PCE-Allocated Identifier", the value MUST be set to the IPv4/IPv6 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]. The PCE ID is carried in the Diversity Identifier Source Address field of the subobject.
- * When the Diversity Identifier Type is set to "Network-Assigned Identifier", the value MUST be set to the IPv4/IPv6 address of the node allocating the Path Affinity Set (PAS).

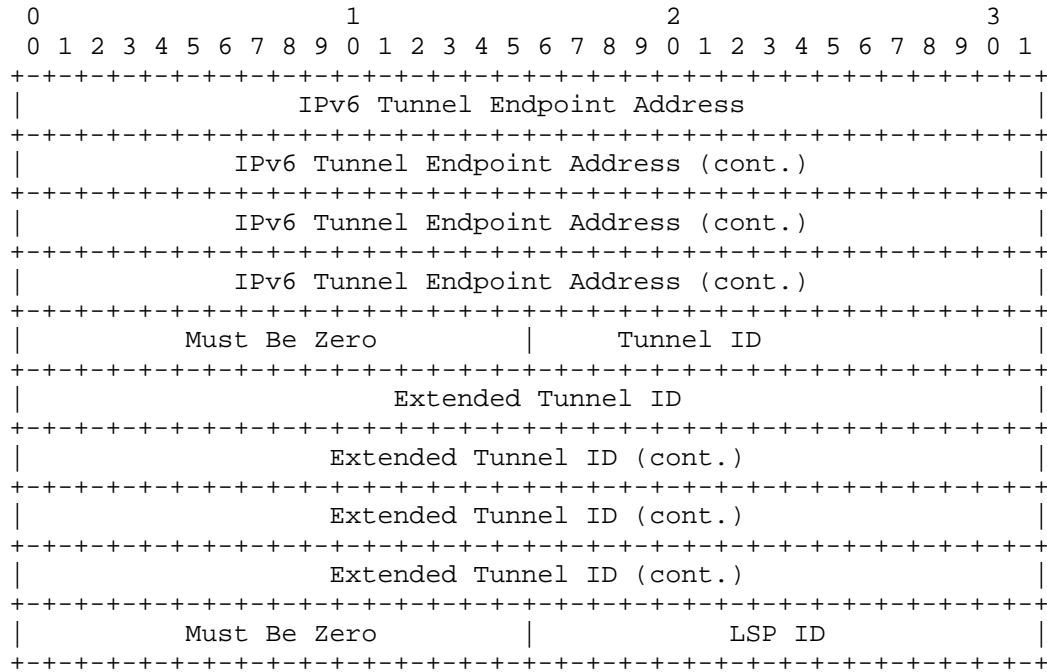
Diversity Identifier Value: Encoding for this field depends on the Diversity Identifier Type, as defined in the following.

When the Diversity Identifier Type is set to "Client-Initiated Identifier" in the IPv4 Diversity XRO subobject, the Diversity Identifier Value MUST be encoded as follows:




The IPv4 Tunnel Endpoint Address, Tunnel ID, Extended Tunnel ID, and LSP ID are as defined in [RFC3209].

When the Diversity Identifier Type is set to "Client-Initiated Identifier" in the IPv6 Diversity XRO subobject, the Diversity Identifier Value MUST be encoded as follows:




The IPv6 Tunnel Endpoint Address, Tunnel ID, IPv6 Extended Tunnel ID, and LSP ID are as defined in [RFC3209].

When the Diversity Identifier Type is set to "PCE-Allocated Identifier" in the IPv4 or IPv6 Diversity XRO subobject, the Diversity Identifier Value MUST be encoded as follows:



The Path Key is defined in [RFC5553].

When the Diversity Identifier Type is set to "Network-Assigned Identifier" in the IPv4 or IPv6 Diversity XRO subobject, the Diversity Identifier Value MUST be encoded as follows:



The Path Affinity Set (PAS) Identifier field is a 32-bit value that is scoped by (i.e., is only meaningful when used in combination with) the Diversity Identifier Source Address field. There are no restrictions on how a node selects a PAS identifier value. Section 1.3 defines the PAS term and provides context on how values may be selected.

2.2. 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 a Diversity subobject as specified in this document. The same type values 38 and 39 MUST be used.

2.3. Processing Rules for the Diversity XRO and EXRS Subobjects

The procedure defined in [RFC4874] for processing the XRO and EXRS is not changed by this document. The processing rules for the Diversity XRO and EXRS subobjects are similar unless the differences are explicitly described. Similarly, IPv4 and IPv6 Diversity XRO subobjects and IPv4 and IPv6 Diversity EXRS subobjects follow the same processing rules.

If the processing node cannot recognize the Diversity XRO/EXRS subobject, the node is expected to follow the procedure defined in [RFC4874].

An XRO/EXRS object MAY contain multiple Diversity subobjects of the same DI Type. For example, in order to exclude multiple Path Keys, a node MAY include multiple Diversity XRO subobjects, each with a different Path Key. Similarly, in order to exclude the routes taken by multiple LSPs, a node MAY include multiple Diversity XRO/EXRS subobjects, each with a different LSP identifier. Likewise, to exclude multiple PAS identifiers, a node MAY include multiple

Diversity XRO/EXRS subobjects, each with a different PAS identifier. However, all Diversity subobjects in an XRO/EXRS MUST contain the same Diversity Identifier Type. If a Path message contains an XRO/EXRS with multiple Diversity subobjects of different DI Types, the processing node MUST return a PathErr with the error code "Routing Problem" (24) and error sub-code "XRO/EXRS Too Complex" (68/69).

If the processing node recognizes the Diversity XRO/EXRS subobject but does not support the DI Type, it MUST return a PathErr with the error code "Routing Problem" (24) and error sub-code "Unsupported Diversity Identifier Type" (36).

In the case of DI Type "Client-Initiated Identifier", all nodes along the path SHOULD process the diversity information signaled in the XRO/EXRS Diversity subobjects to verify that the signaled diversity constraint is satisfied. If a diversity violation is detected, crankback signaling MAY be initiated.

In the case of DI Type "PCE-Allocated Identifier" and "Network-Assigned Identifier", the nodes in the domain that perform path computation SHOULD process the diversity information signaled in the XRO/EXRS Diversity subobjects as follows. In the PCE case, the ingress node of a domain sends a path computation request for a path from ingress node to egress node, including diversity constraints to a PCE. Or, in the PAS case, the ingress node is capable of calculating the path for the new LSP from ingress node to the egress node, taking the diversity constraints into account. The calculated path is then carried in the Explicit Route Object (ERO). Hence, the transit nodes in a domain and the domain egress node SHOULD NOT process the signaled diversity information unless path computation is performed.

While processing the EXRS object, 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 A-Flags affect the processing of the Diversity XRO/EXRS subobject as follows:

- o When the "Processing node exception" flag is set, the exclusion MUST be ignored for the node processing the XRO or EXRS subobject.
- o When the "Destination node exception" flag is set, the exclusion MUST be ignored for the destination node in processing the XRO subobject. The destination node exception for the EXRS subobject applies to the explicit node identified by the ERO subobject that

identifies the next abstract node. When the "Destination node exception" flag is set in the EXRS subobject, exclusion MUST be ignored for said node (i.e., the next abstract node).

- o When the "Penultimate node exception" flag is set in the XRO subobject, the exclusion MUST be ignored for the penultimate node on the path of the LSP being established.

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. When the "Penultimate node exception" flag is set in the EXRS subobject, the exclusion MUST be ignored for said node (i.e., the node before the next abstract node).

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

- o If the Diversity Identifier Type is set to "Client-Initiated Identifier", the processing node MUST ensure that the path calculated/expanded for the signaled LSP is diverse from the route taken by the LSP identified in the Diversity Identifier Value field.
- o If the Diversity Identifier Type is set to "PCE-Allocated Identifier", 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 Diversity Identifier Source Address in the subobject for route computation. The processing node MAY use the Path Key resolution mechanisms described in [RFC5553].
- o If the Diversity Identifier Type is set to "Network-Assigned Identifier", the processing node MUST ensure that the path calculated for the signaled LSP is diverse with respect to the values associated with the PAS Identifier and Diversity Identifier Source Address fields.
- o 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 is diverse from the requested Exclusion Flags.
- o 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 Resv for the signaled LSP, the

processing node MUST return a PathErr with the error code "Notify Error" (25) and error sub-code "Route of XRO LSP identifier unknown" (14) for the signaled LSP.

- o 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 or Diversity EXRS subobject is set, the processing node proceeds as follows:

- o If the Diversity Identifier Type is set to "Client-Initiated Identifier", the processing node SHOULD ensure that the path calculated/expended for the signaled LSP is diverse from the route taken by the LSP identified in the Diversity Identifier Value field.
- o If the Diversity Identifier Type is set to "PCE-Allocated Identifier", the processing node SHOULD ensure that the path calculated for the signaled LSP is diverse from the route identified by the Path Key.
- o If the Diversity Identifier Type is set to "Network-Assigned Identifier", the processing node SHOULD ensure that the path calculated for the signaled LSP is diverse with respect to the values associated with the PAS Identifier and Diversity Identifier Source Address fields.
- o 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 MUST return a PathErr message with error code "Notify Error" (25) and error sub-code "Failed to satisfy Exclude Route" (15) 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 MUST re-evaluate the exclusion and diversity constraints requested by the diverse LSP to determine whether they are still satisfied.

- o In the case where the L-flag was not set in the initial setup message, the exclusion and diversity constraints were satisfied at the time of the initial setup. If the processing node re-evaluating the exclusion and diversity constraints for a diverse LSP detects that the exclusion and diversity constraints are no

longer met, it 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 Path_State_Removed (PSR) flag [RFC3473] MUST NOT be set. A source node receiving a PathErr message with this error code and sub-code combination SHOULD take appropriate actions and move the diverse LSP to a new path that meets the original constraints.

- o In the case where the L-flag was set in the initial setup message, the exclusion and diversity constraints may or may not be satisfied at any given time. If the exclusion constraints for a diverse LSP were satisfied before, and if the processing node re-evaluating the exclusion and diversity constraints for a diverse LSP detects that exclusion and diversity constraints are no longer met, it MUST send a PathErr message for the diverse LSP with the error code "Notify Error" (25) and error sub-code "Failed to satisfy Exclude Route" (15). The PSR flag MUST NOT be set. The source node MAY take no consequent action and keep the LSP along the path that does not meet the original constraints. Similarly, if the exclusion constraints for a diverse LSP were not satisfied before, and if the processing node re-evaluating the exclusion and diversity constraints for a diverse LSP detects that the exclusion constraints are met, it MUST send a PathErr message for the diverse LSP with the error code "Notify Error" (25) and a new error sub-code "Compliant path exists" (16). The PSR flag MUST NOT be set. A source node receiving a PathErr message with this error code and sub-code combination MAY move the diverse LSP to a new path that meets the original constraints.

3. Security Considerations

This document does not introduce any additional security issues in addition to those identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473], [RFC2747], [RFC4874], [RFC5520], and [RFC5553].

The diversity mechanisms defined in this document rely on the new diversity subobject that is carried in the XRO or EXRS, respectively. In Section 7 of [RFC4874], it is noted that some administrative boundaries may remove the XRO due to security concerns on explicit route information exchange. However, when the diversity subobjects specified in this document are used, removing at the administrative boundary an XRO containing these diversity subobjects would result in the request for diversity being dropped at the boundary, and path computation would be unlikely to produce the requested diverse path. As such, diversity subobjects MUST be retained in an XRO crossing an administrative boundary, even if other subobjects are removed. This

retention would be based on operator policy. The use of diversity subobjects is based on mutual agreement. This avoids the need to share the identity of network resources when supporting diversity.

4. IANA Considerations

IANA has assigned new values defined in this document and summarized in this section.

4.1. New XRO Subobject Types

In the IANA registry for RSVP parameters, under "Class Names, Class Numbers, and Class Types", this document defines two new subobjects for the EXCLUDE_ROUTE object [RFC4874], C-Type 1 (see "Class Types or C-Types - 232 EXCLUDE_ROUTE" on <<https://www.iana.org/assignments/rsvp-parameters>>).

Description	Value
IPv4 Diversity	38
IPv6 Diversity	39

4.2. New EXRS Subobject Types

The Diversity XRO subobjects are also defined as new EXRS subobjects (see "Class Types or C-Types - 20 EXPLICIT_ROUTE" on <<https://www.iana.org/assignments/rsvp-parameters>>). The same numeric values have been assigned:

Description	Value
IPv4 Diversity	38
IPv6 Diversity	39

4.3. New RSVP Error Sub-codes

In the IANA registry for RSVP parameters, under "Error Codes and Globally Defined Error Value Sub-Codes", for Error Code "Routing Problem" (24) (see [RFC3209]), the following sub-codes are defined (see "Sub-Codes - 24 Routing Problem" on <https://www.iana.org/assignments/rsvp-parameters>).

Value	Description	Reference
36	Unsupported Diversity Identifier Type	RFC 8390

For Error Code "Notify Error" (25) (see [RFC3209]), the following sub-codes are defined (see "Sub-Codes - 25 Notify Error" on <https://www.iana.org/assignments/rsvp-parameters>).

Value	Description	Reference
14	Route of XRO LSP identifier unknown	RFC 8390
15	Failed to satisfy Exclude Route	RFC 8390
16	Compliant path exists	RFC 8390

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, DOI 10.17487/RFC2747, January 2000, <https://www.rfc-editor.org/info/rfc2747>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <<https://www.rfc-editor.org/info/rfc3473>>.
- [RFC4202] Kompella, K., Ed. and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, DOI 10.17487/RFC4202, October 2005, <<https://www.rfc-editor.org/info/rfc4202>>.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, DOI 10.17487/RFC4874, April 2007, <<https://www.rfc-editor.org/info/rfc4874>>.
- [RFC4920] Farrel, A., Ed., Satyanarayana, A., Iwata, A., Fujita, N., and G. Ash, "Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE", RFC 4920, DOI 10.17487/RFC4920, July 2007, <<https://www.rfc-editor.org/info/rfc4920>>.
- [RFC5553] Farrel, A., Ed., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, DOI 10.17487/RFC5553, May 2009, <<https://www.rfc-editor.org/info/rfc5553>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

5.2. Informative References

- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<https://www.rfc-editor.org/info/rfc2205>>.
- [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, DOI 10.17487/RFC4208, October 2005, <<https://www.rfc-editor.org/info/rfc4208>>.
- [RFC5251] Fedyk, D., Ed., Rekhter, Y., Ed., Papadimitriou, D., Rabbat, R., and L. Berger, "Layer 1 VPN Basic Mode", RFC 5251, DOI 10.17487/RFC5251, July 2008, <<https://www.rfc-editor.org/info/rfc5251>>.

- [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, DOI 10.17487/RFC5520, April 2009, <<https://www.rfc-editor.org/info/rfc5520>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC8001] Zhang, F., Ed., Gonzalez de Dios, O., Ed., Margaria, C., Hartley, M., and Z. Ali, "RSVP-TE Extensions for Collecting Shared Risk Link Group (SRLG) Information", RFC 8001, DOI 10.17487/RFC8001, January 2017, <<https://www.rfc-editor.org/info/rfc8001>>.

Acknowledgements

The authors would like to thank Xihua Fu for his contributions. The authors also would like to thank Luyuan Fang and Walid Wakim for their review and comments.

Contributors

Igor Bryskin
Huawei Technologies
Email: Igor.Bryskin@huawei.com

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

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

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

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

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

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

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

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

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

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

Lieven Levrau
Nokia
Email: Lieven.Levrau@nokia.com

Cyril Margaria
Email: 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 (editor)
Cisco Systems.

Email: zali@cisco.com

George Swallow (editor)
Southend Technical Center

Email: swallow.ietf@gmail.com

Fatai Zhang (editor)
Huawei Technologies

Email: zhangfatai@huawei.com

Dieter Beller (editor)
Nokia

Email: Dieter.Beller@nokia.com

