

TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires February 24, 2017

F. Zhang, Ed.
Huawei
O. Gonzalez de Dios, Ed.
Telefonica Global CTO
M. Hartley
Z. Ali
Cisco
C. Margaria

August 24, 2016

RSVP-TE Extensions for Collecting SRLG Information
draft-ietf-teas-rsvp-te-srlg-collect-08

Abstract

This document provides extensions for the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE), including GMPLS, to support automatic collection of Shared Risk Link Group (SRLG) information for the TE link formed by a Label Switched Path (LSP).

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 February 24, 2017.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

1. Introduction	2
1.1. Applicability Example: Dual Homing	3
2. Requirements Language	5
3. RSVP-TE Requirements	5
3.1. SRLG Collection Indication	5
3.2. SRLG Collection	5
3.3. SRLG Update	5
3.4. SRLG ID definition	6
4. Encodings	6
4.1. SRLG Collection Flag	6
4.2. RRO SRLG sub-object	6
5. Signaling Procedures	8
5.1. SRLG Collection	8
5.2. SRLG Update	10
5.3. Domain Boundaries	10
5.4. Compatibility	10
6. Manageability Considerations	10
6.1. Policy Configuration	11
6.2. Coherent SRLG IDs	11
7. Security Considerations	11
8. IANA Considerations	11
8.1. RSVP Attribute Bit Flags	11
8.2. ROUTE_RECORD Object	12
8.3. Policy Control Failure Error subcodes	12
9. Contributors	12
10. Acknowledgements	13
11. References	13
11.1. Normative References	13
11.2. Informative References	13
Authors' Addresses	14

1. Introduction

It is important to understand which Traffic Engineering (TE) links in the network might be at risk from the same failures. In this sense, a set of links can constitute a 'shared risk link group' (SRLG) if they share a resource whose failure can affect all links in the set [RFC4202].

On the other hand, as described in [RFC4206] and [RFC6107], H-LSP (Hierarchical LSP) or S-LSP (stitched LSP) can be used for carrying one or more other LSPs. Both of the H-LSP and S-LSP can be formed as a TE link. In such cases, it is important to know the SRLG information of the LSPs that will be used to carry further LSPs.

This document provides a signaling mechanism to collect the SRLGs used by a LSP, which can then be advertised as properties of the TE-link formed by that LSP.

1.1. Applicability Example: Dual Homing

An interesting use case for the SRLG collection procedures defined in this document is achieving LSP diversity in a dual homing scenario. The use case is illustrated in Figure 1, when the overlay model is applied as defined in RFC 4208 [RFC4208]. In this example, the exchange of routing information over the User-Network Interface (UNI) is prohibited by operator policy.




Figure 1: Dual Homing Configuration

Single-homed customer edge (CE) devices are connected to a single provider edge (PE) device via a single UNI link (which could be a bundle of parallel links, typically using the same fiber cable). This single UNI link can constitute a single point of failure. Such a single point of failure can be avoided if the CE device is connected to two PE devices via two UNI interfaces as depicted in Figure 1 above for CE1 and CE2, respectively.

For the dual-homing case, it is possible to establish two connections

(LSPs) from the source CE device to the same destination CE device where one connection is using one UNI link to PE1, for example, and the other connection is using the UNI link to PE2. In order to avoid single points of failure within the provider network, it is necessary to also ensure path (LSP) diversity within the provider network in order to achieve end-to-end diversity for the two LSPs between the two CE devices CE1 and CE2. This use case describes how it is possible to achieve path diversity within the provider network based on collected SRLG information. As the two connections (LSPs) enter the provider network at different PE devices, the PE device that receives the connection request for the second connection needs to know the additional path computation constraints such that the path of the second LSP is disjoint with respect to the already established first connection.

As SRLG information is normally not shared between the provider network and the client network, i.e., between PE and CE devices, the challenge is how to solve the diversity problem when a CE is dual-homed. The RSVP extensions for collecting SRLG information defined in this document make it possible to retrieve SRLG information for an LSP and hence solve the dual-homing LSP diversity problem. For example, CE1 in Figure 1 may have requested an LSP1 to CE2 via PE1 that is routed via PE3 to CE2. CE1 can then subsequently request an LSP2 to CE2 via PE2 with the constraint that it needs to be maximally SRLG disjoint with respect to LSP1. PE2, however, does not have any SRLG information associated with LSP1, which is needed as input for its constraint-based path computation function. If CE1 is capable of retrieving the SRLG information associated with LSP1 from PE1, it can pass this discovered information to PE2 as part of the LSP2 setup request (RSVP PATH message) in an EXCLUDE_ROUTE Object (XRO) or Explicit Exclusion Route Subobject (EXRS) as described in [RFC4874], and PE2 can now calculate a path for LSP2 that is SRLG disjoint with respect to LSP1. The SRLG information associated with LSP1 can be retrieved when LSP1 is established or at any time before LSP2 is setup.

When CE1 sends the setup request for LSP2 to PE2, it can also request the collection of SRLG information for LSP2 and send that information to PE1 by re-signaling LSP1 with SRLG-exclusion based on LSP2's discovered SRLGs. This will ensure that the two paths for the two LSPs remain mutually diverse, which is important when the provider network is capable of restoring connections that failed due to a network failure (fiber cut) in the provider network.

Note that the knowledge of SRLG information even for multiple LSPs does not allow a CE device to derive the provider network topology based on the collected SRLG information. It would, however, be possible for an entity controlling multiple CE devices to derive some

information related to the topology. This document therefore allows PE devices to control the communication of SRLGs outside the provider network if desired.

2. 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 RFC 2119 [RFC2119].

3. RSVP-TE Requirements

The SRLG-collection process takes place in three stages:

- o The LSP's ingress node requests that SRLG collection take place;
- o SRLG data is added to the Path and Resv ROUTE_RECORD Objects (RROs) by all nodes during signaling;
- o Changes to previously-signaled SRLG data are made by sending updated Path and Resv messages as required.

3.1. SRLG Collection Indication

The ingress node of the LSP needs be capable of indicating whether the SRLG information of the LSP is to be collected during the signaling procedure of setting up an LSP. There is no need for SRLG information to be collected without an explicit request for it being made by the ingress node.

It may be preferable for the SRLG collection request to be understood by all nodes along the LSP's path, or it may be more important for the LSP to be established successfully even if it traverses nodes that cannot supply SRLG information or have not implemented the procedures specified in this document. It is desirable for the ingress node to make the SRLG collection request in a manner that best suits its own policy.

3.2. SRLG Collection

If requested, the SRLG information is collected during the setup of an LSP. SRLG information is added by each hop to the Path RRO during Path message processing. The same information is also added to the Resv RRO during Resv processing at each hop.

3.3. SRLG Update

When the SRLG information of an existing LSP for which SRLG information was collected during signaling changes, the relevant nodes of the LSP need to be capable of updating the SRLG information of the LSP. This means that the signaling procedure needs to be capable of updating the new SRLG information.

3.4. SRLG ID definition

The identifier of an SRLG (SRLG ID) is defined as a 32-bit quantity in [RFC4202]. This definition is used in this document.

4. Encodings

4.1. SRLG Collection Flag

In order to indicate to nodes that SRLG collection is desired, this document defines a new flag in the Attribute Flags TLV (see RFC 5420 [RFC5420]). This document defines a new SRLG collection flag in the Attribute Flags TLV (see RFC 5420 [RFC5420]). A node that wishes to indicate that SRLG collection is desired MUST set this flag in an Attribute Flags TLV in an LSP_REQUIRED_ATTRIBUTES Object if collection is to be mandatory, or an LSP_ATTRIBUTES Object if collection is desired but not mandatory.


- o Bit Number (specified in Section 8.1): SRLG Collection flag

The SRLG Collection flag is meaningful on a Path message. If the SRLG Collection flag is set to 1, it means that the SRLG information SHOULD be reported to the ingress and egress node along the setup of the LSP.

The rules for the processing of the Attribute Flags TLV are not changed.

4.2. RRO SRLG sub-object

This document defines a new RRO sub-object (ROUTE_RECORD sub-object) to record the SRLG information of the LSP. Its format is modeled on the RRO sub-objects defined in RFC 3209 [RFC3209].



```

~                               .....                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SRLG ID n (4 octets)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type (8 bits)

The type of the sub-object. The value is specified in Section 8.2.

Length (8 bits)

The Length field contains the total length of the sub-object in octets, including the Type and Length fields. The Length depends on the number of SRLG IDs.

Direction bit (D-bit) (1 bit)

If not set, the SRLGs contained in this sub-object apply to the downstream direction. If set, they apply to the upstream direction.

Reserved (15 bits)

This 15-bit field is reserved. It SHOULD be set to zero on transmission and MUST be ignored on receipt.

SRLG ID (4 octets)

This field contains one SRLG ID. There is one SRLG ID field per SRLG collected. There MAY be multiple SRLG ID fields in an SRLG sub-object.

A node MUST NOT push a SRLG sub-object in the RECORD_ROUTE without also pushing either a IPv4 sub-object, a IPv6 sub-object, a Unnumbered Interface ID sub-object or a Path Key sub-object.

As described in RFC 3209 [RFC3209], the RECORD_ROUTE object is managed as a stack. The SRLG sub-object MUST be pushed by the node before the node IP address or link identifier. The SRLG-sub-object SHOULD be pushed after the Attribute sub-object, if present, and after the LABEL sub-object, if requested. It MUST be pushed within the hop to which it applies.

RFC 5553 [RFC5553] describes mechanisms to carry a PKS (Path Key Sub-object) in the RRO so as to facilitate confidentiality in the signaling of inter-domain TE LSPs, and allows the path segment that needs to be hidden (that is, a Confidential Path Segment (CPS)) to be replaced in the RRO with a PKS. If the CPS contains SRLG Sub-objects, these MAY be retained in the RRO by adding them again after

the PKS Sub-object in the RRO. The CPS is defined in RFC 5520 [RFC5520].

The rules for the processing of the LSP_REQUIRED_ATTRIBUTES, LSP_ATTRIBUTE and ROUTE_RECORD Objects are not changed.

5. Signaling Procedures

The ingress node of the LSP MUST be capable of indicating whether the SRLG information of the LSP is to be collected during the signaling procedure of setting up an LSP.

5.1. SRLG Collection

Per RFC 3209 [RFC3209], an ingress node initiates the recording of the route information of an LSP by adding a RRO to a Path message. If an ingress node also desires SRLG recording, it MUST set the SRLG Collection Flag in the Attribute Flags TLV which MAY be carried either in an LSP_REQUIRED_ATTRIBUTES Object when the collection is mandatory, or in an LSP_ATTRIBUTES Object when the collection is desired, but not mandatory.

A node MUST NOT add SRLG information without an explicit request for it being made by the ingress node in the Path message.

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object with the SRLG Collection Flag set, if local policy determines that the SRLG information is not to be provided to the endpoints, it MUST return a PathErr message with:

- o Error Code 2 (policy) and
- o Error subcode "SRLG Recording Rejected" (see Section 8.3 for value)

to reject the Path message.

When a node receives a Path message which carries an LSP_ATTRIBUTES Object with the SRLG Collection Flag set, if local policy determines that the SRLG information is not to be provided to the endpoints, the Path message MUST NOT be rejected due to the SRLG recording restriction and the Path message MUST be forwarded without any SRLG sub-object(s) added to the RRO of the corresponding outgoing Path message.

If local policy permits the recording of the SRLG information, the processing node SHOULD add local SRLG information, as defined below, to the RRO of the corresponding outgoing Path message. The processing node MAY add multiple SRLG sub-objects to the RRO if necessary. It then forwards the Path message to the next node in the downstream direction. The processing node MUST retain a record of the

SRLG recording request for reference during Resv processing described below.

If the addition of SRLG information to the RRO would result in the RRO exceeding its maximum possible size or becoming too large for the Path message to contain it, the requested SRLGs MUST NOT be added. If the SRLG collection request was contained in an LSP_REQUIRED_ATTRIBUTES Object, the processing node MUST behave as specified by RFC 3209 [RFC3209] and drop the RRO from the Path message entirely. If the SRLG collection request was contained in an LSP_ATTRIBUTES Object, the processing node MAY omit some or all of the requested SRLGs from the RRO; otherwise it MUST behave as specified by [RFC3209] and drop the RRO from the Path message entirely. Subsequent processing of the LSP proceeds as further specified in RFC 3209 [RFC3209].

Following the steps described above, the intermediate nodes of the LSP can collect the SRLG information in the RRO during the processing of the Path message hop by hop. When the Path message arrives at the egress node, the egress node receives SRLG information in the RRO.

Per RFC 3209 [RFC3209], when issuing a Resv message for a Path message which contains an RRO, an egress node initiates the RRO process by adding an RRO to the outgoing Resv message. The processing for RROs contained in Resv messages then mirrors that of the Path messages.

When a node receives a Resv message for an LSP for which SRLG Collection was specified in the corresponding Path message, then when local policy allows recording SRLG information, the node MUST add SRLG information to the RRO of the corresponding outgoing Resv message as specified below. When the Resv message arrives at the ingress node, the ingress node can extract the SRLG information from the RRO in the same way as the egress node.

Note that a link's SRLG information for the upstream direction cannot be assumed to be the same as that in the downstream.

- o For Path and Resv messages for a unidirectional LSP, a node SHOULD include SRLG sub-objects in the RRO for the downstream data link only.
- o For Path and Resv messages for a bidirectional LSP, a node SHOULD include SRLG sub-objects in the RRO for both the upstream data link and the downstream data link from the local node. In this case, the node MUST include the information in the same order for both Path messages and Resv messages. That is, the SRLG sub-object for the upstream link is added to the RRO before the SRLG

sub-object for the downstream link.

If SRLG data is added for both the upstream and downstream links, the two sets of SRLG data MUST be added in separate SRLG sub-objects. A single SRLG sub-object MUST NOT contain a mixture of upstream and downstream SRLGs. When adding a SRLG sub-object to an RRO, the D-bit MUST be set appropriately to indicate the direction of the SRLGs. If an SRLG ID applies in both directions, it SHOULD be added to both the upstream and downstream SRLG sub-objects.

Based on the above procedure, the endpoints can get the SRLG information automatically. Then the endpoints can for instance advertise it as a TE link to the routing instance based on the procedure described in [RFC6107] and configure the SRLG information of the Forwarding Adjacency (FA) automatically.

5.2. SRLG Update

When the SRLG information of a link is changed, the endpoints of LSPs using that link need to be made aware of the changes. When a change to the set of SRLGs associated with a link occurs, the procedures defined in Section 4.4.3 of RFC 3209 [RFC3209] MUST be used to refresh the SRLG information for each affected LSP if the SRLG change is to be communicated to other nodes according to the local node's policy.

5.3 Domain Boundaries

If mandated by local policy as specified by the network operator, a node MAY remove SRLG information from any RRO in a Path or Resv message being processed. It MAY add a summary of the removed SRLGs or map them to other SRLG values. However, this SHOULD NOT be done unless explicitly mandated by local policy.

5.4. Compatibility

A node that does not recognize the SRLG Collection Flag in the Attribute Flags TLV is expected to proceed as specified in RFC 5420 [RFC5420]. It is expected to pass the TLV on unaltered if it appears in a LSP_ATTRIBUTES object, or reject the Path message with the appropriate Error Code and Value if it appears in a LSP_REQUIRED_ATTRIBUTES object.

A node that does not recognize the SRLG RRO sub-object is expected to behave as specified in RFC 3209 [RFC3209]: unrecognized sub-objects are to be ignored and passed on unchanged.

6. Manageability Considerations

6.1. Policy Configuration

In a border node of inter-domain or inter-layer network, the following SRLG processing policy MUST be capable of being configured:

- o Whether the node is allowed to participate in SRLG collection and notify changes to collected SRLG information to endpoint nodes as described in section 5.2.
- o Whether the SRLG IDs of the domain or specific layer network can be exposed to the nodes outside the domain or layer network, or whether they SHOULD be summarized, mapped to values that are comprehensible to nodes outside the domain or layer network, or removed entirely as described in section 5.3.

A node using RFC 5553 [RFC5553] and PKS MAY apply the same policy.

6.2. Coherent SRLG IDs

In a multi-layer multi-domain scenario, SRLG IDs can be configured by different management entities in each layer/domain. In such scenarios, maintaining a coherent set of SRLG IDs is a key requirement in order to be able to use the SRLG information properly. Thus, SRLG IDs SHOULD be unique. Note that current procedure is targeted towards a scenario where the different layers and domains belong to the same operator, or to several coordinated administrative groups. Ensuring the aforementioned coherence of SRLG IDs is beyond the scope of this document.

Further scenarios, where coherence in the SRLG IDs cannot be guaranteed are out of the scope of the present document and are left for further study.

7. Security Considerations

This document builds on the mechanisms defined in [RFC3473], which also discusses related security measures. In addition, [RFC5920] provides an overview of security vulnerabilities and protection mechanisms for the GMPLS control plane. The procedures defined in this document permit the transfer of SRLG data between layers or domains during the signaling of LSPs, subject to policy at the layer or domain boundary. As described in section 5.3 and section 6.1, local policy as specified by the network operator will explicitly mandate the processing of information at domain or layer boundaries.

8. IANA Considerations

8.1. RSVP Attribute Bit Flags

IANA has created a registry and manages the space of the Attribute bit flags of the Attribute Flags TLV, as described in section 11.3 of RFC 5420 [RFC5420], in the "Attribute Flags" section of the "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters" registry located in <http://www.iana.org/assignments/rsvp-te-parameters>.

This document introduces a new Attribute Bit Flag:

Bit No	Name	Attribute Flags Path	Attribute Flags Resv	ERO	RRO	Reference
-----	-----	-----	-----	---	---	-----
TBD; suggested value: 12	SRLG Collection Flag	Yes	No	No	Yes	This I-D

8.2. ROUTE_RECORD Object

IANA manages the "RSVP PARAMETERS" registry located at <http://www.iana.org/assignments/rsvp-parameters>. This document introduces a new RRO sub-object:

Value	Description	Reference
-----	-----	-----
TBD; suggested value: 34	SRLG sub-object	This I-D

8.3. Policy Control Failure Error subcodes

IANA manages the assignments in the "Error Codes and Globally-Defined Error Value Sub-Codes" section of the "RSVP PARAMETERS" registry located at <http://www.iana.org/assignments/rsvp-parameters>.

This document introduces a new Policy Control Failure Error sub-code:

Value	Description	Reference
-----	-----	-----
TBD; suggested value: 21	SRLG Recording Rejected	This I-D

9. Contributors

Dan Li
Huawei
F3-5-B RD Center
Bantian, Longgang District, Shenzhen 518129

P.R.China
Email: danli@huawei.com

10. Acknowledgements

The authors would like to thank Dieter Beller, Vishnu Pavan Beeram, Lou Berger, Deborah Brungard, Igor Bryskin, Ramon Casellas, Niclas Comstedt, Alan Davey, Elwyn Davies, Dhruv Dhody, Himanshu Shah and Xian Zhang for their useful comments and improvements to this document.

11. References

11.1. 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.
- [RFC4202] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.
- [RFC5553] Farrel, A., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, May 2009.

11.2. Informative References

- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching

(GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

- [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.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [RFC6107] Shiimoto, K. and A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", RFC 6107, February 2011.

Authors' Addresses

Fatai Zhang (editor)
Huawei
F3-5-B RD Center
Bantian, Longgang District, Shenzhen 518129
P.R.China
Email: zhangfatai@huawei.com

Oscar Gonzalez de Dios (editor)
Telefonica Global CTO
Distrito Telefonica, edificio sur, Ronda de la Comunicacion 28045
Madrid 28050
Spain
Phone: +34 913129647
Email: oscar.gonzalezdedios@telefonica.com

Cyril Margaria
Suite 4001, 200 Somerset Corporate Blvd.
Bridgewater, NJ 08807
US
Email: cyril.margaria@gmail.com

Matt Hartley
Cisco
Email: mhartley@cisco.com

Zafar Ali
Cisco
Email: zali@cisco.com