

TEAS Working Group
Internet Draft
Intended status: Standard Track
Expires: Jan 22, 2020

Z. Ali, Ed.
C. Filsfils
Cisco Systems
J. Meuric
France telecom
K. Kumaki
KDDI Corporation
R. Kunze
Deutsche Telekom AG
July 22, 2019

**Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
extension for recording TE Metric of a Label Switched Path
draft-ietf-teas-te-metric-recording-09**

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 <https://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 Jan 22, 2020.

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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

There are many scenarios in which Traffic Engineering (TE) metrics such as cost, delay and delay variation associated with the TE link formed by Label Switched Path (LSP) are not available to the ingress and egress nodes. This draft provides extensions for the Resource ReserVation Protocol- Traffic Engineering (RSVP-TE) to support automatic collection of cost, delay and delay variation information for the TE link formed by a LSP.

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	3
1.1. Use Cases	4
1.1.1. Inter-domain TE LSPs	4
1.1.2. Inter-area tunnels with loose-hops	4
2. RSVP-TE Requirement	4
2.1. Cost, Delay and Delay Variation Collection Indication	4
2.2. Cost, Delay and Delay Variation Collection	5
2.3. Cost, Delay and Delay Variation Update	5
2.4. Cost Definition	5
3. Encoding	5
3.1. Cost, Delay and Delay Variation Collection Flags	5
3.2. RRO Cost Subobject	6
3.3. RRO Delay Subobject	7
3.4. RRO Delay Variation Subobject	7
4. Signaling Procedures	8
4.1. Cost, Delay and Delay Variation Collection	9
4.2. Metric Update	12

4.3.	Domain Boundaries	12
4.4.	Endpoint processing	12
4.5.	Compatibility	13

5. Manageability Considerations	13
5.1. Policy Configuration	13
6. Security Considerations	14
7. IANA Considerations	14
7.1. RSVP Attribute Bit Flags	14
7.2. ROUTE_RECORD subobject	15
7.3. Policy Control Failure Error subcodes	15
8. References	16
8.1. Normative References	16
8.2. Informative References	16
Acknowledgements	16
Contributors	17
Authors' Addresses	17

1. Introduction

In certain networks, such as financial information networks, network performance information (e.g. delay, delay variation) is becoming as critical to data path selection as other metrics [[RFC7471](#)], [[RFC7810](#)]. If cost, delay or delay variation associated with a Forwarding Adjacency (FA) or a Routing Adjacency (RA) LSP is not available to the ingress or egress node, it cannot be advertised as an attribute of the TE link (FA or RA). There are scenarios in packet and optical networks where the route information of an LSP may not be provided to the ingress node for confidentiality reasons and/or the ingress node may not run the same routing instance as the intermediate nodes traversed by the path. Similarly, there are scenarios in which measuring delay and/ or delay variation on a TE link formed by a LSP is not supported. In such scenarios, the ingress node cannot determine the cost, delay and delay variation properties of the LSP's route.

One possible way to address this issue is to configure cost, delay and delay variation values manually. However, in the event of an LSP being rerouted (e.g. due to re-optimization), such configuration information may become invalid. Consequently, in cases where an LSP is advertised as a TE-Link, the ingress and/or egress nodes cannot provide the correct delay, delay variation and cost information associated with the TE-Link automatically.

In summary, there is a requirement for the ingress and egress nodes to learn the cost, delay and delay variation information of the TE link formed by a LSP. This document provides a mechanism to collect the cost, delay and delay variation

information of a LSP, which can then be advertised as properties of the TE-link formed by that LSP.

1.1. Use Cases

This section describes some of the use cases for the TE metric recording.

1.1.1. Inter-domain TE LSPs

The cost, delay, delay-variation information of a LSP collected by procedures defined in this document can be advertised as properties of TE-link formed by that LSP. This information is very helpful in multi-domain or multi-layer network. A GMPLS User-Network Interface (UNI) [[RFC4208](#)] is also an example of such networks.

1.1.2. Inter-area tunnels with loose-hops

When a LSP is established over multiple IGP-areas using loose hops in the ERO, the ingress node may only have knowledge of the first IGP-area traversed by the LSP. In this case, it cannot determine the cost, delay and delay variation properties of the LSP path.

2. RSVP-TE Requirement

This section outlines RSVP-TE requirements for the support of the automatic collection of cost, delay and delay variation information of an LSP.

As RSVP-TE requirements for cost, delay and delay variation collection are similar, many parts of this section are written such that they apply equally to cost, delay and delay variation collection. There is also very strong similarity of these RSVP-requirements with SRLG recording [[RFC8001](#)].

The Cost, Delay, Delay variation collection process takes place in three stages:

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

2.1. Cost, Delay and Delay Variation Collection Indication

The ingress node of the LSP needs be capable of indicating whether the cost and/or delay and/ or delay variation information of the LSP is to be collected during the signaling procedure of setting up an LSP. A separate collection indication flag for each of these attributes is required. There is no need for cost and/or delay and/ or delay variation to be collected without an explicit request for it being made by the ingress node.

It may be preferable for the cost and/ or delay and/ or delay variation 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 the requested information or have not implemented the procedures specified in this document. It is desirable for the ingress node to make the cost, delay and delay variation collection request in a manner that best suits its own policy.

2.2. Cost, Delay and Delay Variation Collection

If requested, the cost and/or delay and/ or delay variation information is collected during the setup of an LSP. Each of the cost, delay or delay variation can be collected independently. Cost and/ or delay and/ or delay variation information is added by each hop to the Path RRO during Path message processing. The corresponding information is also added to the Resv RRO during Resv processing at each hop.

2.3. Cost, Delay and Delay Variation Update

When the cost and/or delay and/ or delay variation information of an existing LSP for which corresponding information was collected during signaling changes, the relevant nodes of the LSP need to be capable of updating the associated information of the LSP. This means that the signaling procedure needs to be capable of updating the new cost and/or delay and/ or delay variation information.

2.4. Cost Definition

Although the terms delay and delay variation are well understood, "cost" may be ambiguous; in particular, in the context of a LSP that traverse nodes and links operated by different entities, there may be no common definition of cost. However, there are situations in which the entire LSP may be within a single AS (e.g. inter-area LSPs) in which cost discovery is useful.

The precise meaning and interpretation of numerical costs is a matter for the network operator. For the purposes of this document, two constraints are assumed:

- . A higher cost represents an inferior path.
- . Simple addition of costs for different sections of a path must make sense.

3. Encoding

3.1. Cost, Delay and Delay Variation Collection Flags

In order to indicate nodes that cost and/or Delay and/or Delay variation collection is desired, this document defines the following new flags in the Attribute Flags TLV (see [RFC 5420](#))

Ali, Swallow, Filsfils

Expires Jan 2020

[Page 5]

[RFC5420]). A node that wishes to indicate Cost and/or Delay and/or Delay Variation collection is desired MUST set corresponding flag in Attribute Flags TLV in an LSP_REQUIRED_ATTRIBUTES object (if collection is mandatory) or LSP_ATTRIBUTES Object (if collection is desired but not mandatory):

- Cost Collection flag (Bit number to be assigned by IANA)
- Delay Collection flag (Bit number to be assigned by IANA)
- Delay Variation Collection flag (Bit number to be assigned by IANA)

The Cost, Delay and Delay Variation Collection flags are meaningful on a Path message. If the Cost Collection flag is set to 1, it means that the cost information SHOULD be reported to the ingress and egress node along the setup of the LSP. Similarly, if the Delay Collection flag is set to 1, it means that the Delay information SHOULD be reported to the ingress and egress node along the setup of the LSP. Likewise, if the Delay Variation Collection flag is set to 1, it means that the Delay Variation information SHOULD be reported to the ingress and egress node along the setup of the LSP.

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

3.2. RRO Cost Subobject

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

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   | Length |D| Reserved (must be zero) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Cost                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: The type of the sub-object (value to be assigned by IANA).

Length: The Length field contains the total length of the sub-object in bytes, including the Type and Length fields. The Length value is set to 8.

Direction bit (D-bit)

If not set, the cost contained in this sub-object applies to the downstream direction. If set, it applies to the upstream direction.

Ali, Swallow, Filsfils

Expires Jan 2020

[Page 6]

Reserved: This field is reserved for future use. It MUST be set to 0 on transmission and MUST be ignored when received.

Cost: Cost of the local TE link along the route of the LSP.

3.3. RRO Delay Subobject

This document defines a new RRO sub-object (ROUTE_RECORD sub-object) to record the delay information of the LSP. Its format is modeled on the RRO sub-objects defined in [RFC 3209](https://datatracker.ietf.org/doc/rfc3209) [[RFC3209](https://datatracker.ietf.org/doc/rfc3209)].

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
Type								Length								D Reserved (must be zero)															
A Reserved																Delay															

Type: The type of the sub-object (value to be assigned by IANA).

Length: The Length field contains the total length of the sub-object in bytes, including the Type and Length fields. The Length value is set to 8.

Direction bit (D-bit)

If not set, the Delay contained in this sub-object applies to the downstream direction. If set, it applies to the upstream direction.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay respectively, as defined in [RFC 7471](https://datatracker.ietf.org/doc/rfc7471) [[RFC7471](https://datatracker.ietf.org/doc/rfc7471)].

Reserved: These fields are reserved for future use. They MUST be set to 0 when sent and MUST be ignored when received.

Delay: Delay of the local TE link along the route of the LSP, encoded as 24-bit integer, as defined in [RFC 7471](https://datatracker.ietf.org/doc/rfc7471) [[RFC7471](https://datatracker.ietf.org/doc/rfc7471)]. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

3.4. RRO Delay Variation Subobject

This document defines a new RRO sub-object (ROUTE_RECORD sub-

object) to record the delay variation information of the LSP.

Ali, Swallow, Filsfils

Expires Jan 2020

[Page 7]

Its format is modeled on the RRO sub-objects defined in [RFC 3209](https://datatracker.ietf.org/doc/rfc3209) [[RFC3209](https://datatracker.ietf.org/doc/rfc3209)].

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type           |  Length       |D|  Reserved (must be zero)  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|A|  Reserved     |                               Delay Variation  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: The type of the sub-object (value to be assigned by IANA).

Length: The Length field contains the total length of the sub-object in bytes, including the Type and Length fields. The Length value is set to 8.

Direction bit (D-bit)

If not set, the Delay Variation contained in this sub-object applies to the downstream direction. If set, it applies to the upstream direction.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay Variation respectively, as defined in [RFC 7471](https://datatracker.ietf.org/doc/rfc7471) [[RFC7471](https://datatracker.ietf.org/doc/rfc7471)].

Reserved: These fields are reserved for future use. It SHOULD be set to 0 when sent and MUST be ignored when received.

Delay Variation: Delay Variation of the local TE link along the route of the LSP, encoded as 24-bit integer, as defined in [RFC 7471](https://datatracker.ietf.org/doc/rfc7471) [[RFC7471](https://datatracker.ietf.org/doc/rfc7471)]. When set to the maximum value 16,777,215 (16.777215 sec), the delay variation is at least that value and may be larger.

4. Signaling Procedures

As signaling procedure for cost, delay and delay variation collection is similar, many parts of this section are written such that they apply equally to cost, delay and delay variation collection. There is also very strong similarity of these procedures with SRLG recording [[RFC8001](https://datatracker.ietf.org/doc/rfc8001)].

The ingress node of the LSP MUST be capable of indicating whether the Cost and/ or Delay and/ or Delay Variation information of the LSP is to be collected during the signaling procedure of setting up an LSP.

A node MUST NOT push Cost and/ or Delay and/ or Delay Variation sub-object(s) in the RECORD_ROUTE without also pushing either an IPv4 sub-object, an IPv6 sub-object, an Unnumbered Interface ID sub-object or a Path Key sub-object or an SRLG sub-object.

As described in [RFC 3209](#) [[RFC3209](#)], the RECORD_ROUTE object is managed as a stack. The Cost and/ or Delay and/ or Delay Variation sub-object(s) SHOULD be pushed by the node before the node IP address or link identifier. These sub-object(s) SHOULD be pushed after the Attribute sub-object, if present, and after the LABEL sub-object, if requested, and after the SRLG sub-object, if requested. These sub-object(s) 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 Cost and/ or Delay and/ or Delay Variation 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 of the processing of the LSP_REQUIRED_ATTRIBUTES, LSP_ATTRIBUTE and ROUTE_RECORD Objects are not changed.

4.1. Cost, Delay and Delay Variation 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 Cost and/or Delay and/or Delay Variation recording, it MUST set the appropriate flag(s) 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.

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object with the Cost Collection Flag set, if local policy determines that the Cost 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 "Cost Recording Rejected" (value to be assigned by IANA)

to reject the Path message. Similarly, when a node receives a

Path message which carries an LSP_REQUIRED_ATTRIBUTES Object
with the Delay Collection Flag set, if local policy determines
Ali, Swallow, Filsfils Expires Jan 2020 [Page 9]

that the Delay 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 "Delay Recording Rejected" (value to be assigned by IANA)

to reject the Path message. Likewise, when a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object with the Delay Variation Collection Flag set, if local policy determines that the Delay Variation 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 "Delay Variation Recording Rejected" (value to be assigned by IANA)

to reject the Path message.

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

If local policy permits the recording of the Cost and/or Delay and/or Delay Variation information, the processing node SHOULD add corresponding information for the local TE link, as defined below, to the RRO of the corresponding outgoing Path message. The A-bit for the Delay MUST be set as described in [RFC 7471](#) [[RFC7471](#)]. Similarly, the A-bit for the Delay Variation MUST be set as described in [RFC 7471](#) [[RFC7471](#)]. It then forwards the Path message to the next node in the downstream direction. The processing node MUST retain a record of the Cost and/ or Delay and/ or Delay Variation Collection request for reference during Resv processing described below.

If the addition of Cost and/or Delay and/or Delay Variation 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 information MUST NOT be added. If the Cost and/or Delay and/or Delay Variation 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 Cost
Ali, Swallow, Filsfils Expires Jan 2020 [Page 10]

and/or Delay and/or Delay Variation collection request was contained in an LSP_ATTRIBUTES Object, the processing node MAY omit some or all of the corresponding information from the RRO; otherwise it MUST behave as specified by [RFC 3209](#) [[RFC3209](#)] and drop the RRO from the Path message entirely.

Following the steps described above, the intermediate nodes of the LSP can collect the Cost and/or Delay and/or Delay Variation 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 the corresponding 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 Cost and/or Delay and/or Delay Variation Collection was specified, then when local policy allows recording of the requested information, the node SHOULD add corresponding information, to the RRO of the outgoing Resv message, as specified below. The A-bit for the Delay MUST be set as described in [RFC 7471](#) [[RFC7471](#)]. Similarly, the A-bit for the Delay Variation MUST be set as described in [RFC 7471](#) [[RFC7471](#)]. When the Resv message arrives at the ingress node, the ingress node can extract the requested information from the RRO in the same way as the egress node.

Note that a link's Cost and/ or Delay and/ or Delay Variation 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 Cost and/ or Delay and/ or Delay Variation 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 Cost and/ or Delay and/ or Delay Variation 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 metric information in the same order for both Path messages and Resv messages. That is, the Cost and/ or Delay and/ or Delay Variation sub-object(s) for the upstream link is added to the RRO before the corresponding sub-object for the downstream link.

If Cost and/ or Delay and/ or Delay Variation data is added
for both the upstream and downstream links, the two sets of
the data MUST be added in separate corresponding sub-
Ali, Swallow, Filsfils Expires Jan 2020 [Page 11]

object(s). A single Cost or Delay or Delay Variation sub-object MUST NOT contain a mixture of the applicable data for upstream and downstream directions. When adding a Cost or Delay or Delay Variation sub-object to an RRO, the D-bit MUST be set appropriately to indicate the direction of the TE Link. If the same value applies in both directions, it SHOULD be added to both the corresponding upstream and downstream sub-objects.

Based on the local policy, a transit node may edit a Path or Resv RRO to remove route information (e.g. node or interface identifier information) before forwarding it. A node that does this SHOULD summarize the cost, Delay and Delay Variation data. How a node that performs the RRO edit operation calculates the Cost and/or Delay and/or Delay variation metric is beyond the scope of this document.

A node SHOULD NOT add Cost or Delay or Delay Variation information without an explicit request for the corresponding information being made by the ingress node in the Path message.

4.2. Metric Update

When the Cost and/or Delay and/or Delay Variation information of a link is changed, the endpoints of LSPs using that link need to be aware of the changes. When a change to Cost or Delay or Delay Variation information associated with a link occurs, the procedures defined in [Section 4.4.3 of RFC 3209](#) [[RFC3209](#)] MUST be used to refresh the corresponding metric information if the change is to be communicated to other nodes according to the local node's policy. If local policy is that the Cost and/or Delay and/or Delay Variation change SHOULD be suppressed or would result in no change to the previously signaled information, the node SHOULD NOT send an update.

4.3. Domain Boundaries

If mandated by local policy, a node MAY remove Cost and/or Delay and/or Delay Variation information from any RRO in a Path or Resv message being processed. A node that does this SHOULD summarize the Cost, Delay and Delay Variation data. How a node that performs the RRO edit operation calculates the Cost, Delay and/or Delay variation metric is beyond the scope of this document.

4.4. Endpoint processing

Based on the procedures described above, the endpoints can get the Cost and/or Delay and/or Delay Variation 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 corresponding TE metric

Ali, Swallow, Filsfils Expires Jan 2020 [Page 12]

information of the Forwarding Adjacency (FA) or Routing Adjacency (RA) automatically. How the end point uses the collected information is outside the scope of this document.

The ingress and egress nodes of a LSP may calculate the end-to-end Cost, Delay and/or Delay variation properties of the LSP from the supplied values in the Resv or Path RRO, respectively.

Typically, Cost and Delay are additive metrics, but Delay variation is not an additive metric. The means by which the ingress and egress nodes compute the end-to-end Cost, Delay and Delay variation metric from information recorded in the RRO is a local decision and is beyond the scope of this document.

Based on the local policy, the ingress and egress nodes can advertise the calculated end-to-end Cost, Delay and/or Delay variation properties of the FA or RA LSP in TE link advertisement to the routing instance based on the procedure described in [RFC 7471](#) [[RFC7471](#)], [[RFC7810](#)].

4.5. Compatibility

A node that does not recognize the Cost and/or Delay and/or Delay Variation Collection Flag in the Attribute Flags TLV is expected to proceed as specified in [RFC 5420](#) [[RFC5420](#)]. Specifically, the node is expected to pass the TLV on unaltered if it appears in a LSP_ATTRIBUTES object. On the other hand, if the TLV appears in a LSP_REQUIRED_ATTRIBUTES object, the node is expected to reject the Path message with the Error Code and Value defined in [RFC 5420](#) [[RFC5420](#)].

A node that does not recognize the Cost and/or Delay and/or Delay Variation RRO sub-object is expected to behave as specified in [RFC 3209](#) [[RFC3209](#)]: unrecognized sub-objects are to be ignored and passed on unchanged.

5. Manageability Considerations

5.1. Policy Configuration

In a border node of inter-domain or inter-layer network, the following Cost and/or Delay and/or Delay Variation processing policy SHOULD be capable of being configured:

- o Whether the node is allowed to participate in Cost or Delay or Delay Variation collection.
- o Whether the node should notify changes to collected Cost and/ or Delay and/ or Delay Variation information to endpoint nodes as described in [section 4.2](#).

- o Whether the Cost and/or Delay and/or Delay Variation 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.

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

6. 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 Cost and/or Delay and/or Delay Variation data between layers or domains during the signaling of LSPs, subject to policy at the layer or domain boundary. It is recommended that domain/layer boundary policies take the implications of releasing Cost and/or Delay and/or Delay Variation information into consideration and behave accordingly during LSP signaling.

7. IANA Considerations

7.1. 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 the following three new Attribute Bit Flags:

Bit No	Name	Attribute Flags Path	Attribute Flags Resv	RRO	Reference
-----	-----	-----	-----	---	-----
TBA by IANA	Cost Collection Flag	Yes	No	Yes	This I-D
TBA by	Delay	Yes	No	Yes	This I-D

IANA	Collection		
	Flag		
Ali, Swallow, Filsfils		Expires Jan 2020	[Page 14]

TBA by IANA	Delay Variation Collection Flag	Yes	No	Yes	This I-D
-------------	---------------------------------	-----	----	-----	----------

7.2. ROUTE_RECORD sub-object

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

Type	Name	Reference
-----	-----	-----
TBA by IANA	Cost sub-object	This I-D
TBA by IANA	Delay sub-object	This I-D
TBA by IANA	Delay Variation sub-object	This I-D

7.3. 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 the following three new Policy Control Failure Error sub-code:

Value	Description	Reference
-----	-----	-----
TBA by IANA	Cost Recoding Rejected	This I-D
TBA by IANA	Delay Recoding Rejected	This I-D
TBA by IANA	Delay Variation Recoding Rejected	This I-D

8. References

8.1. 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.
- [RFC5420] Farrel, A., Ed., 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.
- [RFC7471] S. Giacalone, D. Ward, J. Drake, A. Atlas, S. Previdi., "OSPF Traffic Engineering (TE) Metric Extensions", [RFC 7471](#), March 2015.
- [RFC7810] S. Previdi, S. Giacalone, D. Ward, J. Drake, A. Atlas, C. Filsfils, "IS-IS Traffic Engineering (TE) Metric Extensions", [draft-ietf-isis-te-metric-extensions](#), work in progress.

8.2. 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.
- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", [RFC 2209](#), September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.
- [RFC8001] F. Zhang, O. Gonzalez de Dios, M. Hartley, Z. Ali, C. Margaria,, "RSVP-TE Extensions for Collecting SRLG Information", [draft-ietf-teas-rsvp-te-srlg-collect.txt](#), work in progress.

Acknowledgements

Authors would like to thank Ori Gerstel, Gabriele Maria Galimberti, Luyuan Fang and Walid Wakim for their review comments.

Ali, Swallow, Filsfils

Expires Jan 2020

[Page 16]

Contributors

Sajal Agarwal
Cisco Systems
Email: sajaagar@cisco.com

George Swallow
Individual

Matt Hartley
Individual

Authors' Addresses

Zafar Ali (editor)
Cisco Systems, Inc.
Email: zali@cisco.com

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

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

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

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

Ali, Swallow, Filsfils

Expires Jan 2020

[Page 17]