

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 13, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Cisco Systems

Kenji Kumaki
KDDI Corporation

Ruediger Kunze
Deutsche Telekom AG
July 14, 2013

**Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
extension for recording TE Metric of a Label Switched Path
draft-ietf-ccamp-te-metric-recording-02.txt**

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

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

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

There are many scenarios in which Traffic Engineering (TE) metrics such as cost, latency and latency variation associated with a Forwarding Adjacency (FA) or Routing Adjacency (RA) 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) for the support of the discovery of cost, latency and latency variation of an 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

Copyright Notice.....	1
1 . Introduction.....	3
2 . RSVP-TE Requirement.....	3
2.1 . Cost, Latency and Latency Variation Collection Indication....	4
2.2 . Cost, Latency and Latency Variation Collection.....	4
2.3 . Cost, Latency and Latency Variation Update.....	4
3 . RSVP-TE signaling extensions.....	4
3.1 . Cost, Latency and Latency Variation Collection Flags.....	4
3.2 . Cost Subobject.....	5
3.3 . Latency Subobject.....	6
3.4 . Latency Variation Subobject.....	7
3.5 . Signaling Procedures.....	8
4 . Security Considerations.....	9
5 . IANA Considerations.....	9
5.1 . RSVP Attribute Bit Flags.....	9

5.2.	New RSVP error sub-code.....	10
6.	Acknowledgments.....	11
7.	References.....	11
7.1.	Normative References.....	11
7.2.	Informative References.....	12

1. Introduction

There are many 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. In such scenarios, the ingress node cannot determine the cost, latency and latency variation properties of the LSP's route. Similarly, in Generalized Multi-Protocol Label Switching (GMPLS) networks signaling bidirectional LSP, the egress node cannot determine the cost, latency and latency variation properties of the LSP route. A multi-domain or multi-layer network is an example of such networks. Similarly, a GMPLS User-Network Interface (UNI) [[RFC4208](#)] is also an example of such networks.

In certain networks, such as financial information networks, network performance information (e.g. latency, latency variation) is becoming as critical to data path selection as other metrics [[DRAFT-OSPF-TE-METRIC](#)], [[DRAFT-ISIS-TE-METRIC](#)]. If cost, latency or latency variation associated with an FA or an RA LSP is not available to the ingress or egress node, it cannot be advertised as an attribute of the FA or RA. One possible way to address this issue is to configure cost, latency and latency 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 case where that an LSP is advertised as a TE-Link, the ingress and/or egress nodes cannot provide the correct latency, latency variation and cost attribute associated with the TE-Link automatically.

In summary, there is a requirement for the ingress and egress nodes to learn the cost, latency and latency variation attributes of an FA or RA LSP. This draft provides extensions to the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) for the support of the automatic discovery of these attributes.

2. RSVP-TE Requirement

This section outlines RSVP-TE requirements for the support of the automatic discovery of cost, latency and latency variation attributes of an LSP. These requirements are very similar to the requirement of discovering the Shared Risk Link Groups (SRLGs)

associated with the route taken by an LSP [DRAFT-SRLG-RECORDING].

Ali, Swallow, Filsfils

Expires January 2014

[Page 3]

2.1. Cost, Latency and Latency Variation Collection Indication

The ingress node of the LSP must be capable of indicating whether the cost, latency and latency variation attributes of the LSP should be collected during the signaling procedure of setting up the LSP. No cost, latency or latency variation information is collected without an explicit request being made by the ingress node.

2.2. Cost, Latency and Latency Variation Collection

If requested, cost, latency and latency variation is collected during the setup of an LSP. The endpoints of the LSP may use the collected information for routing, flooding and TE link configuration and other purposes.

2.3. Cost, Latency and Latency Variation Update

When the cost, latency and latency variation property of a TE link along the route of a LSP for which that property was collected changes, e.g., if the administrator changes cost of a TE link, the node where the change occurred needs to be capable of updating the cost, latency and latency variation information of the path and signaling this to the end-points. Similarly, if a path segment of the LSP is rerouted, the endpoints of the rerouted segment need to be capable of updating the cost, latency and latency variation information of the path. Any node, which adds cost, latency or latency variation information to an LSP during initial setup, needs to signal changes to these values to both endpoints.

3. RSVP-TE signaling extensions

3.1. Cost, Latency and Latency Variation Collection Flags

Three Attribute flags are defined in the Attribute Flags TLV, which can be set and carried in either the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Objects.

- Cost Collection flag (to be assigned by IANA)
- Latency Collection flag (to be assigned by IANA)
- Latency Variation Collection flag (to be assigned by IANA)

These flags are meaningful in a Path message. If the Cost Collection flag is set to 1, the transit nodes SHOULD report the cost information in the Record Route Objects (RRO) of both the Path and Resv messages.

Ali, Swallow, Filsfils

Expires January 2014

[Page 4]

If the Cost Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.

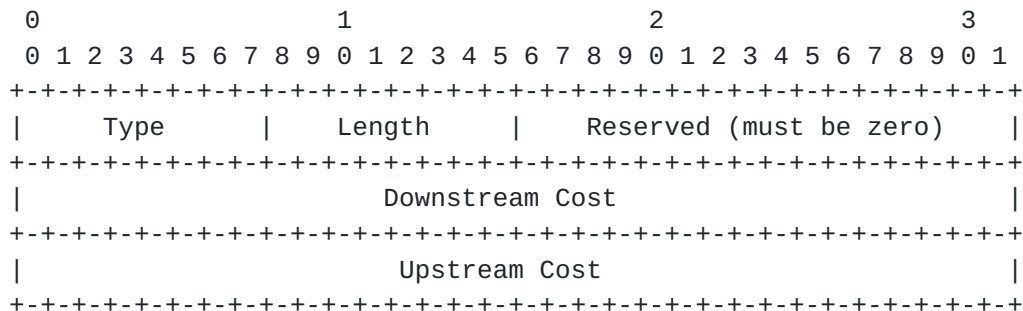
If the Latency Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.

If the Latency Variation Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.

The procedure for the processing the Attribute Flags TLV follows [RFC5420].

3.2. Cost Subobject

The cost subobject is defined for the RRO to record the cost information of the LSP. Its format is similar to the RRO subobjects (ROUTE_RECORD sub-object) defined in [\[RFC3209\]](#).



Type: TBA1 - Cost subobject (to be assigned by IANA).

Length: The Length value is set to 8 or 12 depending on the presence of Upstream Cost information.

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

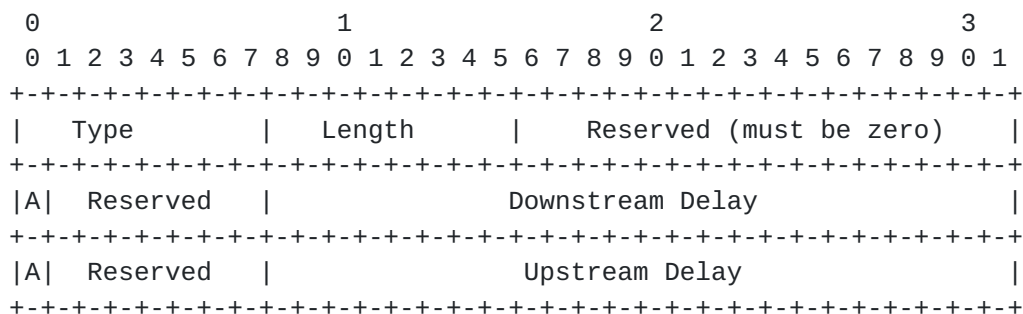
Downstream Cost: Cost of the local link along the route of the LSP in the direction of the tail-end node, encoded as a 32-bit integer. Based on the policy at the recording node, the cost value can be set to the Interior Gateway Protocol (IGP) metric or TE metric of the link in question. This approach has been taken to avoid defining a flag for each cost type in the Attribute-Flags TLV. It is assumed that, based on policy, all nodes report the same cost-type and that

the ingress and egress nodes know the cost type reported in the RRO.

Upstream Cost: Cost of the local link along the route of the LSP in the direction of the head-end node, encoded as a 32-bit integer. Based on the policy at the recording node, the cost value can be set to the Interior Gateway Protocol (IGP) metric or TE metric of the link in question. This approach has been taken to avoid defining a flag for each cost type in the Attribute-Flags TLV. It is assumed that, based on policy, all nodes report the same cost-type and that the ingress and egress nodes know the cost type reported in the RRO.

3.3. Latency Subobject

The Latency subobject is defined for RR0 to record the latency information of the LSP. Its format is similar the RR0 subobjects defined in [\[RFC3209\]](#).



Type: TBA2 - Latency subobject (to be assigned by IANA).

Length: 8 or 12 depending on the presence of Upstream Cost information.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay respectively, as defined in [[DRAFT-OSPF-TE-METRIC](#)].

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

Downstream Delay: Delay of the local link along the route of the LSP in the direction of the tail-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

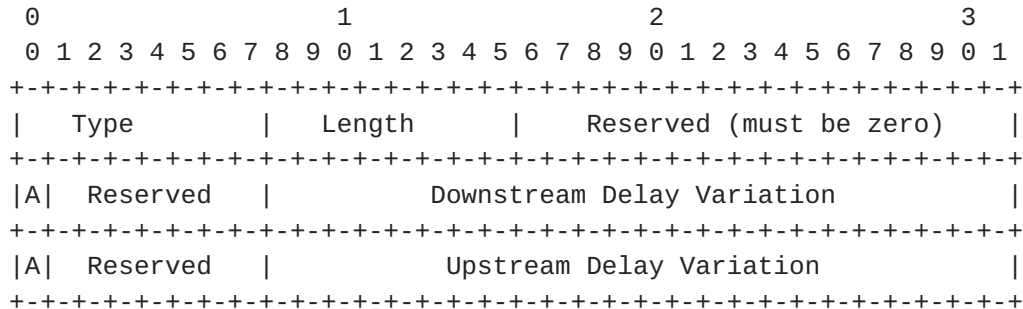
Upstream Delay: Delay of the local link along the route of

the LSP in the direction of the head-end node, encoded as 24-
Ali, Swallow, Filsfils Expires January 2014 [Page 6]

bit integer. When set to 0, it has not been measured. 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. Latency Variation Subobject

The Latency Variation subobject is defined for RRO to record the Latency Variation information of the LSP. Its format is similar to the RRO subobjects defined in [RFC3209].



Type: TBA3 - Latency Variation subobject (to be assigned by IANA).

Length: 8 or 12 depending on the presence of Upstream Latency Variation information.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay respectively, as defined in [DRAFT-OSPF-TE-METRIC].

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

Downstream Delay Variation: Delay Variation of the local link along the route of the LSP in the direction of the tail-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

Upstream Delay Variation: Delay Variation of the local link along the route of the LSP in the direction of the head-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

3.5. Signaling Procedures

Typically, the ingress node learns the route of an LSP by adding a RRO in the Path message. If an ingress node also desires cost, latency and/or latency variation recording, it sets the appropriate flag(s) in the Attribute Flags TLV of the LSP_ATTRIBUTES (if recording is desired but not mandatory) or LSP_REQUIRED_ATTRIBUTES (if recording is mandatory) Object. None, all or any of the Cost Collection, Latency Collection or Latency Variation Collection flags may be set in the Attribute Flags TLV of the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Object. The rules for processing the LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES Objects and RRO are not changed. The corresponding sub-objects MUST be included in the RRO, with the Downstream (only) information filled in.

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object and the Cost, Latency and/or Latency Variation Collection Flag(s) is (are) set, if local policy disallows providing the requested information to the endpoints, the node MUST return a Path Error message with error code "Policy Control Failure (2)" and one of the following error subcodes:

- . "Cost Recording Rejected" (value to be assigned by IANA, suggested value 105) if Cost Collection Flag is set.
- . "Latency Recording Rejected" (value to be assigned by IANA, suggested value 106) if Latency Collection Flag is set.
- . "Latency Variation Recording Rejected" (value to be assigned by IANA, suggested value 107) if Latency Variation Collection Flag is set.

When a node receives a Path message which carries an LSP_ATTRIBUTES Object and the Cost, Latency and/or Latency Variation Collection Flag(s) is (are) set, if local policy disallows providing the requested information to the endpoints, the Path message SHOULD NOT be rejected due to Metric recording restriction and the Path message is forwarded without the appropriate sub-object(s) in the Path RRO.

If local policy permits the recording of the requested information, the processing node SHOULD add the requested subobject(s) with the cost, latency and/or latency variation metric value(s) associated with the local hop to the Path RRO. If the LSP being setup is bidirectional, both Downstream and Upstream information SHOULD be included. If the LSP is unidirectional, only Downstream information SHOULD be included.

Following the steps described above, the intermediate nodes of
the LSP provide the requested metric value(s) associated with
Ali, Swallow, Filsfils Expires January 2014 [Page 8]

the local hop in the Path RRO. When the egress node receives the Path message, it can calculate the end-to-end cost, latency and/or latency variation properties of the LSP.

Before the Resv message is sent to the upstream node, the egress node adds the requested subobject(s) with the downstream cost, latency and/or latency variation metric value(s) associated with the local hop to the Resv RRO in a similar manner to that specified above for the addition of Path RRO sub-objects by transit nodes.

Similarly, the intermediate nodes of the LSP provide the requested metric value(s) associated with the local hop in the Resv RRO. When the ingress node receives the Resv message, it can calculate the end-to-end cost, latency and/or latency variation properties of the LSP.

Typically, cost and latency are additive metrics, but latency variation is not an additive metric. The means by which the ingress and egress nodes compute the end-to-end cost, latency and latency variation metric from information recorded in the RRO 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, latency and/or latency variation properties of the FA or RA LSP in TE link advertisement to the routing instance based on the procedure described in [[DRAFT-OSPF-TE-METRIC](#)], [[DRAFT-ISIS-TE-METRIC](#)].

Based on the local policy, a transit node (e.g. the edge node of a domain) 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, latency and latency variation data removed as a single value for each for the loose hop that is summarized by the transit node. How a transit node calculates the cost, latency and/or latency variation metric for the segment summarized by the transit node is beyond the scope of this document.

4. Security Considerations

This document does not introduce any additional security issues above those identified in [[RFC5920](#)], [[RFC5420](#)], [[RFC2205](#)], [[RFC3209](#)], and [[RFC3473](#)].

5. IANA Considerations

5.1. RSVP Attribute Bit Flags

The IANA has created a registry and manages the space of

attributes bit flags of Attribute Flags TLV as described in
[section 11.3 of \[RFC5420\]](#). It is requested that the IANA makes
Ali, Swallow, Filsfils Expires January 2014 [Page 9]

assignments from the Attribute Bit Flags defined in this document.

This document introduces the following three new Attribute Bit Flag:

- Bit number: TBD (recommended bit position 11)
- Defining RFC: this I-D
- Name of bit: Cost Collection Flag

- Bit number: TBD (recommended bit position 12)
- Defining RFC: this I-D
- Name of bit: Latency Collection Flag

- Bit number: TBD (recommended bit position 13)
- Defining RFC: this I-D
- Name of bit: Latency Variation Flag

5.2. ROUTE_RECORD subobject

This document introduces the following three new RRO subobject:

Type	Name	Reference
-----	-----	-----
TBD (35)	Cost subobject	This I-D
TBD (36)	Latency subobject	This I-D
TBD (37)	Latency Variation subobject	This I-D

5.2. New RSVP error sub-code

For Error Code = 2 "Policy Control Failure" (see [[RFC2205](https://datatracker.ietf.org/doc/rfc2205)]) the following sub-code is defined.

Sub-code	Value
----------	-------

Ali, Swallow, Filsfils

Expires January 2014

[Page 10]

Cost Recoding Rejected	To be assigned by IANA. Suggested Value: 105.
Latency Recoding Rejected	To be assigned by IANA. Suggested Value: 106.
Latency Variation Recoding Rejected IANA.	To be assigned by Suggested Value: 107.

6. Acknowledgments

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

7. References

7.1. 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.
- [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.
- [DRAFT-OSPF-TE-METRIC] S. Giacalone, D. Ward, J. Drake, A. Atlas, S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", [draft-ietf-ospf-te-metric-extensions](#), work in progress.
- [DRAFT-ISIS-TE-METRIC] 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.
- [DRAFT-SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria,, "RSVP-TE Extensions for Collecting SRLG Information", [draft-ietf-ccamp-rsvp-te-srlg-collect.txt](#), work in progress.

7.2. 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.

Authors' Addresses

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

George Swallow
Cisco Systems, Inc.
swallow@cisco.com

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

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

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

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

