

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 15, 2013

Zafar Ali
George Swallow
Clarence Filsfils
Cisco Systems
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG

July 16, 2012

**Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
extension for recording TE Metric of a Label Switched Path
draft-ali-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 15, 2013.

Copyright Notice

Copyright (c) 2012 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.

Ali, Swallow, Filsfils, et al Expires January 2013 [Page 1]

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, 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 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 Collection Flag.....	4
3.2 . Latency Collection Flag.....	4
3.3 . Latency Variation Collection Flag.....	5
3.4 . Cost subobject.....	5
3.5 . Latency subobject.....	6
3.6 . Latency Variation subobject.....	7
3.7 . Signaling Procedures.....	7

[4. Security Considerations.....9](#)
[5. IANA Considerations.....9](#)
[5.1. RSVP Attribute Bit Flags.....9](#)
[5.2. New RSVP error sub-code.....10](#)
Ali, Swallow, Filsfils Expires January 2013 [Page 2]

[6.](#) Acknowledgments.....[10](#)
[7.](#) References.....[10](#)
[7.1.](#) Normative References.....[10](#)
[7.2.](#) Informative References.....[11](#)

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 get the cost, latency and latency variation properties of the LSP's route. Similarly, in Generalized Multi-Protocol Label Switching (GMPLS) networks signaling bidirectional Label Switched Path (LSP), the egress node cannot get 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 2013

[Page 3]

2.1. Cost, Latency and Latency Variation Collection Indication

The ingress and egress nodes 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.

2.2. Cost, Latency and Latency Variation Collection

The endpoints of the LSP may collect the cost, latency and latency variation information and use it for routing, flooding, and TE link configuration purposes.

2.3. Cost, Latency and Latency Variation Update

When the cost, latency and latency variation property of a TE link along the LSP route changes, e.g., if the administrator changes cost of a TE link, the endpoints of the LSP need to be capable of updating the cost, latency and latency variation information of the path. Similarly, if a path segment of the LSP is rerouted, the endpoints of the LSP need to be capable of updating the cost, latency and latency variation information of the path. In summary, the signaling should be capable of updating the new cost, latency and latency variation information to the endpoints.

3. RSVP-TE signaling extensions

3.1. Cost Collection Flag

In order to indicate that cost collection is desired, a new flag in the Attribute Flags TLV which can be carried in an LSP_REQUIRED_ATTRIBUTES Object is required:

Cost Collection flag (to be assigned by IANA, recommended bit position 9)

The Cost Collection flag is meaningful in a Path message. If the Cost Collection flag is set to 1, the transit nodes SHOULD report the cost information to the ingress and egress nodes in the Path Record Route Object (RRO) and the Resv RRO.

The rules of the processing of the Attribute Flags TLV follows [[RFC5420](#)].

3.2. Latency Collection Flag

In order to indicate that latency collection is desired, a new flag in the Attribute Flags TLV which can be carried in an LSP_REQUIRED_ATTRIBUTES Object is required:

Latency Collection flag (to be assigned by IANA, recommended bit position 10)

The Latency Collection flag is meaningful on a Path message. If the Latency Collection flag is set to 1, the transit nodes SHOULD report the latency information to the ingress and egress nodes in the Path RRO and the Resv RRO.

The rules of the processing of the Attribute Flags TLV follows [\[RFC5420\]](#).

3.3. Latency Variation Collection Flag

In order to indicate that latency variation collection is desired, a new flag in the Attribute Flags TLV which can be carried in an LSP_REQUIRED_ATTRIBUTES Object is required:


Latency Variation Collection flag (to be assigned by IANA, recommended bit position 11)

The Latency Variation Collection flag is meaningful on a Path message. If the Latency Variation Collection flag is set to 1, the transit nodes SHOULD report the latency variation information to the ingress and egress nodes in the Path RRO and the Resv RRO.

The rules of the processing of the Attribute Flags TLV follows [\[RFC5420\]](#).

3.4. Cost subobject

A new cost subobject is defined for the RRO to record the cost information of the LSP. Its format is similar to the RRO subobjects defined in [\[RFC3209\]](#).



Type: The type of the subobject, to be assigned by IANA (recommended value 35).

Length: The Length value is set to 8.

Reserved: This field is reserved for future use. It MUST be

set to 0 when sent and MUST be ignored when received.

Ali, Swallow, Filsfils

Expires January 2013


[Page 5]

Cost Value: Cost of the link along the route of the LSP. 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 LSP_REQUIRED_ATTRIBUTES subobject. It is assumed that, based on policy, all nodes reports the same cost-type and that the ingress and egress nodes know the cost type reported in the RRO.

The rules of the processing of the LSP_REQUIRED_ATTRIBUTES Object and RRO are not changed.

3.5. Latency subobject

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



Type: The type of the subobject, to be assigned by IANA (recommended value 36).

Length: The Length value is set to 8.

A-bit: This field represents the Anomalous (A) bit, 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.

Delay Value: This 24-bit field carries the average link delay over a configurable interval in micro-seconds, encoded as an integer value. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value and may be larger.

The rules of the processing of the LSP_REQUIRED_ATTRIBUTES Object and RRO are not changed.

3.6. Latency Variation subobject

A new Latency Variation subobject is defined for RRO to record the Latency information of the LSP. Its format is similar to the RRO subobjects defined in [[RFC3209](#)].



Type: The type of the subobject, to be assigned by IANA (recommended value 37).

Length: The Length value is set to 8.

A-bit: This field represents the Anomalous (A) bit, 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.

Delay Variation Value: This 24-bit field carries the average link delay variation over a configurable interval in microseconds, encoded as an integer value. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value and may be larger.

The rules of the processing of the LSP_REQUIRED_ATTRIBUTES Object and RRO are not changed.

3.7. 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 or latency variation recording, it sets the Cost Collection flag, Latency Collection flag or Latency Variation Collection flag in the Attribute Flags TLV of LSP_REQUIRED_ATTRIBUTES Object, respectively. None, all or any of the Cost Collection, Latency Collection or Latency Variation Collection flags may be set in the Attribute Flags TLV of LSP_REQUIRED_ATTRIBUTES Object.

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object and the Cost, Latency or/ and

Latency Variation Collection Flag(s) is (are) set, if local
policy disallows providing the requested information to the
Ali, Swallow, Filsfils Expires January 2013 [Page 7]

endpoints, the node SHOULD return a Path Error message with error code "Policy Control Failure (2)" and one of the following error subcodes:

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

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object and the Cost, Latency or/ and Latency Variation Collection Flag(s) is (are) set, if local policy allows providing the requested information to the endpoints, the node MUST add the requested subobject(s) with the cost, latency or/ and latency variation metric value(s) associated with the local hop to the Path RRO. Then it forwards the Path message to the next node in the downstream direction.

Following the steps described above, the intermediate nodes of the LSP provide the requested metric value(s) associated with the local hop in the Path RRO. When the Path message is received by the egress node, the egress node can calculate end-to-end the cost, latency or/ and latency variation properties of the LSP.

Before the Resv message is sent to the upstream node, the egress node MUST add the requested subobject(s) with the cost, latency or/ and latency variation metric value(s) associated with the local hop to the Resv RRO. Similarly, the intermediate nodes of the LSP provide the requested metric value(s) associated with the local hop in the Resv RRO. When the Resv message is received by the Ingress node, the Ingress node can calculate end-to-end the cost, latency or/ and latency variation properties of the LSP.

Typically, cost and latency are additive metrics, but latency variation is not an additive metric. How the ingress and egress nodes computes the end-to-end cost, latency or/ 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 end-to-end the cost, latency or/ and latency variation properties of the FA/ RA LSP in TE link advertisement to the routing instance based on the procedure described in

[\[DRAFT-OSPF-TE-METRIC\]](#), [\[DRAFT-ISIS-TE-METRIC\]](#).

Ali, Swallow, Filsfils

Expires January 2013

[Page 8]

Based on the local policy, a transit node (e.g. the edge node of a domain) may edit the RRO to remove the route information (e.g. node, interface identifier information) before forwarding it and can summarize the cost, latency or/ and latency variation as a single number for the loose hop that is summarized by the edge node. How a transit node calculates the cost, latency or/ and 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 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 9)
- Defining RFC: this I-D
- Name of bit: Cost Collection Flag

- Bit number: TBD (recommended bit position 9)
- Defining RFC: this I-D
- Name of bit: Latency Collection Flag

- Bit number: TBD (recommended bit position 9)
- 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](#)]) the following sub-code is defined.

Sub-code	Value
-----	-----
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	To be assigned by IANA. Suggested Value: 107.

6. Acknowledgments

Authors would like to thanks Matt Hartley, 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.

Internet-Draft [draft-ali-ccamp-te-metric-recording-02.txt](#)

[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-previdi-isis-te-metric-extensions](#), work in progress.

[DRAFT-SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria, C. "RSVP-TE Extensions for Configuration SRLG of an FA", [draft-zhang-ccamp-srlg-fa-configuration.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
Ali, Swallow, Filsfils

Expires January 2013

[Page 11]

Internet-Draft [draft-ali-ccamp-te-metric-recording-02.txt](#)

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

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

