

Network Working Group
Internet Draft
Intended status: Experimental
Expires: October 26, 2011

A. Zamfir
Z. Ali
Cisco Systems
D. Papadimitriou
Alcatel-Lucent
April 27, 2011

Component Link Recording and Resource Control for TE Links

draft-ietf-mppls-explicit-resource-control-bundle-10.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 26, 2011.

Abstract

Record Route is a useful administrative tool that has been used extensively by the service providers. However, when TE links are bundled, identification of label resource in Record Route object (RRO) is not sufficient to determine the component link within a TE link that is being used by a given LSP. In other words, when link bundling is used, resource recording requires mechanisms to specify the component link identifier, along with the TE link identifier and Label. As it is not possible to record component link in the RRO, this document defines the extensions to RSVP-TE [RFC3209] and [RFC3473] to specify component link identifiers for resource recording purposes.

A.Zamfir et al. Expires October 26, 2011 [Page 1]
Component Link Recording & Resource Control for TE Links April 27, 2011

This document also defines the Explicit Route object (ERO) counterpart of the RRO extension. The ERO extensions are needed to perform explicit label/resource control over bundled TE link. Hence, this document defines the extensions to RSVP-TE [RFC3209] and [RFC3473] to specify component link identifiers for explicit

resource control and recording over TE link bundles.

Table of Contents

1. Introduction.....	3
Conventions used in this document.....	4
2. Terminology.....	4
3. LSP Resource Recording.....	4
3.1. Component Interface Identifier RRO Subobject.....	5
3.2. Processing of Component Interface ID RRO Subobject.....	6
4. Signaling Component Interface Identifier in ERO.....	7
4.1. Component Interface ID ERO Subobject.....	7
4.2. Processing of Component Interface ID ERO Subobject.....	8
5. Backward Compatibility.....	10
6. Security Considerations.....	10
7. IANA Considerations.....	10
8. References.....	11
8.1. Normative References.....	11
9. Acknowledgments.....	12
10. Copyright.....	13

1. Introduction

In GMPLS networks [RFC3945] where unbundled (being either Packet-Switching Capable, Layer2-Switching Capable, Time Division Multiplexing or Lambda-Switching Capable) Traffic Engineering (TE) Links are used, one of the types of resources that an LSP originator could control and record are the component links used by non-neighboring nodes on the LSP path. The resource control and recording is done by the use of the EXPLICIT_ROUTE object (ERO) and RECORD_ROUTE object (RRO), respectively.

Link Bundling, introduced in [RFC4201], is used to improve routing scalability by reducing the amount of TE related information that need to be flooded and handled by IGP in a TE network. This is accomplished by aggregating and abstracting the TE Link components.

In some cases the component link selection/recording within a TE link is left as a local decision (ERO and RRO contains only TE links). However there are cases when it is desirable for a non-local (e.g., LSP head-end) node to make this selection. The use of such information has found since so far three main applications (while not excluding others unknown at the time of writing): diagnostic, association of component specific attributes for which the bundled information is too coarse (e.g., Shared Risk Link Groups) and thus blocking SRLG-disjoint LSP establishment, allocation of labels at network edges, and notification in case of failures. The latter is useful when a single TE link interconnects two parts of the network. In case one of its components fails notifying a complete TE link failure leaves the network disconnected. In either case, it is required to know which component link within a bundled TE link has been used for a given LSP. For these cases, the TE Link and the Label currently specified in the ERO/RRO are not enough and the component link needs to be specified along with the label. In the case of bi-directional Label Switched Paths (LSP) both upstream and downstream information may be specified. Therefore, explicit resource control and recording over a bundled TE link also requires ability to specify a component link within the TE link.

Another important assumption of this document is that the identifier space used for component link identification are unique for a given node (following [RFC4201]). The reason stems as follows: most experimental developments started with TE links composed by a single component link and then only bundling was added by grouping them. Component links were thus identified such that they could mimic the behavior of TE link processing. This also justifies the experimental status of this document.

This document defines extensions to and describes the use of RSVP-TE

A.Zamfir et al. Expires October 26, 2011 [Page 3]
Component Link Recording & Resource Control for TE Links April 27, 2011

[RFC3209], [RFC3471], [RFC3473] to specify the component link identifier for resource recording and explicit resource control over TE link bundles. Specifically, in this document, component interface identifier RRO and ERO subobjects are defined to complement their Label RRO and ERO counterparts. Furthermore, procedures for processing component interface identifier RRO and ERO subobjects and how they can co-exist with the Label RRO and ERO subobjects are specified.

Conventions used in this document

In examples, "C:" and "S:" indicate lines sent by the client and server respectively.

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].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC 2119 significance.

In this document, the characters ">>" preceding an indented line(s)

indicates a compliance requirement statement using the key words listed above. This convention aids reviewers in quickly identifying or finding the explicit compliance requirements of this RFC.

2. Terminology

o) TE Link: Unless specified otherwise, it refers to a bundled Traffic Engineering link as defined in [RFC4201]. Furthermore, the terms TE Link and bundled TE Link are used interchangeably in this document.

o) Component (interface) link: refers (locally) to a link that is part of a bundled TE link as described in RFC4201.

o) Component Interface Identifier: Refers to an ID used to uniquely identify a Component Interface. on a bundled link a combination of <component link identifier, label> is sufficient to unambiguously identify the appropriate resources used by an LSP. The IDs used for component link identification are unique for a given node [RFC4201].

3. LSP Resource Recording

LSP Resource Recording refers to the ability to record the resources used by an LSP.

A.Zamfir et al. Expires October 26, 2011 [Page 4]
Component Link Recording & Resource Control for TE Links April 27, 2011

The procedure for unbundled numbered TE links is described in [RFC3209] and for unbundled unnumbered TE links in [RFC3477]. For the purpose of recording LSP resources used over bundled TE Links, the Component Interface Identifier RRO sub-object is introduced.

3.1. Component Interface Identifier RRO subobject

A new subobject of the Record Route object (RRO) is used to record component interface identifier of a (bundled) TE Link. This subobject has the following format:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|L|      Type      |      Length      |U| Reserved (must be zero) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|
//   IPv4, IPv6 or unnumbered Component Interface Identifier   //
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

L: 1 bit

This bit must be set to 0.

Type

Type 10 (TBD): Component Interface identifier IPv4
Type 11 (TBD): Component Interface identifier IPv6
Type 12 (TBD): Component Interface identifier Unnumbered

Length

Component Link Record. & Resource Control for TE Link Bundles

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is 8 bytes for the Component Interface identifier IPv4 and Component Interface identifier Unnumbered types. For Component Interface identifier IPv6 type of sub-object, the length field is 20 bytes.

U: 1 bit

This bit indicates the direction of the component interface. It is set to 0 for the downstream interface. It is set to 1 for the upstream interface and is only used for bi-directional LSPs.

A.Zamfir et al. Expires October 26, 2011 [Page 5]
Component Link Recording & Resource Control for TE Links April 27, 2011

3.2. Processing of Component Interface identifier RRO Subobject

If a node desires component link recording, the "Component Link Recording desired" flag (value TBD) should be set in the LSP_ATTRIBUTES object, object that is defined in [RFC5420].

Setting of "Component Link Recording desired" flag is independent of the Label Recording flag in SESSION_ATTRIBUTE object as specified in [RFC3209]. Nevertheless, the following combinations are valid:

- 1) If both Label and Component Link flags are clear, then neither Labels nor Component Links are recorded.
- 2) If Label Recording flag is set and Component Link flag is clear, then only Label Recording is performed as defined in [RFC3209].
- 3) If Label Recording flag is clear and Component Link flag is set, then Component Link Recording is performed as defined in this document.
- 4) If both Label Recording and Component Link flags are set, then Label Recording is performed as defined in [RFC3209] and also Component Link recording is performed as defined in this document.

In most cases, a node initiates recording for a given LSP by adding the RRO to the Path message. If the node desires Component Link recording and if the outgoing TE link is bundled, then the initial RRO contains the Component Link identifier (numbered or unnumbered) as selected by the sender. As well, the Component Link Recording desired flag is set in the LSP_ATTRIBUTE object. If the node also desires label recording, it sets the Label_Recording flag in the SESSION_ATTRIBUTE object.

When a Path message with the "Component Link Recording desired" flag set is received by an intermediate node, if a new Path message is to be sent for a downstream bundled TE link, the node adds a new Component Link subobject to the RECORD_ROUTE object (RRO) and

appends the resulting RRO to the Path message before transmission. Note also that, unlike Labels, Component Link identifiers are always known on receipt of the Path message.

When the destination node of an RSVP session receives a Path message with an RRO and the "Component Link Recording desired" flag set, this indicates that the sender node needs TE route as well as component link recording. The destination node initiates the RRO

process by adding an RRO to Resv messages. The processing mirrors that of the Path messages. The Component Interface Record subobject is pushed onto the RECORD_ROUTE object (RRO) prior to pushing on the node's IP address. A node MUST NOT push on a Component Interface Record subobject without also pushing on the IP address or unnumbered Interface Id subobject that identifies the TE Link.

When component interfaces are recorded for unidirectional LSPs, the downstream interface is the one identified by the Component Interface subobject. For bi-directional LSPs, component interface RRO subobjects for both downstream and upstream interfaces MUST be included.

4. Signaling Component Interface Identifier in ERO

4.1. Component Interface Identifier ERO subobject

A new OPTIoNAL subobject of the EXPLICIT_ROUTE object (ERO) is used to specify component interface identifier of a bundled TE Link. This Component Interface Identifier subobject has the following format:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|   Type   |   Length   |U|   Reserved (MUST be zero)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                     |
|//  IPv4, IPv6 or unnumbered Component Interface Identifier  //
|                                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

L: 1 bit

This bit must be set to 0.

Type

- Type 10 (TBD): Component Interface identifier IPv4
- Type 11 (TBD): Component Interface identifier IPv6
- Type 12 (TBD): Component Interface identifier Unnumbered

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is 8 bytes for the Component Interface identifier types: IPv4

and Component Interface identifier Unnumbered. For Component Interface identifier IPv6 type of sub-object, the length field

A.Zamfir et al. Expires October 26, 2011 [Page 7]
Component Link Recording & Resource Control for TE Links April 27, 2011

is 20 bytes.

U: 1 bit

This bit indicates the direction of the component interface. It is 0 for the downstream interface. It is set to 1 for the upstream interface and is only used for bi-directional LSPs.

4.2. Processing of Component Interface Identifier ERO Subobject

The Component Interface Identifier ERO subobject follows a subobject containing the IP address, or the link identifier [RFC3477], associated with the TE link on which it is to be used. It is used to identify the component of a bundled TE Link.

The following SHOULD result in "Bad EXPLICIT_ROUTE object" error being sent upstream by a node processing an ERO that contains the Component Interface ID sub-object:

- o) The first component interface identifier subobject is not preceded by a sub-object containing an IPv4 or IPv6 address, or an interface identifier [RFC3477], associated with a TE link.
- o) The Component Interface Identifier ERO subobject follows a subobject that has the L-bit set.
- o) on unidirectional LSP setup, there is a Component Interface Identifier ERO subobject with the U-bit set.
- o) Two Component Interface Identifier ERO subobjects with the same U-bit values exist.

If a node implements the component interface identifier subobject, it MUST check if it represents a component interface in the bundled TE Link specified in the preceding subobject that contains the IPv4/IPv6 address or interface identifier of the TE Link. If the content of the component interface identifier subobject does not match a component interface in the TE link, a "Bad EXPLICIT_ROUTE object" error SHOULD be reported as "Routing Problem" (error code 24).

If U-bit of the subobject being examined is cleared (0) and the upstream interface specified in this subobject is acceptable, then the value of the upstream component interface is translated locally in the TLV of the IF_ID RSVP_HOP object [RFC3471]. The local decision normally used to select the upstream component link is bypassed except for local translation into the outgoing interface identifier from the received incoming remote interface identifier.

A.Zamfir et al. Expires October 26, 2011 [Page 8]
Component Link Recording & Resource Control for TE Links April 27, 2011

If this interface is not acceptable, a "Bad EXPLICIT_ROUTE object"

error SHOULD be reported as "Routing Problem" (error code 24).

If the U-bit of the subobject being examined is set (1), then the value represents the component interface to be used for upstream traffic associated with the bidirectional LSP. Again, if this interface is not acceptable or if the request is not one for a bidirectional LSP, then a "Bad EXPLICIT_ROUTE object" error SHOULD be reported as "Routing Problem" (error code 24). otherwise, the component interface IP address/ identifier is copied into a TLV sub-object as part of the IF_ID RSVP_HOP object. The local decision normally used to select the upstream component link is bypassed except for local translation into the outgoing interface identifier from the received incoming remote interface identifier.

The IF_ID RSVP_HOP object constructed as above MUST be included in the corresponding outgoing Path message.

Note that, associated with a TE Link sub-object in the ERO, either the (remote) upstream component interface or the (remote) downstream component interface or both may be specified. As specified in [RFC4201] there is no relationship between the TE Link type (numbered or unnumbered) and the Link type of any one of its components.

The Component Interface Identifier ERO subobject is optional. Similarly, presence of the Label ERO sub-objects is not mandatory [RFC3471], [RFC3473]. Furthermore, component interface identifier ERO subobject and Label ERO subobject may be included in the ERO independently of each other. one of the following alternatives applies:

- o) When both sub-objects are absent, a node may select any appropriate component link within the TE link and any label on the selected component link.
- o) When the Label subobject is only present for a bundled link, then the selection of the component link within the bundle is a local decision and the node may select any appropriate component link, which can assume the label specified in the Label ERO.
- o) When only the component interface identifier ERO subobject is present, a node MUST select the component interface specified in the ERO and may select any appropriate label value at the specified component link.
- o) When both component interface identifier ERO subobject and Label ERO subobject are present, the node MUST select the locally

corresponding component link and the specified label value on that component link. When present, both subobjects may appear in any relative order to each other but they MUST appear after the TE Link subobject that they refer to.

After processing, the component interface identifier subobjects are removed from the ERO.

Inferred from above, the interface subobject should never be the

first subobject in a newly received message. If the component interface subobject is the first subobject in a received ERO, then it SHOULD be treated as a "Bad strict node" error.

Note: Information to construct the Component Interface ERO subobject MAY come from the same mean used to populate the label ERO subobject. Procedures by which an LSR at the head-end of an LSP obtains the information needed to construct the Component Interface subobject are outside the scope of this document.

5. Backward Compatibility

The extensions specified in this document do not affect the processing of the RRO, ERO at nodes that do not support them. A node that does not support the Component Interface RRO subobject but that does support Label subobject SHOULD only insert the Label subobject in the RRO as per [RFC3471] and [RFC3473].

A node that receives an ERO that contains a Component Link ID subobject SHOULD send "Bad EXPLICIT_ROUTE object" if it does not implement this subobject.

Per [RFC3209], Section 4.4.5, a non-compliant node that receives an RRO that contains Component Interface Identifier sub-objects should ignore and pass them on. This limits the full applicability of if nodes traversed by the LSP are compliant with the proposed extensions.

6. Security Considerations

An implementation of the extensions described in this document does exposes the component interface identifiers to other nodes in the network. If this is considered confidential information the mechanisms described in [RFC5920] should be considered.

7. IANA Considerations

This document introduces the following RSVP protocol elements:

A.Zamfir et al. Expires October 26, 2011 [Page 10]
Component Link Recording & Resource Control for TE Links April 27, 2011

o) Component Interface Identifier RRO subobject of the RECORD_ROUTE object (RRO):

- . IANA registry: RSVP PARAMETERS
- . Registry Name: Class Names, Class Numbers, and Class Types
- . Reference: [RFC3936]
- . Following subobjects have been added to the existing entry for:

21 RECORD_ROUTE

- . Type 10 (TBD): Component Interface identifier IPv4
- . Type 11 (TBD): Component Interface identifier IPv6
- . Type 12 (TBD): Component Interface identifier Unnumbered

o) Component Interface Identifier subobject of the EXPLICIT_ROUTE object (ERO).

- . IANA registry: RSVP PARAMETERS

- . Registry Name: Class Names, Class Numbers, and Class Types
- . Reference: [RFC3936]
- . Following subobjects have been added to the existing entry for:

20 EXPLICIT_ROUTE

- . Type 10 (TBD): Component Interface identifier IPv4
- . Type 11 (TBD): Component Interface identifier IPv6
- . Type 12 (TBD): Component Interface identifier Unnumbered

o) A new "Component Link Recording desired" flag (value TBD) of the LSP_ATTRIBUTES object [RFC5420]:

- . Bit Flag: 0x80
- . Name: Local Component Link Recording desired

8. References

8.1. Normative References

- [RFC2205] R.Braden, et al., "Resource ReSerVation Protocol (RSVP) Version 1, Functional Specification", RFC 2205, September 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

A.Zamfir et al. Expires October 26, 2011 [Page 11]
Component Link Recording & Resource Control for TE Links April 27, 2011

- [RFC2234] Crocker, D. and Overell, P. (Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.
- [RFC3209] D.Awduche, et al., "Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3471] L.Berger, et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] L.Berger, et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] K.Kompella, et al., "Signaling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC3945] E.Mannie, et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4201] K.Kompella, et al., "Link Bundling in MPLS Traffic Engineering", RFC 4201, January 2003.

[RFC5420] A.Farrel, et al., "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 5420.

[RFC5920] L.Fang, "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

9. Acknowledgments

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

A.Zamfir et al. Expires October 26, 2011 [Page 12]
Component Link Recording & Resource Control for TE Links April 27, 2011

10. Copyright Notice

Copyright (c) 2011 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.

Authors Addresses

Anca Zamfir
Cisco Systems Inc.
Email: ancaz@cisco.com

Zafar Ali
Cisco systems Inc.
Email: zali@cisco.com

Dimitri Papadimitriou
Alcatel-Lucent Bell
Email: dimitri.papadimitriou@alcatel-lucent.com

