

Internet Engineering Task Force (IETF)
Request for Comments: 5852
Category: Standards Track
ISSN: 2070-1721

D. Caviglia
D. Ceccarelli
D. Bramanti
Ericsson
D. Li
Huawei Technologies
S. Bardalai
Fujitsu Network
April 2010

RSVP-TE Signaling Extension for LSP Handover from the Management Plane
to the Control Plane in a GMPLS-Enabled Transport Network

Abstract

In a transport network scenario, Data Plane connections controlled by either a Generalized Multiprotocol Label Switching (GMPLS) Control Plane (Soft Permanent Connections - SPC) or a Management System (Permanent Connections - PC) may independently coexist. The ability of transforming an existing PC into an SPC and vice versa -- without actually affecting Data Plane traffic being carried over it -- is a requirement. The requirements for the conversion between permanent connections and switched connections in a GMPLS Network are defined in RFC 5493.

This memo describes an extension to GMPLS Resource Reservation Protocol - Traffic Engineering (RSVP-TE) signaling that enables the transfer of connection ownership between the Management and the Control Planes. Such a transfer is referred to as a Handover. This document defines all Handover-related procedures. This includes the handling of failure conditions and subsequent reversion to original state. A basic premise of the extension is that the Handover procedures must never impact an already established Data Plane connection.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5852>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	4
1.1. Dedication	4
2. Terminology	4
3. Motivation	4
4. Procedures	5
4.1. MP-to-CP Handover: LSP Ownership Transfer from Management Plane to Control Plane	6
4.2. MP-to-CP Handover Procedure Failure Handling	7
4.2.1. MP-to-CP Handover Failure - Path Failure	8
4.2.1.1. MP-to-CP Handover Failure - Path Message and Data Plane Failure	8
4.2.1.2. MP-to-CP Handover Failure - Path Message and Communication Failure	8
4.2.2. MP-to-CP Handover Failure - Resv Error	9
4.2.2.1. MP-to-CP Handover Failure - Resv Error and Data Plane Failure	9
4.2.2.2. MP-to-CP Handover Failure - Resv Error and Communication Failure	10
4.2.2.3. MP-to-CP Handover Failure - Node Graceful Restart	12
4.3. CP-to-MP Handover: LSP Ownership Transfer from Control Plane to Management Plane	15
4.4. CP-to-MP Handover Procedure Failure	16
5. Minimum Information for MP-to-CP Handover	17
6. RSVP Message Formats	19
7. Objects Modification	19
7.1. Administrative Status Object	19
7.2. Error Spec Object	19
8. Compatibility	20
9. Security Considerations	20
10. IANA Considerations	20
11. Acknowledgments	21
12. Contributors	21
13. References	21
13.1. Normative References	21
13.2. Informative References	22

1. Introduction

In a typical traditional transport network scenario, Data Plane (DP) connections between two endpoints are controlled by means of a Network Management System (NMS) operating within the Management Plane (MP). NMS/MP is the owner of such transport connections, being responsible for their setup, teardown, and maintenance.

The adoption of a Generalized MPLS (GMPLS) [RFC3945] Control Plane (CP) in a network that is already in service -- controlled by the NMS at the MP level -- introduces the need for a procedure able to coordinate a controlled Handover of a Data Plane connection from the MP to the CP.

In addition, the control Handover in the opposite direction, from CP to MP should be possible as well. The procedures described in this memo can be applied to a Label Switched Path (LSP) in any DP switching technology and any network architecture.

This memo describes an extension to GMPLS Resource reSerVation Protocol - Traffic Engineering (RSVP-TE) [RFC3471] [RFC3473] signaling that enables the Handover of connection ownership between the Management and the Control Planes. All Handover-related procedures are defined below. This includes the handling of failure conditions and subsequent reversion to original state. A basic premise of the extension is that the Handover procedures must never impact the exchange of user data on LSPs that are already established in the DP.

1.1. Dedication

We would like to dedicate this work to our friend and colleague Dino Bramanti, who passed away at the early age of 38. Dino has been involved in this work since its beginning.

2. Terminology

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 [RFC2119].

3. Motivation

The main motivation behind this work is the definition of a simple and very low-impact procedure that satisfies the requirements defined in [RFC5493]. Such a procedure is aimed at giving the transport network operators the chance to hand over the ownership of existing LSPs provisioned by NMS from the MP to the CP without disrupting user

traffic flowing on them. Handover from the MP to the CP (i.e., when existing DP connection ownership and control is passed from the MP to the CP) has been defined as a mandatory requirement, while the opposite operation, CP-to-MP Handover, has been considered as a nice-to-have feature that can be seen as an emergency procedure to disable the CP and take manual control of the LSP. For more details on requirements and motivations, please refer to [RFC5493].

4. Procedures

The modification defined in this document refers only to the ADMIN_STATUS Object, that is, the message flow is left unmodified for both LSP setup and deletion. Moreover, a new Error Value is defined to identify the failure of a Handover procedure.

The following paragraphs give detailed description of the "MP-to-CP Handover" and "CP-to-MP Handover" procedures, based on the use of a newly defined bit called "H bit".

Just as when setting up an LSP using the CP [RFC3473], the Path message may contain full information about the explicit route including the links and labels traversed by the LSP. This information is encoded in the Explicit Route Object (ERO), and must be supplied by the MP using details recorded when the LSP was provisioned, or collected by the MP by inspecting the nodes along the path.

Alternatively, and also just as when setting up an LSP using the CP [RFC3473], the ERO may include less information such that the details of the next hop have to be determined by each node along the LSP as it processes the Path message. This approach may be desirable when the full information is not available to the MP or cannot be passed to the head-end node when initiating the Handover from the MP to the CP.

This section (Section 4) describes the general procedures and protocol extensions for MP-to-CP Handover, and it uses the case of a fully detailed ERO to describe the mechanism. Section 5 describes how each node behaves in the case of a limited amount of information in the ERO.

Note that when Handover is being performed for a bidirectional LSP and the ERO contains full information including labels, the ERO SHOULD include both upstream and downstream labels. Per Section 5.1.1 of [RFC3473], the labels are indicated on an output basis; this means that the labels are used by the upstream node to create the LABEL_SET Object and, for bidirectional LSPs, the UPSTREAM_LABEL Object used in the outgoing Path message.

4.1. MP-to-CP Handover: LSP Ownership Transfer from Management Plane to Control Plane

The MP-to-CP Handover procedure MUST create an RSVP-TE session along the path of the LSP to be moved from the MP to the CP, associating it with the existing cross-connected resources owned by the MP (e.g., lambdas, time slots, or reserved bandwidth) and at the same time transferring their ownership to the CP.

The operator instructs the ingress node to hand over control of the LSP from the MP to the CP. In this Handover mode, it supplies the exact path of the LSP including any resource reservation and label information.

The ingress MUST check that no corresponding Path state exists and that corresponding Data Plane state does exist. If there is an error, this MUST be reported to the operator and further protocol action MUST NOT be taken.

The ingress signals the LSP using a Path message with the H bit and R bit set in the ADMIN_STATUS Object. In this mode of Handover, the Path message also carries an ERO that includes Label subobjects indicating the labels used by the LSP at each hop. The ingress MUST start the Expiration timer (see Section 4.2.1.2 for expiration of this timer). Such a timer SHOULD be configurable per LSP and have a default value of 30 seconds.

Each Label Switching Router (LSR) that receives a Path message with the H bit set checks to see whether there is any matching Path state.

- o If matching Path state is found with the H bit set, this is a Path refresh and should be treated accordingly [RFC3473].
- o If matching Path state is found with the H bit clear, this is an error and MUST be treated according to the error case description in Section 4.2.1.1.
- o If no Path state is found, the LSR goes on to check whether there is any matching Data Plane state.
- o If no matching Data Plane state is found (including only partially matching Data Plane state), this is an error or a race condition. It MUST be handled according to the description in Section 4.2.1.1.

- o If matching Data Plane state is found, the LSR MUST save the Path state (including the set H bit), and it MUST forward the Path message to the egress. The LSR MUST retain any MP state associated with the LSP at this point.

An egress LSR MUST act as any other LSR, except that there is no downstream node to which to forward the Path message. If all checks are passed, the egress MUST respond with a Resv with the H bit set.

A transit LSR MUST process each Resv according to the normal rules of [RFC3473].

When an ingress LSR receives a Resv message carrying the H bit set, it checks the Expiration timer.

- o If the timer is not running, the Resv is treated as a refresh and no special action is taken [RFC3473].
- o If the timer is running, the ingress MUST cancel the timer and SHOULD notify the operator that the first stage of Handover is complete. The ingress MUST send a Path message that is no different from the previous message except that the H bit MUST be clear.

The Path message with the H bit clear will travel the length of the LSP and will result in the return of a Resv with the H bit clear according to normal processing [RFC3473]. As a result, the H bit will be cleared in the stored Path state at each transit LSR and at the egress LSR. Each LSR SHOULD release any associated MP state associated with the LSP when it receives the Path message with H bit clear, but MAY retain the information according to local policy for use in future MP processing.

When the ingress receives a Resv with the H bit clear, the Handover is completed. The ingress SHOULD notify the operator that the Handover is correctly completed.

4.2. MP-to-CP Handover Procedure Failure Handling

In the case of MP-to-CP Handover, two different failure scenarios can happen: Path Failure and Resv Failure. Moreover, each failure can be due to two different causes: DP Failure or Communication Failure. In any case, the LSP ownership MUST be immediately rolled back to the one previous to the Handover procedure. A section for each combination will be analyzed in the following.

4.2.1. MP-to-CP Handover Failure - Path Failure

4.2.1.1. MP-to-CP Handover Failure - Path Message and Data Plane Failure

In the following paragraph, we will analyze the case where the Handover procedure fails during the Path message processing.




Figure 1: MP2CP - Path Msg and DP Failure

If an error occurs, the node detecting the error MUST respond to the received Path message with a PathErr message, and MUST abort the Handover procedure. The PathErr message SHOULD have the Path_State_Removed flag set [RFC3473], but implementations MAY retain their local state and wait for Path state timeout as per normal RSVP processing.

Nodes receiving a PathErr message MUST follow standard PathErr message processing and the associated DP resources MUST NOT be impacted. If the local CP state indicates that a Handover is in progress (based on the H bit in the Path message), the LSR MUST revert the LSP ownership to the MP.

4.2.1.2. MP-to-CP Handover Failure - Path Message and Communication Failure

Other possible scenarios are shown in the following figures and are based on the inability to reach a node along the path of the LSP.

The below scenario postulates the use of a reliable message delivery based on the mechanism defined in [RFC2961].




Figure 2: MP2CP - Path Msg and Communication Failure
(Reliable Delivery)

The Path message sent from LSR A towards LSR B is lost or does not reach the destination for any reason. As a reliable delivery mechanism is implemented, LSR A retransmits the Path message for a configurable number of times, and if no ack is received, the Handover procedure will be aborted (via the Expiration timer).

In the next scenario RSVP-TE messages are sent without reliable message delivery, that is, no [RFC2961] MessageID procedure is used.




Figure 3: MP2CP - Path Msg and Communication Failure
(No Reliable Delivery)

If the Resv message is not received before the expiration of the Expiration timer, the Handover procedure is aborted as described in Section 4.2.1.1. Please note that any node that has forwarded a Path (LSR A), i.e., has installed local path state, will send a PathTear when that state is removed (according to [RFC2205]).

4.2.2. MP-to-CP Handover Failure - Resv Error

4.2.2.1. MP-to-CP Handover Failure - Resv Error and Data Plane Failure

In the case of a failure occurrence during the Resv message processing (in case there has been any change in the Data Plane during the signaling), the node MUST send a PathErr message [RFC2205] in the upstream direction. The PathErr message is constructed and

processed as defined above in Section 4.2.1.1. The failure detection node MUST also send a PathTear message downstream. The PathTear message is constructed and processed as defined above in Section 4.2.1.1.




Figure 4: MP2CP - Resv Error and DP Failure

In the case shown in Figure 4, the failure occurs in LSR A. A PathTear message is sent towards B and a PathErr message (with ErrorCode set to "Handover Procedure Failure") is sent in the upstream direction. The PathErr and PathTear messages remove the Path state established by the Path messages along the nodes of the LSP and abort the Handover procedure.

Please note that the failure occurred after the Handover procedure was successfully completed in LSR B, but Handover state will still be maintained locally as, per Section 4.1, a Path message with the H bit clear will have not yet been sent or received. A node that receives a PathTear when it has Path state with the H bit set MUST remove Path state, but MUST NOT change Data Plane state. It MUST return LSP ownership to the MP.

4.2.2.2. MP-to-CP Handover Failure - Resv Error and Communication Failure

When a Resv message cannot reach one or more of the upstream nodes, the procedure is quite similar to the one previously seen about the Path message. Even in this case, it is possible to distinguish two different scenarios.

In the first scenario we consider the utilization of a reliable message delivery based on the mechanism defined in [RFC2961]. After a correct forwarding of the Path message along the nodes of the LSP, the Egress LSR sends a Resv message in the opposite direction. It might happen that the Resv message does not reach the ingress Label Edge Router (LER) or an LSR, say LSR A. LSR B MUST send a Resv message again for a configurable number of times and then, if the delivery does not succeed, the adoption procedure will be aborted (via the Expiration timer).

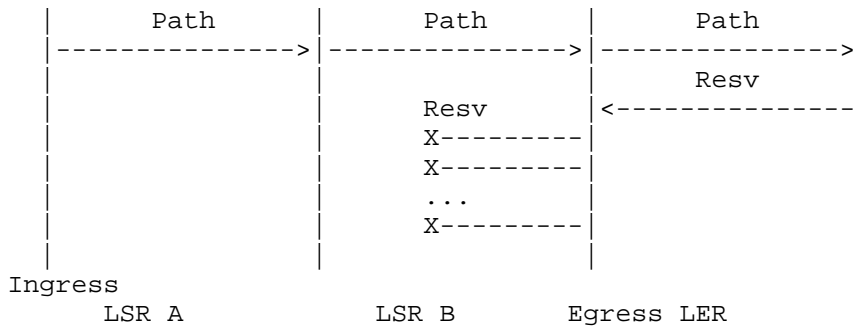


Figure 5: MP2CP - Resv Error and Communication Failure (Reliable Delivery)

Considering that the Resv message did not manage to reach LSR A, it is highly probable that the PathErr would fail too. Due to this fact, the Expiration timer is used on the ingress LER after sending the path and on LSR A after forwarding it. When the timer expires, if no Resv or PathErr message is received, the Handover procedure is aborted as described in Section 4.2.1.1 and the LSP ownership is returned to the Management Plane.

Figure 6, on the other hand, shows the scenario in which no reliable delivery mechanism is implemented.




Figure 6: MP2CP - Resv Error and Communication Failure
(No Reliable Delivery)

If no Resv message is received before the Expiration timer expires, the ingress LER follows the same procedure defined in Section 4.1.

4.2.2.3. MP-to-CP Handover Failure - Node Graceful Restart

If node restart and graceful restart are enabled, then one of the following scenarios will happen.

Case I - Finite Restart Time

In this case, the Restart Time (see [RFC3473]) is finite, i.e., not a value of 0xffffffff. In the sequence diagram below, assume LSR A restarts. If the ingress LER does not receive the Resv message in time, it MUST abort the Handover process by generating a PathTear message downstream. Also, if LSR A does not complete the restart process within the restart time interval, then LSR B MUST start tearing down all LSPs between LSR A and LSR B and this includes the LSP that is being used to carry out the Handover of MP resources to the CP. LSR B MUST generate a PathTear message downstream and a PathErr message upstream. Both LSR B and the egress LER MUST NOT release the DP resources because, in both nodes, the H bit is set in the local Path state.

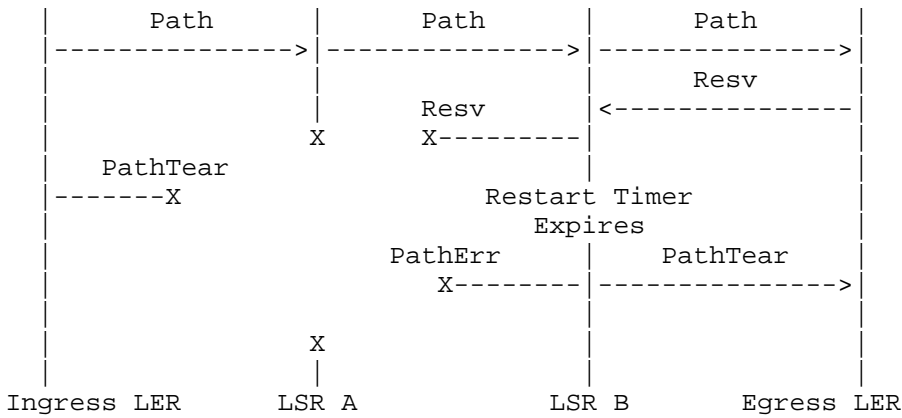


Figure 7: MP2CP - Node Graceful Restart - Case I

Case II - Infinite Restart Time

In this case, the Restart Time (see [RFC3473]) indicates that the restart of the sender's Control Plane may occur over an indeterminate interval, i.e., is 0xffffffff. The sequence is quite similar to the previous one. In this sequence, the restart timer will not expire in LSR B since it is run infinitely. Instead, after LSR A restarts, LSR B MUST start the recovery timer. The recovery timer will expire since there will be no Path message with the RECOVERY LABEL received from LSR A given the ingress node had already removed the local Path state after it aborts the Handover process. Thus, LSR B MUST tear down the specific LSP that is being used to convert the MP resources to CP by generating a PathTear message downstream and PathErr message upstream. Similarly to the previous case, both LSR B and the egress LER MUST NOT release the DP resources because the H bit is set in the local Path state.

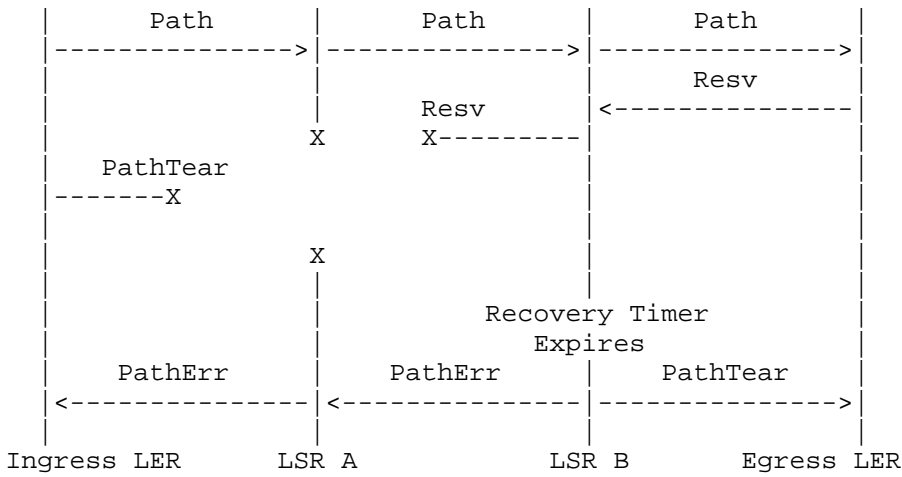


Figure 8: MP2CP - Node Graceful Restart - Case II

Case III

In this case, the ingress LER does not abort the Handover process. When LSR A restarts, the ingress LER detects the restart and MUST re-generate the Path message with the H bit set in order to restart the Handover.

When LSR B receives the Path message, it sees the H-bit set on the message and also sees that it has the H-bit set in its own state and that it has sent the Resv. But it is also aware that LSR A has restarted and could have sent a Path message with a RECOVERY LABEL object. LSR B may attempt to resume the Handover process or may abort the Handover. This choice is made according to local policy.

If resuming the Handover, LSR B MUST treat the received Path message as a retransmission, and MUST retransmit its Resv. If aborting Handover, LSR B MUST return a PathErr and MUST send a PathTear downstream. In both cases, LSR B MUST NOT modify the DP state.




Figure 9: MP2CP - Node Graceful Restart - Case III

4.3. CP-to-MP Handover: LSP Ownership Transfer from Control Plane to Management Plane

Let's now consider the case of LSP ownership transfer from Control Plane to Management Plane. Also in this section, we will analyze the Handover procedure success and failure.

The scenario is still a DP connection between two nodes acting as ingress and egress for a LSP, but in this case, the CP has the ownership and control of the LSP. The CP-to-MP Handover procedure MUST delete the existing RSVP-TE session information and MUST NOT affect the cross-connected resources, but just move their ownership to the MP.

In other words, after LSP ownership transfer from CP to MP, the LSP is no longer under the control of RSVP-TE, which is no more able to "see" the LSP itself. The CP-to-MP Handover procedure MUST be a standard LSP deletion procedure as described in Section 7.2.1 of [RFC3473]. The procedure is initiated at the ingress node of the LSP by an MP entity. The ingress node and MP exchange the relevant information for this task and then propagate it over CP by means of RSVP-TE tear down signaling as described below.

The ingress node MUST send a Path message in the downstream direction with Handover and Reflect bits set in the ADMIN_STATUS Object. No action is taken over the DP and transit LSRs must forward such message towards the egress node. All of the nodes MUST keep track of the procedure storing the H bit in their local Path and Resv states. Then, every node waits for the H bit to be received within the related Resv message. After the Resv message is received by the ingress LER, it MUST send a PathTear message in order to clear the

whole LSP information recorded on the RSVP-TE data structures of the nodes. Downstream nodes processing a PathTear message that follows a Path message with the H bit set, MUST NOT remove any associated Data Plane state. In other words, a normal LSP tear down signaling is exchanged between nodes traversed by the LSP, but the H bit set in the Path message indicates that no DP action has to correspond to CP signaling.

4.4. CP-to-MP Handover Procedure Failure

Failures during CP-to-MP Handover procedure MUST NOT result in the removal of any associated Data Plane state. To that end, when a Resv message containing an ADMIN_STATUS Object with the H bit not received during the period of time described in Section 7.2.2. of [RFC3473] different processing is required. While the H bit is set in the Path state, a node MUST NOT send a PathTear when a failure is detected. Instead, the failure is reported upstream using a PathErr. The only node that can send a PathTear is the ingress node, and it can only do this as a step in the procedures specified in this document. This applies to both MP-to-CP and CP-to-MP Handover. The ingress node MAY choose to report the failure in the CP-to-MP Handover procedure via the MP.

The CP-to-MP Handover procedure can also fail due to two causes: PathTear lost or node down. In the former case, if the LSP is not under MP control after the Expiration timer elapses, a manual intervention from the network operator is requested, while in the latter case, different scenarios may happen:

- CASE I - Path message and node down




Figure 10: Case I - Path Message and Node Down

Per [RFC3473], Section 7.2.2, the ingress node should wait for a configurable amount of time (30 seconds by default) and then send a PathTear message. In this case, the normal deletion procedure MUST

NOT be followed. When the Expiration timer elapses, a manual intervention from network operator is requested and normal, i.e., pre-CP-to-MP Handover, LSP processing continues.

- CASE II - Resv message and node down




Figure 11: Case II - Resv Message and Node Down

The procedure to be followed is the same depicted in CASE I. The network operator can ask for the automatic CP-to-MP procedure again after the failed node comes back up. Per [RFC3473], section 7.2, the nodes will forward the new Path and Resv messages correctly.

- CASE III - PathTear message and node down




Figure 12: Case III - PathTear Message and Node Down

This scenario can be managed as a normal PathTear lost described above in this section.

5. Minimum Information for MP-to-CP Handover

As described in Section 4, it is also possible for the ERO to contain less than the full set of path information for the LSP being handed over. This arises when only a minimal set of information is handed

to the CP by the MP at the LSP's head-end. Instead of collecting all of the LSP information (including the labels) and formatting it into an ERO, as described in Section 4, it is possible to start with a minimum amount of information. The full ERO method and the partial/no ERO method are not mutually exclusive; support of both methods is required.

At the ingress node, the information needed to specify the LSP is the outgoing interface ID, upstream label, and downstream label of this interface and egress node ID. The remaining information about an existing LSP can then be collected hop by hop, as the signaling is going on, by looking up the cross-connection table in the DP at each node along the LSP path.

Starting from the information available at the ingress LER about the outgoing interface ID of that ingress node, the incoming interface ID of the next hop can be found by looking up the link resource table/database in the LER itself.

The Path message is hence built with the LABEL_SET Object ([RFC3473]) and the UPSTREAM_LABEL Object ([RFC3473]), where the upstream label and downstream label of ingress outgoing interface of the LSP are included in these two objects. In addition to the above mentioned objects, the Path message MUST include the ADMIN_STATUS Object with the H bit set, as already defined in previous chapters for the detailed ERO-based way of proceeding. Such a Handover Path is sent to the incoming interface of the next hop. When this Path message reaches the second node along the LSP, the information about incoming interface ID and the upstream and downstream labels of this interface is extracted from it, and it is used to find next hop outgoing interface ID and the upstream/downstream labels by looking up the DP cross-connection table.

After having determined, in this way, the parameters describing the LSPs next hop, the outgoing Path message to be sent is built replacing the LABEL_SET Object and UPSTREAM_LABEL Object content with the looked-up values of upstream and downstream labels.

By repeating this procedure for each transit node along the LSP, it is possible to make the Handover Path message reach the egress node, exactly following the LSP that is in place over DP. The ERO MAY, in this case, be included in the Path message as an optional object, and MAY be filled with the LSP-relevant information down to either the port level with the interface ID or the label level with upstream and downstream labels. The ERO can be used to check the consistency of resource in the DP down to the port level or label level at each intermediate node along the LSP.

Where the DP path continues beyond the egress, by indicating the Egress label at the head-end of an LSP, the traffic can be directed to the right destination. The GMPLS signaling procedure for egress control is described in [RFC4003]

6. RSVP Message Formats

This memo does not introduce any modification in RSVP messages object composition.

7. Objects Modification

The modifications required concern two RSVP objects: the ADMIN_STATUS and ERROR_SPEC Objects.

7.1. Administrative Status Object

This memo introduces a new flag into the ADMIN_STATUS Object. The ADMIN_STATUS Object is defined in [RFC3473]. This document uses the H bit of the ADMIN_STATUS Object. The bit is bit number 25.

7.2. Error Spec Object

It is possible that a failure, such as the loss of the Data Communication Network (DCN) connection or the restart of a node, occurs during the LSP ownership handing over. In this case, the LSP Handover procedure is interrupted, the ownership of the LSP must remain with the ownership prior to the initiation of the Handover procedure. It is important that the transaction failure not affect the DP. The LSP is kept in place and no traffic hit occurs.

The failure is signaled by a PathErr message in the upstream direction and PathTear messages in the downstream direction. The PathErr messages include an ERROR_SPEC Object specifying the causes of the failure.

This memo introduces a new Error Code (with different Error Values) into the ERROR_SPEC Object, defined in [RFC2205].

The defined Error Code is "Handover Procedure Failure", and its value is 35. For this Error Code, the following Error Value sub-codes are defined:

- 1 = Cross-connection mismatch
- 2 = Other failure

8. Compatibility

The main requirement for the Handover procedure to work is that all nodes along the path MUST support the extension defined in this document. This requirement translates to an administrative requirement as it is not enforced at the protocol level. As defined, non-supporting nodes will simply propagate the H bit without setting local state. This may result in an impact on data traffic during the Handover procedure.

9. Security Considerations

The procedures described in this document rely completely on RSVP-TE messages and mechanism. The use of the H bit being set in the ADMIN_STATUS Object basically informs the receiving entity that no operations are to be done over the DP as a consequence of such special signaling flow. Using specially flagged signaling messages, we want to limit the function of setup and teardown messages to the CP, making them ineffective over related DP resource usage.

However, the Handover procedures allow the Control Plane to be used to take an LSP out of the control of the Management Plane. This could cause considerable disruption and could introduce a new security concern. As a consequence, the use of GMPLS security techniques is more important. For RSVP-TE security, please refer to [RFC3473], for the GMPLS security framework, please refer to [sec-fwk].

10. IANA Considerations

IANA manages the bit allocations for the ADMIN_STATUS Object ([RFC3473]). This document requires the allocation of the Handover bit: the H bit. IANA has allocated a bit for this purpose.

Bit Number	Hex Value	Name	Reference
25	0x00000040	Handover (H)	[RFC5852]

IANA has also allocated a new Error Code:

35 Handover failure

This Error Code has the following globally defined Error Value sub-codes:

- 1 = Cross-connection mismatch
- 2 = Other failure

11. Acknowledgments

We wish to thank Adrian Farrel, Lou Berger, Alan Elder, and Ben Campbell for their assistance and precious advice to prepare this document for publication. We also wish to thank Nicola Ciulli (Nextworks) who contributed to the initial stage of this document.

12. Contributors

Shan Zhu
Fujitsu Network Communications Inc.
2801 Telecom Parkway,
Richardson, TX 75082
USA
EMail: Shan.Zhu@us.fujitsu.com
Tel: +1-972-479-2041

Igor Bryskin
ADVA Optical Networking Inc
7926 Jones Branch Drive, Suite 615
McLean, VA 22102
USA
EMail: ibryskin@advaoptical.com

Francesco Fondelli
Ericsson
Via Negrone 1A
Genova - 16145
Italy
EMail: francesco.fondelli@ericsson.com

Lou Berger
LabN Consulting, LLC
Phone: +1 301 468 9228
EMail: lberger@labn.net

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

- [RFC2961] Berger, L., Gan, D., Swallow, G., Pan, P., Tommasi, F., and S. Molendini, "RSVP Refresh Overhead Reduction Extensions", RFC 2961, April 2001.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4003] Berger, L., "GMPLS Signaling Procedure for Egress Control", RFC 4003, February 2005.

13.2. Informative References

- [RFC5493] Caviglia, D., Bramanti, D., Li, D., and D. McDysan, "Requirements for the Conversion between Permanent Connections and Switched Connections in a Generalized Multiprotocol Label Switching (GMPLS) Network", RFC 5493, April 2009.
- [sec-fwk] Fang, L. and M. Behringer, "Security Framework for MPLS and GMPLS Networks", Work in Progress, March 2010.

Authors' Addresses

Diego Caviglia
Ericsson
Via A. Negrone 1A
Genova - Sestri Ponente 16153
Italy

E-Mail: diego.caviglia@ericsson.com

Daniele Ceccarelli
Ericsson
Via A. Negrone 1A
Genova - Sestri Ponente 16153
Italy

E-Mail: daniele.ceccarelli@ericsson.com

Dino Bramanti
Ericsson

Dan Li
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Shenzhen 518129
P.R. China

E-Mail: danli@huawei.com

Snigdho Bardalai
Fujitsu Network
2801 Telecom Parkway
Richardson, TX 75082
USA

E-Mail: sbardalai@gmail.com

