

Network Working Group
Request for Comments: 4090
Category: Standards Track

P. Pan, Ed.
Hammerhead Systems
G. Swallow, Ed.
Cisco Systems
A. Atlas, Ed.
Avici Systems
May 2005

Fast Reroute Extensions to RSVP-TE for LSP Tunnels

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document defines RSVP-TE extensions to establish backup label-switched path (LSP) tunnels for local repair of LSP tunnels. These mechanisms enable the re-direction of traffic onto backup LSP tunnels in 10s of milliseconds, in the event of a failure.

Two methods are defined here. The one-to-one backup method creates detour LSPs for each protected LSP at each potential point of local repair. The facility backup method creates a bypass tunnel to protect a potential failure point; by taking advantage of MPLS label stacking, this bypass tunnel can protect a set of LSPs that have similar backup constraints. Both methods can be used to protect links and nodes during network failure. The described behavior and extensions to RSVP allow nodes to implement either method or both and to interoperate in a mixed network.

Table of Contents

1.	Introduction	3
1.1.	Background	4
2.	Terminology	4
3.	Local Repair Techniques	6
3.1.	One-to-One Backup	6
3.2.	Facility Backup	7
4.	RSVP Extensions	8
4.1.	FAST_REROUTE Object	8
4.2.	DETOUR Object	11
4.2.1.	DETOUR Object for IPv4 Address	11
4.2.2.	DETOUR Object for IPv6 Address	12
4.3.	SESSION_ATTRIBUTE Flags	13
4.4.	RRO IPv4/IPv6 Sub-object Flags	14
5.	Head-End Behavior	15
6.	Point of Local Repair (PLR) Behavior	16
6.1.	Signaling a Backup Path	17
6.1.1.	Backup Path Identification: Sender Template-Specific	19
6.1.2.	Backup Path Identification: Path-Specific	19
6.2.	Procedures for Backup Path Computation	20
6.3.	Signaling Backups for One-to-One Protection	21
6.3.1.	Make-before-Break with Detour LSPs	22
6.3.2.	Message Handling	23
6.3.3.	Local Reroute of Traffic onto Detour LSP	23
6.4.	Signaling for Facility Protection	24
6.4.1.	Discovering Downstream Labels	24
6.4.2.	Procedures for the PLR before Local Repair	24
6.4.3.	Procedures for the PLR during Local Repair	25
6.4.4.	Processing Backup Tunnel's ERO	26
6.5.	PLR Procedures during Local Repair	26
6.5.1.	Notification of Local Repair	26
6.5.2.	Revertive Behavior	27
7.	Merge Node Behavior	28
7.1.	Handling Backup Path Messages before Failure	28
7.1.1.	Merging Backup Paths using the Sender Template-Specific Method	29
7.1.2.	Merging Detours using the Path-Specific Method	29
7.1.3.	Message Handling for Merged Detours	31
7.2.	Handling Failures	31
8.	Behavior of All LSRs	32
8.1.	Merging Detours in the Path-Specific Method	32
9.	Security Considerations	33
10.	IANA Considerations	33
11.	Contributors	35
12.	Acknowledgments	36
13.	Normative References	36

1. Introduction

This document extends RSVP [RSVP] to establish backup label-switched path (LSP) tunnels for the local repair of LSP tunnels. This extension will meet the needs of real-time applications such as voice over IP, for which user traffic should be redirected onto backup LSP tunnels in 10s of milliseconds. This timing requirement can be satisfied by computing and signaling backup LSP tunnels in advance of failure and by re-directing traffic as close to the failure point as possible. In this way, the time for redirection includes no path computation and no signaling delays, including delays to propagate failure notification between label-switched routers (LSRs). Speed of repair is the primary advantage of the methods and extensions described here. The term local repair is used when referring to techniques that re-direct traffic to a backup LSP tunnel in response to a local failure.

A protected LSP is an explicitly-routed LSP that is provided with protection. The repair methods described here are applicable only to explicitly-routed LSPs. Application of these methods to LSPs that dynamically change their routes, such as LSPs used in unicast IGP routing, is beyond the scope of this document.

Section 2 covers new terminology used in this document. Section 3 describes two basic methods for creating backup LSPs. Section 4 describes the RSVP protocol extensions to support local protection. Section 5 presents the behavior of an LSR that seeks to request local protection for an LSP. The behavior of a potential point of local repair (PLR) is given in Section 6, which describes how to determine the appropriate strategy for protecting an LSP and how to implement each of the strategies. Section 7 describes the behavior of a merge node, the LSR where a protected LSP and its backup LSP rejoin. Finally, Section 8 discusses the required behavior of other nodes in the network.

The methods discussed in this document depend upon three assumptions:

- o An LSR that is on the path of a protected LSP should always assume that it is a merge point. This is necessary because the facility backup method does not signal backups through a bypass tunnel before failure.
- o If the one-to-one backup method is used and a DETOUR object is included, the LSRs in the traffic-engineered network should support the DETOUR object. This is necessary so that the Path message containing the DETOUR object is not rejected.

- o Understanding the DETOUR object is required to support the path-specific method, which requires that LSRs in the traffic-engineered network be capable of merging detours.

1.1. Background

Several years before work began on this document, operational networks had deployed two independent methods of doing fast reroute; these methods are called here one-to-one backup and facility backup. Vendors trying to support both methods experienced compatibility problems in attempting to produce a single implementation capable of interoperating with both methods. There are technical tradeoffs between the methods. These tradeoffs are so topologically dependent that the community has not converged on a single approach.

This document rationalizes the RSVP signaling for both methods so that any implementation can recognize all fast reroute requests and clearly respond. The response may be positive if the method can be performed, or it may be a clear error to inform the requester to seek alternate backup means. This document also allows a single implementation to support both methods, thereby providing a range of capabilities. The described behavior and extensions to RSVP allow LERs and LSRs to implement either method or both.

While the two methods could in principle be used in a single network, it is expected that operators will continue to deploy either one or the other. The goal of this document is to standardize the RSVP signaling so that a network composed of LSRs that implement both methods or a network composed of some LSRs that support one method and others that support both can properly signal among those LSRs to achieve fast restoration.

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 [RFC-WORDS].

The reader is assumed to be familiar with the terminology in [RSVP] and [RSVP-TE].

LSR: Label-Switch Router.

LSP: An MPLS Label-Switched Path. In this document, an LSP will always be explicitly routed.

Local Repair: Techniques used to repair LSP tunnels quickly when a node or link along the LSP's path fails.

PLR: Point of Local Repair. The head-end LSR of a backup tunnel or a detour LSP.

One-to-One Backup: A local repair method in which a backup LSP is separately created for each protected LSP at a PLR.

Facility Backup: A local repair method in which a bypass tunnel is used to protect one or more protected LSPs that traverse the PLR, the resource being protected, and the Merge Point in that order.

Protected LSP: An LSP is said to be protected at a given hop if it has one or multiple associated backup tunnels originating at that hop.

Detour LSP: The LSP that is used to re-route traffic around a failure in one-to-one backup.

Bypass Tunnel: An LSP that is used to protect a set of LSPs passing over a common facility.

Backup Tunnel: The LSP that is used to backup up one of the many LSPs in many-to-one backup.

NHOP Bypass Tunnel: Next-Hop Bypass Tunnel. A backup tunnel that bypasses a single link of the protected LSP.

NNHOP Bypass Tunnel: Next-Next-Hop Bypass Tunnel. A backup tunnel that bypasses a single node of the protected LSP.

Backup Path: The LSP that is responsible for backing up one protected LSP. A backup path refers to either a detour LSP or a backup tunnel.

MP: Merge Point. The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure. The same LSR may be both an MP and a PLR simultaneously.

DMP: Detour Merge Point. In the case of one-to-one backup, this is an LSR where multiple detours converge. Only one detour is signaled beyond that LSR.

Reroutable LSP: Any LSP for which the head-end LSR requests local protection. See Section 5 for more detail.

CSPF: Constraint-based Shortest Path First.


SRLG Disjoint: A path is considered to be SRLG disjoint from a given link or node if the path does not use any links or nodes which belong to the same SRLG as that given link or node.

3. Local Repair Techniques

Two different methods for local protection are described. In the one-to-one backup method, a PLR computes a separate backup LSP, called a detour LSP, for each LSP that the PLR protects. In the facility backup method, the PLR creates a single bypass tunnel that can be used to protect multiple LSPs.

3.1. One-to-One Backup

In the one-to-one backup method, a label-switched path is established that intersects the original LSP somewhere downstream of the point of link or node failure. A separate backup LSP is established for each LSP that is backed up.



```

Protected LSP:  [R1->R2->R3->R4->R5]
R1's Backup:    [R1->R6->R7->R8->R3]
R2's Backup:    [R2->R7->R8->R4]
R3's Backup:    [R3->R8->R9->R5]
R4's Backup:    [R4->R9->R5]

```

Example 1. One-to-One Backup Technique

In the simple topology shown in Example 1, the protected LSP runs from R1 to R5. R2 can provide user traffic protection by creating a partial backup LSP that merges with the protected LSP at R4. We refer to a partial one-to-one backup LSP [R2->R7->R8->R4] as a detour.


To protect an LSP that traverses N nodes fully, there could be as many as (N - 1) detours. Example 1 shows the paths for the detours necessary to protect fully the LSP in the example. To minimize the number of LSPs in the network, it is desirable to merge a detour back to its protected LSP, when feasible. When a detour LSP intersects its protected LSP at an LSR with the same outgoing interface, it will be merged.

When a failure occurs along the protected LSP, the PLR redirects traffic onto the local detour. For instance, if the link [R2->R3] fails in Example 1, R2 will switch traffic received from R1 onto the protected LSP along link [R2->R7], using the label received when R2 created the detour. When R4 receives traffic with the label provided for R2's detour, R4 will switch that traffic onto link [R4->R5], using the label received from R5 for the protected LSP. At no point does the depth of the label stack increase as a result of the detour. While R2 is using its detour, traffic will take the path [R1->R2->R7->R8->R4->R5].

3.2. Facility Backup

The facility backup method takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. We call such an LSP tunnel a bypass tunnel.

The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the PLR. Naturally, this constrains the set of LSPs being backed up via that bypass tunnel to those that pass through some common downstream node. All LSPs that pass through the point of local repair and through this common node that do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.



Example 2. Facility Backup Technique

In Example 2, R2 has built a bypass tunnel that protects against the failure of link [R2->R3] and node [R3]. The doubled lines represent this tunnel. This technique provides a scalability improvement, in that the same bypass tunnel can also be used to protect LSPs from any of R1, R2, or R8 to any of R4, R5, or R9. Example 2 describes three different protected LSPs that are using the same bypass tunnel for protection.

As with the one-to-one method, there could be as many as (N-1) bypass tunnels to fully protect an LSP that traverses N nodes. However, each of those bypass tunnels could protect a set of LSPs.

When a failure occurs along a protected LSP, the PLR redirects traffic into the appropriate bypass tunnel. For instance, if link [R2->R3] fails in Example 2, R2 will switch traffic received from R1 on the protected LSP onto link [R2->R6]. The label will be switched for one which will be understood by R4 to indicate the protected LSP, and the bypass tunnel's label will then be pushed onto the label-stack of the redirected packets. If penultimate-hop-popping is used, the merge point in Example 2, R4, will receive the redirected packet with a label indicating the protected LSP that the packet is to follow. If penultimate-hop-popping is not used, R4 will pop the bypass tunnel's label and examine the label underneath to determine the protected LSP that the packet is to follow. When R2 is using the bypass tunnel for protected LSP 1, the traffic takes the path [R1->R2->R6->R7->R4->R5]; the bypass tunnel is the connection between R2 and R4.

4. RSVP Extensions

This specification defines two additional objects, FAST_REROUTE and DETOUR, to extend RSVP-TE for fast-reroute signaling. These new objects are backward compatible with LSRs that do not recognize them (see section 3.10 in [RSVP]). Both objects can only be carried in RSVP Path messages.

The SESSION_ATTRIBUTE and RECORD_ROUTE objects are also extended to support bandwidth and node protection features.

4.1. FAST_REROUTE Object

The FAST_REROUTE object is used to control the backup used for the protected LSP. This specifies the setup and hold priorities, session attribute filters, and bandwidth to be used for protection. It also allows a specific local protection method to be requested. This object MUST only be inserted into the PATH message by the head-end LER and MUST NOT be changed by downstream LSRs. The FAST_REROUTE object has the following format:

Class-Num = 205

C-Type = 1

0	1	2	3
Length (bytes)		Class-Num	C-Type
Setup Prio	Hold Prio	Hop-limit	Flags
Bandwidth			
Include-any			
Exclude-any			
Include-all			

Setup Priority

The priority of the backup path with respect to taking resources, in the range 0 to 7. The value 0 is the highest priority. Setup Priority is used in deciding whether this session can preempt another session. See [RSVP-TE] for the usage on priority.

Holding Priority

The priority of the backup path with respect to holding resources, in the range 0 to 7. The value 0 is the highest priority. Holding Priority is used in deciding whether this session can be preempted by another session. See [RSVP-TE] for the usage on priority.

Hop-limit

The maximum number of extra hops the backup path is allowed to take, from current node (a PLR) to an MP, with PLR and MP excluded from the count. For example, hop-limit of 0 means that only direct links between PLR and MP can be considered.

Flags

0x01 One-to-One Backup Desired

Requests protection via the one-to-one backup method.

0x02 Facility Backup Desired

Requests protection via the facility backup method.

Bandwidth

Bandwidth estimate; 32-bit IEEE floating point integer, in bytes per second.

Exclude-any

A 32-bit vector representing a set of attribute filters associated with a backup path, any of which renders a link unacceptable.

Include-any

A 32-bit vector representing a set of attribute filters associated with a backup path, any of which renders a link acceptable (with respect to this test). A null set (all bits set to zero) automatically passes.

Include-all

A 32-bit vector representing a set of attribute filters associated with a backup path, all of which must be present for a link to be acceptable (with respect to this test). A null set (all bits set to zero) automatically passes.

The two high-order bits of the Class-Num (11) cause nodes that do not understand the object to ignore it and pass it forward unchanged.

For informational purposes, a different C-Type value and format for the FAST_REROUTE object are specified below. This is used by legacy implementations. The meaning of the fields is the same as that described for C-Type 1.

Class-Num = 205

C-Type = 7

0	1	2	3
Length (bytes)	Class-Num	C-Type	
Setup Prio	Hold Prio	Hop-limit	Reserved
Bandwidth			
Include-any			
Exclude-any			

Unknown C-Types should be treated as specified in [RSVP] Section 3.10.

4.2. DETOUR Object

The DETOUR object is used in the one-to-one backup method to identify detour LSPs.

4.2.1. DETOUR Object for IPv4 Address

Class-Num = 63

C-Type = 7

0	1	2	3
Length (bytes)	Class-Num	C-Type	
PLR_ID 1			
Avoid_Node_ID 1			
....			
PLR_ID n			
Avoid_Node_ID n			

PLR_ID (1 - n)

IPv4 address identifying the PLR that is the beginning point of the detour. Any local address on the PLR can be used.

Avoid_Node_ID (1 - n)

IPv4 address identifying the immediate downstream node that the PLR is trying to avoid. Any local address of the downstream node can be used. This field is mandatory and is used by the MP for the merging rules discussed below.

4.2.2. DETOUR Object for IPv6 Address

Class-Num = 63

C-Type = 8

0	1	2	3
+-----+-----+-----+-----+			
	Length (bytes)		Class-Num
+-----+-----+-----+-----+			
	PLR_ID 1		
+-----+-----+-----+-----+			
	PLR_ID 1 (continued)		
+-----+-----+-----+-----+			
	PLR_ID 1 (continued)		
+-----+-----+-----+-----+			
	PLR_ID 1 (continued)		
+-----+-----+-----+-----+			
	Avoid_Node_ID 1		
+-----+-----+-----+-----+			
	Avoid_Node_ID 1 (continued)		
+-----+-----+-----+-----+			
	Avoid_Node_ID 1 (continued)		
+-----+-----+-----+-----+			
	Avoid_Node_ID 1 (continued)		
+-----+-----+-----+-----+			
//		//
+-----+-----+-----+-----+			

PLR_ID (1 - n)

An IPv6 128-bit unicast host address identifying the PLR that is the beginning point of the detour. Any local address on the PLR can be used.

Avoid_Node_ID (1 - n)

An IPv6 128-bit unicast host address identifying the immediate downstream node that the PLR is trying to avoid. Any local address on the downstream node can be used. This field is

mandatory and is used by the MP for the merging rules discussed below.

There can be more than one pair of (PLR_ID, Avoid_Node_ID) entries in a DETOUR object. If detour merging is desired, after each merging operation, the Detour Merge Point should combine all the merged detours in subsequent Path messages.

The high-order bit of the Class-Num is zero; LSRs that do not support the DETOUR objects MUST reject any Path message containing a DETOUR object and send a PathErr to notify the PLR. This PathErr SHOULD be generated as specified in [RSVP] for unknown objects with a Class-Num of the form "0bbbbbbb".

Unknown C-Types should be treated as specified in [RSVP] Section 3.10.

4.3. SESSION_ATTRIBUTE Flags

To request bandwidth and node protection explicitly, two new flags are defined in the SESSION_ATTRIBUTE object.

For both C-Type 1 and 7, the SESSION_ATTRIBUTE object currently has the following flags defined [RSVP-TE]:

Local protection desired: 0x01

This flag permits transit routers to use a local repair mechanism that may result in violation of the explicit route object. When a fault is detected on an adjacent downstream link or node, a transit node may reroute traffic for fast service restoration.

Label recording desired: 0x02

This flag indicates that label information should be included when doing a route record.

SE Style desired: 0x04

This flag indicates that the tunnel ingress node may choose to reroute this tunnel without tearing it down. A tunnel egress node SHOULD use the SE Style when responding with a Resv message. When requesting fast reroute, the head-end LSR SHOULD set this flag; this is not necessary for the path-specific method of the one-to-one backup method.

The following new flags are defined:

Bandwidth protection desired: 0x08

This flag indicates to the PLRs along the protected LSP path that a backup path with a bandwidth guarantee is desired. The bandwidth to be guaranteed is that of the protected LSP, if no FAST_REROUTE object is included in the PATH message; if a FAST_REROUTE object is in the PATH message, then the bandwidth specified therein is to be guaranteed.

Node protection desired: 0x10

This flag indicates to the PLRs along a protected LSP path that a backup path that bypasses at least the next node of the protected LSP is desired.

4.4. RRO IPv4/IPv6 Sub-object Flags

To report whether bandwidth and/or node protection are provided as requested, we define two new flags in the RRO IPv4 sub-object.

The RRO IPv4 and IPv6 address sub-objects currently have the following flags defined [RSVP-TE]:

Local protection available: 0x01

Indicates that the link downstream of this node is protected via a local repair mechanism, which can be either one-to-one or facility backup.

Local protection in use: 0x02

Indicates that a local repair mechanism is in use to maintain this tunnel (usually in the face of an outage of the link it was previously routed over, or an outage of the neighboring node).

Two new flags are defined:

Bandwidth protection: 0x04

The PLR will set this bit when the protected LSP has a backup path that is guaranteed to provide the desired bandwidth that is specified in the FAST_REROUTE object or the bandwidth of the protected LSP, if no FAST_REROUTE object was included. The PLR may set this whenever the desired bandwidth is guaranteed; the PLR MUST set this flag when the desired bandwidth is guaranteed

and the "bandwidth protection desired" flag was set in the SESSION_ATTRIBUTE object. If the requested bandwidth is not guaranteed, the PLR MUST NOT set this flag.

Node protection: 0x08

The PLR will set this bit when the protected LSP has a backup path that provides protection against a failure of the next LSR along the protected LSP. The PLR may set this whenever node protection is provided by the protected LSP's backup path; the PLR MUST set this flag when the node protection is provided and the "node protection desired" flag was set in the SESSION_ATTRIBUTE object. If node protection is not provided, the PLR MUST NOT set this flag. Thus, if a PLR could only set up a link-protection backup path, the "Local protection available" bit will be set, but the "Node protection" bit will be cleared.

5. Head-End Behavior

The head-end of an LSP determines whether local protection should be requested for that LSP and which local protection method is desired for the protected LSP. The head-end also determines what constraints should be requested for the backup paths of a protected LSP.

To indicate that an LSP should be locally protected, the head-end LSR MUST either set the "local protection desired" flag in the SESSION_ATTRIBUTE object or include a FAST_REROUTE object in the PATH message, or both. The "local protection desired" flag in the SESSION_ATTRIBUTE object SHOULD always be set. If a head-end LSR signals a FAST_REROUTE object, it MUST be stored for Path refreshes.

The head-end LSR of a protected LSP MUST set the "label recording desired" flag in the SESSION_ATTRIBUTE object. This facilitates the use of the facility backup method. If node protection is desired, the head-end LSR should set the "node protection desired" flag in the SESSION_ATTRIBUTE object; otherwise, this flag should be cleared. Similarly, if a guarantee of bandwidth protection is desired, then the "bandwidth protection desired" flag in the SESSION_ATTRIBUTE object should be set; otherwise, this flag should be cleared. If the head-end LSR determines that control of the backup paths for the protected LSP is desired, then the LSR should include the FAST_REROUTE object. The PLRs will use the attribute filters, bandwidth, hop-limit, and priorities to determine the backup paths.

If the head-end LSR desires that the one-to-one backup method be used for the protected LSP, then the head-end LSR should include a FAST_REROUTE object and set the "one-to-one backup desired" flag. If

the head-end LSR desires that the protected LSP be protected via the facility backup method, then the head-end LSR should include a FAST_REROUTE object and set the "facility backup desired" flag. The lack of a FAST_REROUTE object, or having both these flags clear, should be treated by PLRs as a lack of preference. If both flags are set, a PLR may use either method or both.

The head-end LSR of a protected LSP MUST support the additional flags defined in Section 4.4 being set or clear in the RRO IPv4 and IPv6 sub-objects. The head-end LSR of a protected LSP MUST support the RRO Label sub-object.

If the head-end LSR of an LSP determines that local protection is newly desired, this SHOULD be signaled via make-before-break.

6. Point of Local Repair (PLR) Behavior

Every LSR along a protected LSP (except the egress) MUST follow the PLR behavior described in this document.

A PLR SHOULD support the FAST_REROUTE object, the "local protection desired", "label recording desired", "node protection desired", and "bandwidth protection desired" flags in the SESSION_ATTRIBUTE object, and the "local protection available", "local protection in use", "bandwidth protection", and "node protection" flags in the RRO IPv4 and IPv6 sub-objects. A PLR MAY support the DETOUR object.

A PLR MUST consider an LSP to have asked for local protection if the "local protection desired" flag is set in the SESSION_ATTRIBUTE object and/or the FAST_REROUTE object is included. If the FAST_REROUTE object is included, a PLR SHOULD consider providing one-to-one protection if the "one-to-one desired" is set, and it SHOULD consider providing facility backup if the "facility backup desired" flag is set. If the "node protection desired" flag is set, the PLR SHOULD try to provide node protection; if this is not feasible, the PLR SHOULD then try to provide link protection. If the "bandwidth protection guaranteed" flag is set, the PLR SHOULD try to provide a bandwidth guarantee; if this is not feasible, the PLR SHOULD then try to provide a backup without a guarantee of the full bandwidth.

The following treatment for the RRO IPv4 or IPv6 sub-object's flags must be followed if an RRO is included in the protected LSP's RESV message. Based on this additional information, the head-end may take appropriate actions.

- Until a PLR has a backup path available, the PLR MUST clear the relevant four flags in the corresponding RRO IPv4 or IPv6 sub-object.
- Whenever the PLR has a backup path available, the PLR MUST set the "local protection available" flag. If no established one-to-one backup LSP or bypass tunnel exists, or if the one-to-one LSP and the bypass tunnel is in "DOWN" state, the PLR MUST clear the "local protection available" flag in its IPv4 (or IPv6) address sub-object of the RRO and SHOULD send the updated RESV.
- The PLR MUST clear the "local protection in use" flag unless it is actively redirecting traffic into the backup path instead of along the protected LSP.
- The PLR SHOULD also set the "node protection" flag if the backup path protects against the failure of the immediate downstream node, and, if the path does not, the PLR SHOULD clear the "node protection" flag. This MUST be done if the "node protection desired" flag was set in the SESSION_ATTRIBUTE object.
- The PLR SHOULD set the "bandwidth protection" flag if the backup path offers a bandwidth guarantee, and, if the path does not, the PLR SHOULD clear the "bandwidth protection" flag. This MUST be done if the "bandwidth protection desired" flag was set in the SESSION_ATTRIBUTE object.

6.1. Signaling a Backup Path

A number of objectives must be met to obtain a satisfactory signaling solution. These are summarized as follows:

1. Unambiguously and uniquely identifying backup paths.
2. Unambiguously associating protected LSPs with their backup paths.
3. Working with both global and non-global label spaces.
4. Allowing merging of backup paths.
5. Maintaining RSVP state during and after fail-over.

LSP tunnels are identified by a combination of the SESSION and SENDER_TEMPLATE objects [RSVP-TE]. The relevant fields are as follows.

IPv4 (or IPv6) tunnel end point address

IPv4 (or IPv6) address of the egress node for the tunnel.

Tunnel ID

A 16-bit identifier used in the SESSION that remains constant over the life of the tunnel.

Extended Tunnel ID

A 32-bit (IPv4) or 128-bit (IPv6) identifier used in the SESSION that remains constant over the life of the tunnel. Normally it is set to all zero. Ingress nodes that wish to narrow the scope of a SESSION to the ingress-egress pair may place their IP address here as a globally unique identifier.

IPv4 (or IPv6) tunnel sender address

IPv4 (or IPv6) address for a sender node.

LSP ID

A 16-bit identifier used in the SENDER_TEMPLATE and the FILTER_SPEC, which can be changed to allow a sender to share resources with itself.

The first three of these are in the SESSION object and are the basic identification for the tunnel. Setting the "Extended Tunnel ID" to an IP address of the head-end LSR allows the scope of the SESSION to be narrowed to only LSPs sent by that LSR. A backup LSP is considered part of the same session as its protected LSP; therefore these three cannot be varied.

The last two are in the SENDER_TEMPLATE. Multiple LSPs in the same SESSION may be protected and may take different routes; this is common when a tunnel is rerouted using make-before-break. A backup path must be clearly identified with its protected LSP to allow correct merging and state treatment. Therefore, a backup path must inherit its LSP ID from the associated protected LSP. Thus, the only field in the SESSION and SENDER_TEMPLATE objects that could be varied between a backup path and a protected LSP is the "IPv4 (or IPv6) tunnel sender address" in the SENDER_TEMPLATE.

There are two different methods to uniquely identify a backup path, described below.

6.1.1. Backup Path Identification: Sender Template-Specific

In this approach, the SESSION object and the LSP_ID are copied from the protected LSP. The "IPv4 tunnel sender address" is set to an address of the PLR. If the head-end of a tunnel is also acting as the PLR, it MUST choose an IP address different from the one used in the SENDER_TEMPLATE of the original LSP tunnel.

When the sender template-specific approach is used, the protected LSPs and the backup paths SHOULD use the Shared Explicit (SE) style. This allows bandwidth sharing between multiple backup paths. The backup paths and the protected LSP MAY be merged by the Detour Merge Points, when the ERO from the MP to the egress is the same on each LSP to be merged, as specified in [RSVP-TE].

6.1.2. Backup Path Identification: Path-Specific

In this approach, rather than vary the SESSION or SENDER_TEMPLATE objects, an implementation uses a new object, the DETOUR object, to distinguish between PATH messages for a backup path and the protected LSP.

Thus, the backup paths use the same SESSION and SENDER_TEMPLATE objects as the ones used in the protected LSP. The presence of a DETOUR object in Path messages signifies a backup path; the presence of a FAST_REROUTE object and/or the "local protection requested" flag in the SESSION_ATTRIBUTE object indicates a protected LSP.

In the path message-specific approach, an LSR merges Path messages that are received with the same SESSION and SENDER_TEMPLATE objects and that also have the same next-hop object. Without this behavior, it would be impossible to associate the multiple RESV messages with the backup paths. However, this merging behavior reduces the total number of RSVP states inside the network at the expense of merging LSPs with different EROs.

6.2. Procedures for Backup Path Computation

Before a PLR can create a detour or a bypass tunnel, the desired explicit route must be determined. This can be done using a CSPF (Constraint-based Shortest Path First) computation. Before this CSPF computation, the following information must be collected at a PLR:

- The list of downstream nodes that the protected LSP passes through. This information is readily available from the RECORD_ROUTE objects during LSP setup. This information is also available from the ERO. However, if the ERO contains loose sub-objects, the ERO may not provide adequate information.
- The downstream links/nodes that we want to protect against. Once again, this information is learned from the RECORD_ROUTE objects. Whether node protection is desired is determined by the "node protection" flag in the SESSION_ATTRIBUTE object and local policy.
- The upstream uni-directional links that the protected LSP passes through. This information is learned from the RECORD_ROUTE objects; it is only needed for setting up one-to-one protection. In the path-specific method, it is necessary to avoid the detour and the protected LSP sharing a common next-hop upstream of the failure. In the sender template-specific mode, this same restriction is necessary to avoid sharing bandwidth between the detour and its protected LSP, where that bandwidth has been reserved only once.
- The link attribute filters to be applied. These are derived from the FAST_REROUTE object, if it is included in the PATH message, or from the SESSION_ATTRIBUTE object otherwise.
- The bandwidth to be used is found in the FAST_REROUTE object, if it is included in the PATH message, or in the SESSION_ATTRIBUTE object otherwise. Local policy may modify the bandwidth to be reserved.
- The hop-limit, if a FAST_REROUTE object was included in the PATH message.

When a CSPF algorithm is used to compute the backup route, the following constraints must be satisfied:

- For detour LSPs, the destination MUST be the tail-end of the protected LSP. For bypass tunnels (Section 7), the destination MUST be the address of the MP.

- When one-to-one protection is set up by using the path-specific method, a detour MUST not traverse the upstream links of the protected LSP in the same direction. This prevents the possibility of early merging of the detour into the protected LSP. When one-to-one protection is set up using the sender-template-specific method, a detour should not traverse the upstream links of the protected LSP in the same direction. This prevents sharing the bandwidth between a protected LSP and its backup upstream of the failure where the bandwidth would be used twice in the event of a failure.
- The backup LSP cannot traverse the downstream node and/or link whose failure is being protected against. Note that if the PLR is the penultimate hop, node protection is not possible, and only the downstream link can be avoided. The backup path may be computed to be SRLG disjoint from the downstream node and/or link being avoided.
- The backup path must satisfy the resource requirements of the protected LSP. This includes the link attribute filters, bandwidth, and hop limits determined from the FAST_REROUTE object and the SESSION_ATTRIBUTE object.

If such computation succeeds, the PLR should attempt to establish a backup path. The PLR may schedule a re-computation at a later time to discover better paths that might have emerged. If for any reason, the PLR is unable to bring up a backup path, it must schedule a retry at a later time.

6.3. Signaling Backups for One-to-One Protection

Once a PLR has decided to protect an LSP locally with one-to-one backup and has identified the desired path, it signals for the detour.

The following describes the transformation to be performed upon the protected LSP's PATH message to create the detour LSP's PATH message.

- If the sender template-specific method is to be used, then the PLR MUST change the "IPv4 (or IPv6) tunnel sender address" of the SENDER_TEMPLATE to an address belonging to the PLR that is not the same as that used for the protected LSP. Additionally, the DETOUR object MAY be added to the PATH message.
- If the path-specific method is to be used, then the PLR MUST add a DETOUR object to the PATH message.

- The SESSION_ATTRIBUTE flags "Local protection desired", "Bandwidth protection desired", and "Node protection desired" MUST be cleared. The "Label recording desired" flag MAY be modified. If the Path Message contained a FAST_REROUTE object and the ERO is not completely strict, the Include-any, Exclude-any, and Include-all fields of the FAST_REROUTE object SHOULD be copied to the corresponding fields of the SESSION_ATTRIBUTE object.
- If the protected LSP's Path message contained a FAST_REROUTE object, this object MUST be removed from the detour LSP's PATH message.
- The PLR MUST generate an EXPLICIT_ROUTE object toward the egress. First, the PLR must remove all sub-objects preceding the first address belonging to the Merge Point. Then the PLR SHOULD add sub-objects corresponding to the desired backup path between the PLR and the MP.
- The SENDER_TSPEC object SHOULD contain the bandwidth information from the received FAST_REROUTE object, if included in the protected LSP's PATH message.
- The RSVP_HOP object containing one of the PLR's IP address.
- The detour LSPs MUST use the same reservation style as the protected LSP. This must be correctly reflected in the SESSION_ATTRIBUTE object.

Detour LSPs operate like regular LSPs. Once a detour path is successfully computed and the detour LSP is established, the PLR need not compute detour routes again, unless (1) the contents of FAST_REROUTE have changed or (2) the downstream interface and/or the nexthop router for a protected LSP has changed. The PLR may recompute detour routes at any time.

6.3.1. Make-before-Break with Detour LSPs

If the sender template-specific method is used, it is possible to do make-before-break with detour LSPs. This is done using two different IP addresses belonging to the PLR (which were not used in the SENDER_TEMPLATE of the protected LSP). If the current detour LSP uses the first IP address in its SENDER_TEMPLATE, then the new detour LSP should be signaled by using the second IP address in its SENDER_TEMPLATE. Once the new detour LSP has been created, the current detour LSP can be torn down. By alternating the use of these IP addresses, the current and new detour LSPs will have different SENDER_TEMPLATES and, thus, different state in the downstream LSRs.

This make-before-break mechanism, which changes the PLR IP address in the DETOUR object instead, is not feasible with the path-specific method, as the PATH messages for new and current detour LSPs may be merged if they share a common next-hop.

6.3.2. Message Handling


LSRs must process the detour LSPs independently of the protected LSPs to avoid triggering the LSP loop detection procedure described in [RSVP-TE].

The PLR MUST not mix the messages for the protected and the detour LSPs. When a PLR receives Resv, ResvTear, and PathErr messages from the downstream detour destination, the messages MUST not be forwarded upstream. Similarly, when a PLR receives ResvErr and ResvConf messages from a protected LSP, it MUST not propagate them onto the associated detour LSP.

A session tear-down request is normally originated by the sender via PathTear messages. When a PLR node receives a PathTear message from upstream, it MUST delete both the protected and the detour LSPs. The PathTear messages MUST propagate to both protected and detour LSPs. During error conditions, the LSRs may send ResvTear messages to fix problems on the failing path. When a PLR node receives the ResvTear messages from downstream for a protected LSP, as long as a detour is up, the ResvTear messages MUST not be sent further upstream. PathErrs should be treated similarly.

6.3.3. Local Reroute of Traffic onto Detour LSP

When the PLR detects a failure on the protected LSP, the PLR MUST rapidly switch packets to the protected LSP's backup LSP instead of to the protected LSP's normal out-segment. The goal of this method is to effect the redirection within 10s of milliseconds.



Protected LSP: [R1->R2->R3->R4->R5]

Detour LSP: [R2->R6->R7->R4]

Example 3. Redirect to Detour

In Example 3, if the link [R2->R3] fails, R2 would do the following. Any traffic received on link [R1->R2] with label L32 would be sent on link [R2->R6] with label L46 (along the detour LSP) instead of on link [R3->R4] with label L34 (along the protected LSP). The merge point R4 would recognize that packets received on link [R7->R4] with label L44 should be sent on link [R4->R5] with label L35 and that they should be merged with the protected LSP.

6.4. Signaling for Facility Protection

A PLR may use one or more bypass tunnels to protect against the failure of a link and/or a node. These bypass tunnels may be set up in advance or may be dynamically created as new protected LSPs are signaled.

6.4.1. Discovering Downstream Labels

To support facility backup, the PLR must determine a label that will indicate to the MP that packets received with that label should be switched along the protected LSP. This can be done without explicitly signaling the backup path if the MP uses a label space global to that LSR.

As described in Section 6, the head-end LSR MUST set the "label recording requested" flag in the SESSION_ATTRIBUTE object for LSPs requesting local protection. This will cause (as specified in [RSVP-TE]) all LSRs to record their INBOUND labels and to note via a flag whether the label is global to the LSR. Thus, when a protected LSP is first signaled through a PLR, the PLR can examine the RRO in the Resv message and learn about the incoming labels that are used by all downstream nodes for this LSP.

When MPs use per-interface label spaces, the PLR must send Path messages (for each protected LSP using a bypass tunnel) via that bypass tunnel prior to the failure in order to discover the appropriate MP label. The signaling procedures for this are in Section 6.4.3 below.

6.4.2. Procedures for the PLR before Local Repair

A PLR that determines to use facility-backup to protect a given LSP should select a bypass tunnel to use, taking into account whether node protection is to be provided, what bandwidth was requested, whether a bandwidth guarantee is desired, and what link attribute filters were specified in the FAST_REROUTE object. The selection of a bypass tunnel for a protected LSP is performed by the PLR when the LSP is first set up.

6.4.3. Procedures for the PLR during Local Repair

When the PLR detects a link or/and node failure condition, it has to reroute the data traffic onto the bypass tunnel and to start sending the control traffic for the protected LSP onto the bypass tunnel.

The backup tunnel is identified by using the sender template-specific method. The procedures to follow are similar to those described in Section 6.3.

- The SESSION is unchanged.
- The SESSION_ATTRIBUTE is unchanged except as follows: The "Local protection desired", "Bandwidth protection desired", and "Node protection desired" flags SHOULD be cleared. The "Label recording desired" MAY be modified.
- The IPv4 (or IPv6) tunnel sender address of the SENDER_TEMPLATE is set to an address belonging to the PLR.
- The RSVP_HOP object MUST contain an IP source address belonging to the PLR. Consequently, the MP will send messages back to the PLR with that IP address as the destination.
- The PLR MUST generate an EXPLICIT_ROUTE object toward the egress. Detailed ERO processing is described below.
- The RRO object may have to be updated as described in Section 6.5.

The PLR sends Path, PathTear, and ResvConf messages via the backup tunnel. The MP sends Resv, ResvTear, and PathErr messages by sending them directly to the address in the RSVP_HOP object, as specified in [RSVP].

If it is necessary to signal the backup prior to failure to determine the MP label to use, then the same Path message is sent. In this case, the PLR SHOULD continue to send Path messages for the protected LSP along the normal route. PathTear messages should be duplicated, with one sent along the normal route and one sent through the bypass tunnel. The MP should duplicate the Resv and ResvTear messages and send them to both the PLR and the LSR indicated by the protected LSP's RSVP_HOP object.

6.4.4. Processing Backup Tunnel's ERO

Procedures for ERO processing are described in [RSVP-TE]. This section describes additional ERO update procedures for Path messages that are sent over bypass tunnels. If normal ERO processing rules were followed, the Merge Point would examine the first sub-object and likely reject it (Bad initial sub-object). This is because the unmodified ERO might contain the IP address of a bypassed node (in the case of a NNHOP Bypass Tunnel) or of an interface that is currently down (in the case of a NHOP Backup Tunnel). For this reason, the PLR invokes the following ERO procedures before sending a Path message via a bypass tunnel.

Sub-objects belonging to abstract nodes that precede the Merge Point are removed, along with the first sub-object belonging to the MP. A sub-object identifying the Backup Tunnel destination is then added.

More specifically, the PLR MUST:

- remove all the sub-objects proceeding the first address belonging to the MP, and
- replace this first MP address with an IP address of the MP.
(Note that this could be same address that was just removed.)

6.5. PLR Procedures during Local Repair

In addition to the method-specific signaling and packet treatment, there is common signaling that should be followed.

During fast reroute, for each protected LSP containing an RRO object, the PLR obtains the RRO from the protected LSP's stored RESV. The PLR MUST update the IPv4 or IPv6 sub-object it inserted into the RRO by setting the "Local protection in use" and "Local Protection Available" flags.

6.5.1. Notification of Local Repair

In many situations, the route used during local repair will be less than optimal. The purpose of local repair is to keep high priority and loss-sensitive traffic flowing while a more optimal re-routing of the tunnel can be effected by the head-end of the tunnel. Thus, the head-end has to know of the failure so that it may re-signal an optimal LSP.

To provide this notification, the PLR SHOULD send a Path Error message with error code of "Notify" (Error code = 25) and an error value field of ss00 cccc cccc cccc, where ss=00 and the sub-code = 3 ("Tunnel locally repaired") (see [RSVP-TE]).

Additionally, a head-end may detect that an LSP has to be moved to a more optimal path by noticing failures reported via the IGP. Note that in the case of inter-area TE LSP (TE LSP spanning areas), the head-end LSR will have to rely exclusively on Path Error messages to be informed of failures in another area.

6.5.2. Revertive Behavior

Upon a failure event, a protected TE LSP is locally repaired by the PLR. There are two basic strategies for restoring the TE LSP to a full working path.

- Global revertive mode: The head-end LSR of each tunnel is responsible for reoptimizing the TE LSPs that used the failed resource. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or ISIS LSPs, and timers. Note that this re-optimization process may proceed as soon as the failure is detected. It is not tied to the restoration of the failed resource.
- Local revertive mode: Upon detecting that the resource is restored, the PLR re-signals each of the TE LSPs that used to be routed over the restored resource. Every TE LSP successfully re-signaled along the restored resource is switched back.

There are several circumstances in which a local revertive mode might not be desirable. In the case of resource flapping (not an uncommon failure type), this could generate multiple traffic disruptions. Therefore, in the local revertive mode, the PLR should implement a means to dampen the re-signaling process in order to limit potential disruptions due to flapping.

In the local revertive mode, any TE LSP will be switched back, without any distinction, whereas in the global revertive mode, the decision to reuse the restored resource is made by the head-end LSR based on the TE LSP attributes. When the head-end learns of the failure, it may reoptimize the protected LSP tunnel along a different and more optimal path, as it has a more complete view of the resources and TE LSP constraints. This means that the old LSP that has been reverted to may no longer be optimal. Note that in the case of inter-area LSP, where the TE LSP path computation might be done on some Path Computation Element, the reoptimization process can

still be triggered on the Head-End LSP. The local revertive mode is optional.

However, there are circumstances in which the head-end does not have the ability to reroute the TE LSP (e.g., if the protected LSP is pinned down, as may be desirable if the paths are determined by using an off-line optimization tool), or if the head-end does not have the complete TE topology information (depending on the path computation scenario). In those cases, the local revertive mode might be an interesting option.

The globally revertive mode SHOULD always be used. Note that a link or node "failure" may be due to the facility being permanently taken out of service. Local revertive mode is optional. When used in combination, the global mode may rely solely on timers to do the reoptimization. When local revertive mode is not used, head-end LSRs SHOULD react to RSVP error messages and/or IGP indications in order to make a timely response.

Interoperability: If a PLR is configured with the local revertive mode but the MP is not, any attempt from the PLR to resignal the TE LSP over the restored resource will fail, as the MP will not send any Resv message. The PLR will still refresh the TE LSP over the backup tunnel. The TE LSP will not revert to the restored resource; instead, it will continue to use the backup until it is re-optimized.

7. Merge Node Behavior

An LSR is a Merge Point if it receives the Path message for a protected LSP and one or more messages for a backup LSP that is merged into that protected LSP. In the one-to-one backup method, the LSR is aware that it is a merge node prior to failure. In the facility backup method, the LSR may not know that it is a Merge Point until a failure occurs and it receives a backup LSP's Path message. Therefore, an LSR that is on the path of a protected LSP SHOULD always assume that it is a merge point.

When a MP receives a backup LSP's Path message through a bypass tunnel, the Send_TTL in the Common Header may not match the TTL of the IP packet within which the Path message was transported. This is expected behavior.

7.1. Handling Backup Path Messages before Failure

There are two circumstances in which a Merge Point will receive Path messages for a backup path prior to failure. In the first case, if a PLR is providing local protection via the one-to-one backup method, the detour will be signaled and must be properly handled by the MP.

In this case, the backup LSP may be signaled via the sender template-specific method or via the path-specific method.

In the second case, if the Merge Point does not provide labels global to the MP and record them in a Label sub-object of the RRO, or if the PLR does not use such recorded information, the PLR may signal the backup path as described in Section 6.4.1. This will determine the label to use if the PLR is providing protection according to the facility backup method. In this case, the backup LSP is signaled via the sender template-specific method.

The reception of a backup LSP's path message does not indicate that a failure has occurred or that the incoming protected LSP will no longer be used.

7.1.1. Merging Backup Paths using the Sender Template-Specific Method

An LSR may receive multiple Path messages for one or more backup LSPs and, possibly, for the protected LSP. Each of these Path messages will have a different SENDER_TEMPLATE. The protected LSP can be recognized because it will include the FAST_REROUTE object or have the "local protection desired" flag set in the SESSION_ATTRIBUTE object, or both.

If the outgoing interface and next-hop LSR are the same, then the Path messages are eligible for merging. Similarly to the specification in [RSVP-TE] for merging of RESV messages, only Path messages whose ERO from that LSR to the egress is the same can be merged. If merging occurs and one of the Path messages merged was for the protected LSP, then the final Path message to be sent MUST be that of the protected LSP. This merges the backup LSPs into the protected LSP at that LSR. Once the final Path message has been identified, the MP MUST start to refresh it downstream periodically.

If merging occurs and all the Path messages were for backup LSPs, then the DETOUR object, if any, should be altered as specified in Section 8.1

7.1.2. Merging Detours using the Path-Specific Method

An LSR (that is, an MP) may receive multiple Path messages from different interfaces with identical SESSION and SENDER_TEMPLATE objects. In this case, Path state merging is REQUIRED. The merging rule is as follows:

If all Path messages have neither a FAST_REROUTE nor a DETOUR object, or if the MP is the egress of the LSP, no merging is required. The messages are processed according to [RSVP-TE].

Otherwise, the MP MUST record the Path state and the incoming interface. If the Path messages do not share an outgoing interface and a next-hop LSR, the MP MUST consider them to be independent LSPs and MUST NOT merge them.


For all the Path messages that share the same outgoing interface and next-hop LSR, the MP runs the following procedure to create a Path message to forward downstream.

1. If one or more of the Path messages is for the protected LSP (a protected LSP is one originated from this node, or with the FAST_REROUTE object, or without the DETOUR object), one of these must become the chosen Path message. There could be more than one; in that case, which one to forward is a local decision. Quit.
2. From the remaining set of Detour Path messages, eliminate from consideration those that traverse nodes that others want to avoid.
3. If several still remain, which one to forward is a local decision. If none remain, then the MP MAY try to find a new route that avoids all nodes that merging Detour Paths want to avoid; it will forward a Path message with that ERO.

Once the final Path message has been identified, the MP MUST start to refresh it downstream periodically. Other LSPs are considered merged at this node. For bandwidth reservations on the outgoing link, any merging should be considered to have occurred before bandwidth is reserved. Thus, even though Fixed Filter style is specified, multiple detours and/or their protected LSP (which are to be merged due to sharing an outgoing interface and next-hop LSR) will reserve only the bandwidth of the final Path message on that outgoing interface.

If no merged Path message can be constructed, the MP SHOULD send a PathErr in response to the most recently received detour Path message. If a protected Path is chosen to be forwarded but it traverses nodes that some detours want to avoid, PathErrs SHOULD be sent in response to those detour Paths which cannot merge.

7.1.2.1. An Example of Path Message Merging



Example 4. Path Message Merging

In Example 4, R8 will receive Path messages that have the same SESSION and SENDER_TEMPLATE from detours for R2 and R3. During merging at R8, because detour R3 has a shorter ERO path length (that is, ERO is [R9->R5->R6], and path length is 3), R8 will select it as the final LSP and will only propagate its Path messages downstream. Upon receiving a Resv (or a ResvTear) message, R8 must relay the messages toward both R2 and R3.

R5 has to merge as well, and it will select the main LSP, since it has the FAST_REROUTE object. Thus, the detour LSP terminates at R5.

7.1.3. Message Handling for Merged Detours

When an LSR receives a ResvTear for an LSP, the LSR must determine whether it has an alternate associated LSP. For instance, if the ResvTear was received for a protected LSP but an associated backup LSP has not received a ResvTear, then the LSR has an alternate associated LSP. If the LSR does not have an alternate associated LSP, then the MP MUST propagate the ResvTear toward the LSP's ingress, and, for each backup LSP merged into that LSP at this LSR, the ResvTear SHOULD also be propagated along the backup LSP.

The MP may receive PathTear messages for some of the merging LSPs. PathTear messages SHOULD NOT be propagated downstream until the MP has received PathTear messages for each of the merged LSPs. However, the fact that one or more of the merged LSPs has been torn down should be reflected in the downstream message, such as by changing the DETOUR object, if there is one.

7.2. Handling Failures

When a downstream LSR detects a local link failure, for any protected LSPs routed over the failed link, Path and Resv state MUST NOT be cleared, and PathTear and ResvErr messages MUST NOT be sent immediately. If this is not the case, then the facility backup method will not work. Furthermore, a downstream LSR SHOULD reset the

refresh timers for these LSPs as if they had just been refreshed. This is to allow time for the PLR to begin refreshing state via the bypass tunnel. State MUST be removed if it has not been refreshed before the refresh timer expires. This allows the facility backup method to work without requiring that it signal backup paths through the bypass tunnel before failure.

After a failure has occurred, the MP must still send Resv messages for the backup LSPs associated with the protected LSPs that have failed. If the backup LSP was sent through a bypass tunnel, then the PHOP object in its Path message will have the IP address of the associated PLR. This will ensure that Resv state is refreshed.

Once the local link has recovered, the MP may or may not accept Path messages for existing protected LSPs that had failed over to their backup.

8. Behavior of All LSRs

The objects and methods defined in this document require behavior from all LSRs in the traffic-engineered network, even if an LSR is not along the path of a protected LSP.

First, if a DETOUR object is included in the backup LSP's path message for the sender template-specific method, the LSRs in the traffic-engineered network should support the DETOUR object.

Second, if the path-specific method is to be supported for the one-to-one backup method, it is necessary that the LSRs in the traffic-engineered network be capable of merging detours as specified in Section 8.1.

It is possible to avoid specific LSRs that do not support this behavior by assigning a link attribute to all the links of those LSPs and then requesting that backup paths exclude this link attribute.

8.1. Merging Detours in the Path-Specific Method

If multiple Path Messages for different detours are received with the same SESSION, SENDER_TEMPLATE, outgoing interface, and next-hop LSR, then the LSR must function as a Detour Merge Point and merge the detour Path Messages. This merging should occur as specified in Section 7.1.2 and shown in Example 4.

In addition, it is necessary to update the DETOUR object to reflect the merging that has taken place. This is done using the following algorithm to format the outgoing DETOUR object for the final LSP:

- Combine all the (PLR_ID, Avoid_Node_ID) pairs from all the DETOUR objects of all merged LSPs into a new object. Ordering is insignificant.

9. Security Considerations

This document does not introduce new security issues. The security considerations pertaining to the original RSVP protocol [RSVP] remain relevant.

Note that the facility backup method requires that a PLR and its selected merge point trust RSVP messages received from each other.

10. IANA Considerations

IANA [RFC-IANA] has assigned the following RSVP Class Numbers for objects defined in this document.

10.1. DETOUR Object

IANA has assigned:

63 DETOUR

Class Types or C-Types:

7	IPv4
8	IPv6

Future C-Types will be assigned using the following guidelines:

C-Types 0 through 127 are assigned by Standards Action.

C-Types 128 through 191 are assigned by Expert Review.

C-Types 192 through 255 are reserved for Vendor Private Use.

For C-Types in the range 192 through 255, the first four octets of the DETOUR object after the C-Type must be the Vendor's SMI Network Management Private Enterprise Code (see [ENT]) in network byte order.

10.2. FAST_REROUTE Object

IANA has assigned:

205 FAST_REROUTE

Class Types or C-Types:

1	FAST_REROUTE Type 1
7	RESERVED

In the FAST_REROUTE object, C-Type 7 is reserved as it is still used by pre-standard implementations. Future C-Types will be assigned using the following guidelines:

C-Types 0 through 127 are assigned by Standards Action.

C-Types 128 through 191 are assigned by Expert Review.

C-Types 192 through 255 are reserved for Vendor Private Use.

For C-Types in the range 192 through 255, the first four octets of the FAST_REROUTE object after the C-Type must be the Vendor's SMI Network Management Private Enterprise Code (see [ENT]) in network byte order.

11. Contributors

This document was written by George Swallow, Ping Pan, Alia Atlas, Jean Philippe Vasseur, Markus Jork, Der-Hwa Gan, and Dave Cooper.

Jean Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719
USA

Phone: +1 978 497 6238
EMail: jpv@cisco.com

Markus Jork
Quarry Technologies
8 New England Executive Park
Burlington, MA 01803
USA

Phone: +1 781 359 5071
EMail: mjork@quarrytech.com

Der-Hwa Gan
Juniper Networks
1194 N.Mathilda Ave
Sunnyvale, CA 94089
USA

Phone: +1 408 745 2074
EMail: dhg@juniper.net

Dave Cooper
Global Crossing
960 Hamlin Court
Sunnyvale, CA 94089
USA

Phone: +1 916 415 0437
EMail: dcooper@gblix.net

12. Acknowledgments

We would like to acknowledge input and helpful comments from Rob Goguen, Tony Li, Yakov Rekhter and Curtis Villamizar. Especially, we thank those, who have been involved in interoperability testing and field trails, and provided invaluable ideas and suggestions. They are Rob Goguen, Carol Iturralde, Brook Bailey, Safaa Hasan, Richard Southern, and Bijan Jabbari.

13. Normative References

- [RSVP] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RSVP-TE] 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.
- [RFC-WORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC-IANA] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [ENT] IANA PRIVATE ENTERPRISE NUMBERS,
<http://www.iana.org/assignments/enterprise-numbers>

Authors' Addresses

George Swallow
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719
USA

Phone: +1 978 244 8143
EMail: swallow@cisco.com

Ping Pan
Hammerhead Systems
640 Clyde Court
Mountain View, CA 94043
USA

EMail: ppan@hammerheadsystems.com

Alia Atlas
Avici Systems
101 Billerica Avenue
N. Billerica, MA 01862
USA

Phone: +1 978 964 2070
EMail: aatlas@avici.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

