Network Working Group
Internet Draft
Intended status: Standards Track

Expires: April 27, 2015

Chandra Ramachandran (Ed)
Yakov Rekhter (Ed)
Juniper Networks

October 27, 2014

Refresh Interval Independent FRR Facility Protection
draft-chandra-mpls-enhanced-frr-bypass-00

Status of this Memo

Copyright Notice

Abstract

This document defines RSVP-TE extensions to facilitate refresh-interval independent FRR facility protection.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Table of Contents

1. Introduction

The facility backup protection mechanism is one of two methods discussed in [RFC4090] for enabling the fast reroute of traffic onto backup LSP tunnels in 10s of milliseconds, in the event of a failure. This document discusses a few shortcomings with some of the refresh-interval reliant procedures proposed for this method in [RFC4090]. These shortcomings come to the fore under scaled conditions and get highlighted even further when large RSVP refresh intervals are used. The RSVP-TE extensions defined in this document will enhance the facility backup protection mechanism by making the corresponding procedures refresh-interval independent.

2. Motivation

The primary bottleneck that needs to be overcome in order to scale RSVP-TE implementation to establish and maintain in the order of multiple 100K Label Switched Paths (LSPs) is the rate of RSVP protocol messages that would be required to handle the scale of LSPs. RSVP protocol message rate is influenced by both triggered and periodic messages. The facility protection mechanism is the FRR method of choice in scaled scenarios. The timely establishment of backup LSP after failure is critical to keep the LSP state refreshed on routers downstream of the failure. It should be noted that while timely establishment of backup LSPs after failure is a problem on its own, the requirement of RSVP protocol to periodically refresh existing LSP states exacerbates the problem.

One common and straightforward mechanism to mitigate the RSVP message rate problem is to increase the refresh interval of LSP states so that the routers may prioritize backup LSP establishment and other triggered messages. If large refresh time can be complemented with RSVP refresh reduction extensions defined in [RFC2961], then RSVP-TE implementations can use these extensions to avoid rapid retransmits to reliably convey any new state or state change to neighboring router and avoid re-sending the entire message during refresh to neighboring router. Even though the combination of large refresh time and reliable message delivery could be a potential solution, there are some shortcomings if this combination is applied to facility protection specified in [RFC4090].

3. Problem Description

In the topology illustrated in Figure 1, consider a large number of LSPs from A to D transiting B and C. Assume that refresh interval has been configured to be large of the order of minutes and refresh reduction extensions are enabled on all routers.

```
                                [E]
                               /   \
                              /     \
                             /       \
                            /         \
                           /           \
                          /             \
                        [A]-----[B]-----[C]-----[D]
                                  \             /
                                   \           /
                                    \         /
                                     \       /
                                      \     /
                                       \   /
                                        [F]
```

Figure 1: Example Topology

Also assume that node protection has been configured for the LSPs and the LSPs are protected by each router in the following way

- A has made node protection available using bypass LSP A -> E -> C; A is the Point of Local Repair (PLR) and C is Node Protecting Merge Point (NP-MP)

- B has made node protection available using bypass LSP B -> F -> D; B is the PLR and D is the NP-MP

- C has made link protection available using bypass LSP C -> B -> F -> D; C is the PLR and D is the LP-MP

In the above condition, assume that B-C link fails. The following is the sequence of events that is expected to occur for all protected LSPs under normal conditions.

1. B performs local repair and re-directs LSP traffic over the bypass LSP B -> F -> D.
2. B also creates backup state for the LSP and triggers sending of backup LSP state to D over the bypass LSP B -> F -> D.
3. D receives backup LSP states and merges the backups with the protected LSPs.
4. As the link on C over which the LSP states are refreshed has failed, C will no longer receive state refreshes. Consequently the protected LSP states on C will time out and C will send tear down message for all LSPs.

While the above sequence of events has been described in [RFC4090], there are a few problems for which no mechanism has been specified explicitly.

- If the protected LSP on C times out before D receives signaling for the backup LSP, then D would receive PathTear from C prior to receiving signaling for the backup LSP, thus resulting in deleting the LSP state. This would be possible at scale even with default refresh time.

- If upon the link failure C is to keep state until its timeout, then with long refresh interval this may result in a large amount of stale state on C. Alternatively, if upon the link failure C is to delete the state and send PathTear to D, this would result in deleting the state on D, thus deleting the LSP. D needs a reliable mechanism to determine whether it is MP or not to overcome this problem.

- If head-end A attempts to tear down LSP after step 1 but before step 2 of the above sequence, then B may receive the tear down message before step 2 and delete the LSP state from its state database. If B deletes its state without informing D, with long refresh interval this could cause (large) buildup of stale state on D.

- If B fails to perform local repair in step 1, then B will delete the LSP state from its state database without informing D. As B deletes its state without informing D, with long refresh interval this could cause (large) buildup of stale state on D.

The purpose of this document is to provide solutions to the above problems which will then make it practical to scale up to a large number of protected LSPs in the network.

4. Solution Aspects

The solution consists of five parts.

- Enhance facility protection method defined in [RFC4090] by introducing MP determination mechanism that enables PLR to signal availability of link or node protection to the MP. See section 4.1 for more details.

- Handle upstream link or node failures by cleaning up LSP states if the node has not found itself as MP through MP determination mechanism. See section 4.2 for more details.

- Introduce extensions to enable a router to send tear down message to downstream router that enables the receiving router to conditionally delete its local state. See section 4.3 for more details.

- Enhance facility protection by allowing a PLR to directly send tear down message to MP without requiring the PLR to either have a working bypass LSP or have already refreshed backup LSP state. See section 4.4 for more details.

- Introduce extensions to enable the above procedures to be backward compatible with routers along the LSP path running implementation that do not support these procedures. See section 4.5 for more details.

4.1. Signaling Protection availability for MP determination

4.1.1. PLR Behavior

When protected LSP comes up and if "local protection desired" is set in SESSION_ATTRIBUTE object, each node along the LSP path attempts to make local protection available for the LSP.

- If "node protection desired" flag is set, then the node tries to become a PLR by attempting to create NP-bypass LSP to NNhop node avoiding the Nhop node on protected LSP path. In case node protection could not be made available after some time out, the node attempts to create a LP-bypass LSP to Nhop node avoiding only the link that protected LSP takes to reach Nhop

- If "node protection desired" flag is not set, then the PLR attempts to create a LP-bypass LSP to Nhop node avoiding the link that protected LSP takes to reach Nhop

While selecting destination address of the bypass LSP, the PLR should attempt to select the router ID of the NNhop or Nhop node. If PLR and MP are in same area, then the PLR may utilize TED to determine the router ID from the interface address in RRO (if NodeID is not included in RRO). If the PLR and MP are in different IGP areas, then the PLR should use the NodeID address of NNhop MP if included in the RRO of RESV. If the NP-MP in different area has not included NodeID in RRO, then the PLR should use NP-MP’s interface address present in the RRO. The PLR should use its router ID as the source address of the bypass LSP. The PLR should also include its router ID as NodeID in PATH RRO unless configured explicitly not to include NodeID. In parallel to the attempt made to create NP-bypass or LP-bypass, the PLR initiates remote Hello to the NNhop or Nhop node respectively to track the reachability of NP-MP or LP-MP after any failure.

- If NP-bypass LSP comes up, then the PLR sets "local protection available" and "NP available" RRO flags and triggers PATH to be sent.

- If LP-bypass LSP comes up, then the PLR sets "local protection available" RRO flag and triggers PATH to be sent.

- After signaling protection availability, if the PLR finds that the protection becomes unavailable then it should attempt to make protection available. The PLR should wait for a time out before resetting RRO flags relating to protection availability and triggering PATH downstream. On the other hand, the PLR need not wait for time out to set RRO flags relating to protection availability and immediately trigger PATH downstream.

4.1.2. Remote Signaling Adjacency

A NodeID based signaling adjacency is one in which NodeID is used in source and destination address fields in RSVP Hello. [RFC4558] formalizes NodeID based Hello messages between two neighboring routers. The new procedures defined in the previous section extends the applicability of NodeID based Hello messages between two routers that may not have an interface connecting them for exchange of RSVP messages.

4.1.3. PATH RRO flags Propagation

As each node along the LSP path can make protection available, propagating PATH immediately due to change in RRO flags on any upstream node would increase control plane message load. So whenever a node receives PATH, it should check if the only change is in RRO flags. If the change is only in PATH RRO flags, then the node should decide whether to propagate the PATH based on the following rule.

- If "NP desired" flag is set and "NP available" flag has changed in Phop’s RRO flags, then PATH is triggered.

- In all other cases the change is not propagated.

4.1.4. MP Behavior

When the NNhop or Nhop node receives the triggered PATH with RRO flag(s) set, the node should check the presence of remote signaling adjacency with PLR (this check is needed to detect network being partitioned). If the flags are set and the signaling adjacency is present, the node concludes that protection has been made available at the PLR. If the PLR has included NodeID in PATH RRO, then that NodeID is the remote neighbor address. Otherwise, the PLR’s interface address in RRO will be remote neighbor address. If "NP available" flag is set by PPhop node, then it is NP-MP. Otherwise, it concludes it is LP-MP.

Once a node concludes it is MP, it should consider a "remote" state having been created from an implicit refresh directly from PLR. The "remote" state is identical to the protected LSP state except for the difference in HOP object that contains the address of remote neighbor address of node signaling adjacency with PLR. The procedures relating to "remote" state are explained in Section "Remote State Teardown". The MP should consider the "remote" state automatically deleted if:

- NP-MP receives PATH later with "NP available" flag reset in PLR’s RRO flags, or

- LP-MP receives PATH later with "local protection available" flag reset in PLR’s RRO flags, or

- Node signaling adjacency with PLR goes down, or

- MP receives backup LSP signaling from PLR overriding the shadow state, or

- MP receives PathTear, or

- MP deletes the LSP state

4.2. Impact of Failures on LSP State

4.2.1. Non-MP Behavior on Phop Link/Node Failure

When a node detects Phop link or Phop node failure and the node is not an MP, then it should send Conditional PathTear (refer to Section "Conditional PathTear" below) and delete LSP state.

4.2.2. LP-MP Behavior on Phop Link Failure

When the link to PLR fails, the link signaling adjacency to PLR will fail whereas the node signaling adjacency to PLR will remain up. So the MP should retain state.

4.2.3. LP-MP Behavior on Phop Node Failure

When the node signaling adjacency with Phop (that is also the PLR) goes down, the node should send normal PathTear and delete the LSP state.

4.2.4. NP-MP Behavior on Phop Link Failure

If the Phop link fails on NP-MP, then NP-MP should start a one shot timer (called "NodeFailureCheck" hereafter) with period greater than the hold time of NodeID neighbor session with Phop node. The purpose of "NodeFailureCheck" timer is to detect whether Phop link fails but the Phop node does not. This timer would expire or time out if the node signaling adjacency timer with Phop does not expire. If the node signaling adjacency hold time expires prior to the new timer, then the node should retain LSP state and delete the new timer. If the "NodeFailureCheck" timer expires, then the node should send Conditional PathTear and delete LSP state.

In the example topology in Figure 1, assume both A has made node protection available and C has concluded it is NP-MP. When B-C link fails then C should delete LSP state and send Conditional PathTear to D. If B has made node protection available and D has concluded it is NP-MP, then D would not delete LSP state on receiving Conditional PathTear from C. On the other hand, if D has not concluded it is NP-MP, then D would delete LSP state.

4.2.5. NP-MP Behavior on Phop Node Failure

When the Phop node fails, the node signaling adjacency with Phop will fail whereas the remote signaling adjacency to PLR will remain up. So the MP should retain state till refresh timeout.

4.2.6. NP-MP Behavior on PLR Link Failure

If the PLR link that is not attached to NP-MP fails and NP-MP receives Conditional PathTear from the Phop node, then the MP should retain state as long as the remote signaling adjacency with PLR is up. This is because the Conditional PathTear from the Phop node will not impact the "remote" state from the PLR. Note that Phop node would send Conditional PathTear if it was not an MP.

In the above example, assume C & D are NP-MP for PLRs A & B respectively. Now when A-B link fails, as B is not MP and its Phop link signaling adjacency has failed, B should delete LSP state (this behavior is required for unprotected LSPs). In the data plane, that would require B delete the label forwarding entry corresponding to the LSP. So if B’s downstream nodes C and D continue to retain state, it would not be correct for D to continue to assume itself as NP-MP for PLR B.

- As B had previously signaled NP availability, one possible solution would be to let B signal lack of NP availability before sending Conditional PathTear to C. B may trigger PATH, wait for ACK and then send Conditional PathTear to C, but this solution would increase control message load
- Or B may include both PATH with updated RRO flags and Conditional PathTear in a message bundle. While this solution would reduce control message load, the assumption that RSVP protocol could ensure two messages bundled in same message may not hold always.
- Alternatively, B may just send Conditional PathTear to C and let C interpret Conditional PathTear as implicit signaling of lack of NP availability. C should then update B’s RRO flags to signal D that node protection is longer available on B. This is the option that does not make any assumption on implementation and also not increase control message load.

The mechanism to accomplish PATH RRO update is given below.

1. B should send Conditional PathTear to C and delete LSP state.

2. When C receives Conditional PathTear, it should decide to retain LSP state as it is NP-MP of PLR A. C also should check whether Phop B had previously signaled availability of node protection. As B had previously signaled NP availability in its PATH RRO flags, C should reset "local protection available" and "NP available" on RRO flags corresponding to B and trigger PATH to D.
3. When D receives triggered PATH, it realizes that it is no longer NP-MP and so deletes the "remote" state. D does not propagate PATH further down because the only change is in PATH RRO flags of B.

4.2.7. Phop Link Failure on Node that is LP-MP and NP-MP

A node may be both LP-MP as well as NP-MP at the same time for Phop and PPhop nodes respectively. If Phop link fails on such node, the node should retain state because its Phop has made link protection available. In this scenario, "NodeFailureCheck" timer should not be started because the node would retain state irrespective of whether Phop node would fail subsequently or not.

4.2.8. Phop Node Failure on Node that is LP-MP and NP-MP

If a node that is both LP-MP and NP-MP detects Phop node failure, then the node should retain state till refresh timeout.

4.3. Conditional Path Tear

In the example provided in the previous section "NP-MP Behavior on PLR link failure", B deletes LSP state once B detects its link to Phop went down as B is not MP. If B were to send PathTear normally, then C would delete LSP state immediately. In order to avoid this, there should be some mechanism by which B could indicate to C that B does not require the receiving node to unconditionally delete the LSP state immediately. For this, B should add a new optional object in PathTear. If node C also understands the new object, then C should delete LSP state only if it is not an NP-MP - in other words C should delete LSP state if there is no "remote" PLR state on C.

4.3.1. Sending Conditional Path Tear

A node should send Conditional PathTear if the node decides to delete the LSP state under the following conditions.

- Ingress has requested node protection for the LSP, and

- PathTear is not received from upstream node, and

- A node is not a MP and Phop link or Phop node signaling adjacency goes down, or a node is an NP-MP and "NodeFailureCheck" timer started after Phop link down expires.

It should be noted that a node sends Conditional PathTear upon deleting its state in order for its Nhop node to retain state if it is NP-MP.

4.3.2. Processing Conditional Path Tear

When a node that is not an NP-MP receives Conditional PathTear, the node should delete LSP state, and process Conditional PathTear by considering it as normal PathTear. Specifically, the node should not propagate Conditional PathTear downstream but remove the optional object and send normal PathTear downstream.

When a node that is an NP-MP receives Conditional PathTear, it should not delete LSP state. The node should check whether the Phop node previously set "NP available" flag in PATH RRO flags. If the flag had been set previously by Phop, then the node should clear "local protection available" and "NP available" flags in Phop’s RRO flags and trigger PATH downstream.

If Conditional PathTear is received from a neighbor that has not advertised support (refer to Section 4.5) for the new procedures defined in this document, then the node should consider the message as normal PathTear. The node should propagate normal PathTear downstream and delete LSP state.

4.3.3. CONDITIONS object

As any implementation that does not support Conditional PathTear should ignore the new object but process the message as normal PathTear without generating any error, the Class-Num of the new object should be 10bbbbbb where ’b’ represents a bit (from Section 3.10 of [RFC2205]).

The new object is called as "CONDITIONS" object that will specify the conditions under which default processing rules of the RSVP message should be invoked.

The object has the following format:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          Length               |  Class        |     C-type    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Reserved                          |M|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Length

This contains the size of the object in bytes and should be set to eight.

Class

TBD

C-type

1

M bit

This bit indicates that the message should be processed based on the condition whether the receiving node is Merge Point or not.

4.4. Remote State Teardown

As the refresh timeout of LSP state may be high, it is essential that LSP state be cleaned up properly even after local repair. If the Ingress intends to tear down the LSP or if PLR is unable to perform local repair, it would not be desirable to wait till backup LSP signaling to perform state cleanup. To enable LSP state cleanup when LSP is being locally repaired, nodes should send "remote" tear down message instructing the receiving node to delete LSP state.

Consider node C in above example topology (Figure 1) has gone down and B has not signaled backup LSP to D. If Ingress A intends to tear down the LSP, then the following text describes the mechanism to clean up LSP state on all nodes along the path of the LSP.

1. Ingress A sends normal PathTear to B.

2. To enable LSP state cleanup, B should send "remote" PathTear with destination IP address set to that of D, and HOP object containing local address used in remote Hello session with D.
3. On D there would be a remote signaling adjacency with B and so D should accept the remote PathTear and delete LSP state.

4.4.1. PLR Behavior on Local Repair Failure

If local repair fails on the PLR after a failure, then this should be considered as a case for cleaning up LSP state from PLR to the Egress. PLR would achieve this using "remote" PathTear to clean up state from MP. If MP has retained state, then it would propagate PathTear downstream thereby achieving state cleanup. Note that in the case of link protection, the PathTear would be directed to LP-MP node IP address rather than the Nhop interface address.

4.4.2. LSP Preemption during Local Repair

If an LSP is preempted when there is no failure along the path of the LSP, the node on which preemption occurs would send PathErr and ResvTear upstream and only delete the forwarding state. But if the LSP is being locally repaired upstream of the node on which the LSP is preempted, then the node should delete LSP state and send normal PathTear downstream. When PLR signals backup LSP, the node that was formerly MP will respond with PathErr.

4.4.2.1. Preemption after Phop Link failure

If LSP is preempted on LP-MP after its Phop or incoming link has already failed but the backup LSP has not been signaled yet, then the node should send normal PathTear and delete LSP state. As the LP-MP has retained LSP state because the PLR would refresh the LSP through backup LSP signaling, preemption would bring down the LSP and the node would not be LP-MP any more requiring the node to clean up LSP state.

4.4.2.2. Preemption after Phop Node failure

If LSP is preempted on NP-MP after its Phop node has already failed but the backup LSP has not been signaled yet, then the node should send normal PathTear and delete LSP state. As the NP-MP has retained LSP state because the PLR would refresh the LSP through backup LSP signaling, preemption would bring down the LSP and the node would not be NP-MP any more requiring the node to clean up LSP state.

Consider node B goes down on the same example topology (Figure 1). As C is NP-MP for PLR A, C should retain LSP state.

1. The LSP is preempted on C.
2. C would delete its reservation on C-D link. But C cannot send PathErr or ResvTear to PLR A because backup LSP has not been signaled yet.
3. As the only reason for C having retained state after Phop node failure was that it was NP-MP, C should send normal PathTear to D and delete LSP state. D would also delete state on receiving PathTear from C.
4. B starts backup LSP signaling to D. But as D does not have the LSP state, it should reject backup LSP PATH and send PathErr to B.
5. B should delete its reservation and send ResvTear to A.

4.5. Backward Compatibility Procedures

The "Enhanced FRR facility protection" referred below in this section refers to the set of changes that have been proposed in previous sections. Any implementation that does not support them has been termed as "existing implementation". Of the proposed extensions, signaling protection using RRO flags is expected to be backward compatible and can work safely irrespective of whether the refresh time is large. This is because the existing implementations would not send error or tear down message in response to the flags in PATH RRO but would simply ignore and propagate them. On the other hand, changes proposed relating to LSP state cleanup namely Conditional and remote PathTear require support from other nodes along the LSP path. So procedures that fall under LSP state cleanup category should be turned on only if nodes involved i.e. PLR, MP and intermediate node in the case of NP, support the extensions.

4.5.1. Detecting Support for Enhanced FRR Facility Protection

An implementation supporting the FRR facility protection extensions specified in previous sections should set a new flag "Enhanced facility protection" in CAPABILITY object in Hello messages.

- As nodes supporting the extensions should initiate Node Hellos with adjacent nodes, a node on the path of protected LSP can determine whether its Phop or Nhop neighbor supports FRR enhancements from the Hello messages sent by the neighbor.

- If a node attempts to make node protection available, then the PLR should initiate remote node signaling adjacency with NNhop. If the NNhop (a) does not reply to remote node Hello message or (b) does not set "Enhanced facility protection" flag in CAPABILITY object in the reply, then the PLR can conclude that NNhop does not support FRR extensions.

- If node protection is requested for an LSP and if (a) PPhop node has not set "local protection available" and "NP available" flags in its RRO flags or (b) PPhop node has not initiated remote node Hello messages, then the node should conclude that PLR does not support FRR extensions. The details are described in the "Procedures for backward compatibility" section below.

The new flag that will be introduced to CAPABILITY object is specified below.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            Length             | Class-Num(134)|  C-Type  (1)  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Reserved                     |E|T|R|S|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

E bit

Indicates that the sender supports Enhanced FRR facility protection

Any node that sets the new E-bit is set in its CAPABILITY object must also set Refresh-Reduction-Capable bit in common header of all RSVP messages.

4.5.2. Procedures for backward compatibility

The procedures defined hereafter are performed on a subset of LSPs that traverse a node, rather than on all LSPs that traverse a node. This behavior is required to support backward compatibility for a subset of LSPs traversing nodes running existing implementations.

4.5.2.1. Lack of support on Downstream Node

- If the Nhop does not support enhanced facility protection FRR, then the node should reduce the "refresh period" in TIME_VALUES object carried in PATH to default small refresh default value.

- If node protection is requested and the NNhop node does not support the enhancements, then the node should reduce the "refresh period" in TIME_VALUES object carried in PATH to small refresh default value.

If the node reduces the refresh time from the above procedures, it should also not send remote PathTear or Conditional PathTear messages.

Consider the example topology in Figure 1. If C does not support scalability improvements, then:

- A and B should reduce the refresh time to default value of 30 seconds and trigger PATH

- If B is not an MP and if Phop link of B fails, B cannot send Conditional PathTear to C but should time out LSP state from A normally. This would be accomplished if A would also reduce the refresh time to default value. So if C does not support enhanced facility protection, then Phop B and PPhop A should reduce refresh time to small default value.

4.5.2.2. Lack of support on Upstream Node

- If Phop node does not support enhanced facility protection, then the node should reduce the "refresh period" in TIME_VALUES object carried in RESV to default small refresh time value.

- If node protection is requested and the Phop node does not support the enhancements, then the node should reduce the "refresh period" in TIME_VALUES object carried in PATH to default value.

- If node protection is requested and PPhop node does not support the enhancements, then the node should reduce the "refresh period" in TIME_VALUES object carried in RESV to default value.

- If the node reduces the refresh time from the above procedures, it should also not execute MP determination procedures.

5. Security Considerations

This document does not introduce new security issues. The security considerations pertaining to the original RSVP protocol [RFC2205] remain relevant.

6. IANA Considerations

TBD

7. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4090] Pan, P., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC2961] Berger, L., "RSVP Refresh Overhead Reduction Extensions", RFC 2961, April 2001.

[RFC3209] Awduche, D., "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC2205] Braden, R., "Resource Reservation Protocol (RSVP)", RFC 2205, September 1997.

[RFC4558] Ali, Z., "Node-ID Based Resource Reservation (RSVP) Hello: A Clarification Statement", RFC 4558, June 2006.

8. Acknowledgments

Thanks to Raveendra Torvi and Yimin Shen for their comments and inputs.

9. Authors' Addresses

Chandra Ramachandran
Juniper Networks
csekar@juniper.net

Yakov Rekhter
Juniper Networks
Email: yakov@juniper.net

Markus Jork
Juniper Networks
Email: mjork@juniper.net

Contributors

Harish Sitaraman
Juniper Networks
Email: hsitaraman@juniper.net

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net