

Internet Engineering Task Force (IETF)
Request for Comments: 5711
Updates: 3209
Category: Standards Track
ISSN: 2070-1721

JP. Vasseur, Ed.
G. Swallow
Cisco Systems, Inc.
I. Minei
Juniper Networks
January 2010

Node Behavior upon Originating and Receiving Resource Reservation
Protocol (RSVP) Path Error Messages

Abstract

The aim of this document is to describe a common practice with regard to the behavior of nodes that send and receive a Resource Reservation Protocol (RSVP) Traffic Engineering (TE) Path Error messages for a preempted Multiprotocol Label Switching (MPLS) or Generalized MPLS (GMPLS) Traffic Engineering Label Switched Path (TE LSP). (For reference to the notion of TE LSP preemption, see RFC 3209.) This document does not define any new protocol extensions.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5711>.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Protocol Behavior	3
2.1. Behavior at Detecting Nodes	4
2.2. Behavior at Receiving Nodes	5
2.3. Data-Plane Behavior	5
3. RSVP PathErr Messages for a Preempted TE LSP	5
4. Security Considerations	5
5. Acknowledgements	6
6. References	6
6.1. Normative References	6
6.2. Informative References	6

1. Introduction

The aim of this document is to describe a common practice with regard to the behavior of a node sending a Resource Reservation Protocol (RSVP) Traffic Engineering (TE) Path Error message and to the behavior of a node receiving an RSVP Path Error message for a preempted Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering Label Switched Path (TE LSP). (For reference to the notion of TE LSP preemption, see [RFC3209]).

[RFC2205] defines two RSVP error message types: PathErr and ResvErr that are generated when an error occurs. Path Error messages (PathErr) are used to report errors and travel upstream toward the head-end of the flow. Resv Error messages (ResvErr) travel downstream toward the tail-end of the flow.

This document describes only PathErr message processing for the specific case of a preempted TE LSP, where the term preemption is defined in [RFC3209].

1.1. Requirements Language

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].

2. Protocol Behavior

PathErr messages are routed hop-by-hop using the path state established when a Path message is routed through the network from the head-end to its tail-end.

As stated in [RFC2205], PathErr messages do not modify the state of any node through which they pass; they are only reported to the head-end of the TE LSP (Traffic Engineering Label Switched Path).

The format of the PathErr message is defined in Section 3. of [RFC2205].

The ERROR_SPEC object includes the IP address of the node that detected the error (Error Node Address), and specifies the error through two fields. The Error Code field encodes the category of the error, for example, Policy Control Failure or Unknown object class. The Error Value field qualifies the error code to indicate the error with more precision. [RFC3209] extends RSVP as defined in [RFC2205] for the management of MPLS-TE LSPs. [RFC3209] specifies several additional conditions that trigger the sending of a RSVP PathErr message for which new error codes and error values have been defined

that extend the list defined in [RFC2205]. The exact circumstances under which a TE LSP is preempted and such PathErr messages are sent are defined in [RFC3209] and will not be repeated here.

Values for the Error Code and Error Value fields defined in [RFC2205], [RFC3209], and other documents are maintained in a registry by the IANA.

The error conditions fall into two categories:

- o Fatal errors represent disruptive conditions for a TE LSP.
- o Non-fatal errors are non-disruptive conditions that have occurred for this TE LSP.

PathErr messages may be used in two circumstances:

- o during TE LSP establishment, and
- o after a TE LSP has been successfully established.

Nodal behavior is dependent on which combination of the four cases listed above applies. The following sections describe the expected behavior at nodes that perform a preemption action for a TE LSP (and therefore report using error PathErr messages), and at nodes that receive PathErr messages. This text is a clarification and restatement of the procedures set out in [RFC3209] and does not define any new behavior.

2.1. Behavior at Detecting Nodes

In the case of fatal errors ("Hard Preemption"; see Section 4.7.3 of [RFC3209]), the detecting node MUST send a PathErr message reporting the error condition, and MUST clear the corresponding Path and Resv (control plane) states. A direct implication is that the data-plane resources of such a TE LSP are also released, thus resulting in traffic disruption. It should be noted, however, that in fatal error cases, the LSP has usually already failed in the data plane, and traffic has already been disrupted. When the error arises during LSP establishment, the implications are different to when it arises on an active LSP since no traffic flows until the LSP has been fully established. In the case of non-fatal errors, the detecting node should send a PathErr message, and must not clear control plane or data plane state.

2.2. Behavior at Receiving Nodes

Nodes that receive PathErr messages are all of the nodes along the path of the TE LSP upstream of the node that detected the error. This includes the head-end node. In accordance with Section 3.7.1 of [RFC2205], a node receiving a PathErr message takes no action upon it, and consequently the node must not clear Path or Resv control-plane or data-plane state. This is true regardless of whether the error condition reported by the PathErr is fatal or non-fatal. RSVP states should only be affected upon receiving a PathTear or ResvTear message, or in the event of a Path or Resv state timeout. Further discussion of the processing of these events is outside the scope of this document.

Note that [RFC3473] defines a Path_State_Removed flag in the ERROR_SPEC object carried on a PathErr message. This field may be set to change the behavior of upstream nodes that receive the PathErr message. When set, the flag indicates that the message sender has removed Path state (and any associated Resv and data-plane state) for the TE LSP. The message receiver should do likewise before forwarding the message, but may retain state and clear the flag before forwarding the message.

2.3. Data-Plane Behavior

Any node clearing either or both the Path or the Resv state of a TE LSP MUST also free up the data-plane resources allocated to the corresponding TE LSP.

3. RSVP PathErr Messages for a Preempted TE LSP

Two Error Codes have been defined to report a preempted TE LSP:

- o As defined in [RFC2750]: Error Code=2: "Policy Control Failure", Error Value=5: "Flow was preempted"
- o As defined in [RFC2205], Error Code=12: "Service preempted"

They are both fatal errors.

4. Security Considerations

This document does not define any new procedures, but clarifies those defined in other documents where security considerations are already specified in [RFC3209] and [RFC3473]. This document does not raise specific security issues beyond those of existing MPLS-TE. By

clarifying the procedures, this document reduces the security risk introduced by non-conformant implementations. See [SEC_FMWK] for further discussion of MPLS security issues.

5. Acknowledgements

The authors would like to thank Carol Iturralde, Ashok Narayanan, Rom Reuther, and Reshad Rahman.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2750] Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

6.2. Informative References

- [SEC_FMWK] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, October 2009.

Authors' Addresses

JP Vasseur (editor)
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

EMail: jpv@cisco.com

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

EMail: swallow@cisco.com

Ina Minei
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
USA

EMail: ina@juniper.net

