Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 9, 2013

E. Crabbe
Google, Inc.
J. Medved
Cisco Systems, Inc.
I. Minei
Juniper Networks, Inc.
R. Varga
Pantheon Technologies SRO

May 8, 2013

Stateful PCE extensions for MPLS-TE LSPs
draft-crabbe-pce-stateful-pce-mpls-te-01

Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients (PCCs) requests.

[I-D.ietf-pce-stateful-pce] describes a set of extensions to PCEP to provide stateful control. This document describes the objects and TLVs to be used with these PCEP extensions to control Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSP) via a stateful PCE.

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 [RFC2119].

Status of this Memo

Table of Contents

1. Introduction

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients (PCCs) requests.

[I-D.ietf-pce-stateful-pce] describes a set of extensions to PCEP to provide stateful control. This document describes the objects and TLVs to be used with these PCEP extensions to control Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSP) via a stateful PCE.

2. Terminology

This document uses the following terms defined in [RFC5440]: PCC, PCE, PCEP Peer.

This document uses the following terms defined in [I-D.ietf-pce-stateful-pce] : Passive Stateful PCE, Active Stateful PCE, Delegation, Delegation Timeout Interval, LSP State Report, LSP Update Request, LSP Priority, LSP State Database, Revocation.

Within this document, when describing PCE-PCE communications, the requesting PCE fills the role of a PCC. This provides a saving in documentation without loss of function.

The message formats in this document are specified using Routing Backus-Naur Format (RBNF) encoding as specified in [RFC5511].

3. MPLS-TE specific descriptors used in PCEP Messages

As defined in [RFC5440], a PCEP message consists of a common header followed by a variable-length body made of a set of objects that can be either mandatory or optional. [I-D.ietf-pce-stateful-pce] describes the messages and objects needed in support of stateful PCE. The following sections contain MPLS-TE specific descriptors used in some of these messages.

3.1. MPLS-TE specific descriptors for the PCRpt Message

The format of the PCRpt message is defined in [I-D.ietf-pce-stateful-pce] as follows, and included here for easy reference:

```
<PCRpt Message> ::= <Common Header>
                    <state-report-list>
```

Where:

```
<state-report-list> ::= <state-report>[<state-report-list>]

<state-report> ::= <LSP>
                   [<path-list>]
```

Where:

```
<path-list>::=<path>[<path-list>]
```

For MPLS-TE LSPs, the path descriptor is defined as follows:

```
<path>::=<ERO><attribute-list>
```

Where:

```
<attribute-list> ::= [<LSPA>]
                     [<BANDWIDTH>]
                     [<RRO>]
                     [<metric-list>]

<metric-list> ::= <METRIC>[<metric-list>]
```

The LSP State Report MAY contain a path descriptor for the primary path and one or more path descriptors for backup paths. A path descriptor MUST contain an ERO object as it was specified by a PCE or an operator. A path descriptor MUST contain the RRO object if a primary or secondary LSP is set up along the path in the network. A path descriptor MAY contain the LSPA, BANDWIDTH, and METRIC objects. The ERO,LSPA, BANDWIDTH, METRIC, and RRO objects are defined in[RFC5440].

3.2. MPLS-TE specific descriptors for the PCUpd Message

A Path Computation LSP Update Request message (also referred to as PCUpd message) is a PCEP message sent by a PCE to a PCC to update attributes of an LSP. A PCUpd message can carry more than one LSP Update Request. The Message-Type field of the PCEP common header for the PCUpd message is set to [TBD].

The format of the PCUpd message is defined in [I-D.ietf-pce-stateful-pce] and included here for easy reference:

```
   <PCUpd Message> ::= <Common Header>
                       <udpate-request-list>
```

Where:

```
   <update-request-list> ::= <update-request>[<update-request-list>]

   <update-request> ::= <LSP>
                        [<path-list>]
```

Where:

```
   <path-list>::=<path>[<path-list>]
```

For MPLS-TE LSPs, the endoding of path descriptor is defined as follows:

```
   <path>::=<ERO><attribute-list>
```

Where:

```
   <path>::=<ERO><attribute-list>
```

Where:

```
   <attribute-list> ::= [<LSPA>]
                        [<BANDWIDTH>]
                        [<metric-list>]

   <metric-list> ::= <METRIC>[<metric-list>]
```

There is one mandatory object that MUST be included within each LSP Update Request in the PCUpd message: the LSP object (see [I-D.ietf-pce-stateful-pce]). If the LSP object is missing, the receiving PCE MUST send a PCErr message with Error-type=6 (Mandatory Object missing) and Error-value=[TBD] (LSP object missing).

The LSP Update Request MUST contain a path descriptor for the primary path, and MAY contain one or more path descriptors for backup paths. A path descriptor MUST contain an ERO object. A path descriptor MAY further contain the BANDWIDTH, IRO, and METRIC objects. The ERO, LSPA, BANDWIDTH, METRIC, and IRO objects are defined in [RFC5440].

Each LSP Update Request results in a separate LSP setup operation at a PCC. An LSP Update Request MUST contain all LSP parameters that a PCC wishes to set for the LSP. A PCC MAY set missing parameters from locally configured defaults. If the LSP specified the Update Request is already up, it will be re-signaled. The PCC will use make-before-break whenever possible in the re-signaling operation.

A PCC MUST respond with an LSP State Report to each LSP Update Request to indicate the resulting state of the LSP in the network. A PCC MAY respond with multiple LSP State Reports to report LSP setup progress of a single LSP.

If the rate of PCUpd messages sent to a PCC for the same target LSP exceeds the rate at which the PCC can signal LSPs into the network, the PCC MAY perform state compression and only re-signal the last modification in its queue.

Note that a PCC MUST process all LSP Update Requests - for example, an LSP Update Request is sent when a PCE returns delegation or puts an LSP into non-operational state. The protocol relies on TCP for message-level flow control.

Note also that it’s up to the PCE to handle inter-LSP dependencies; for example, if ordering of LSP set-ups is required, the PCE has to wait for an LSP State Report for a previous LSP before triggering the LSP setup of a next LSP.

3.3. MPLS-TE specific encoding for the PCReq Message for stateful PCE

A PCC MAY include the LSP object defined in [I-D.ietf-pce-stateful-pce] in the PCReq message if the stateful PCE capability has been negotiated on a PCEP session between the PCC and a PCE. The definition of the PCReq message (see [RFC5440], Section 6.4) is then extended as follows:

```
      <PCReq Message>::= <Common Header>
                         [<svec-list>]
                         <request-list>
```

Where:

```
      <svec-list>::=<SVEC>[<svec-list>]
      <request-list>::=<request>[<request-list>]

      <request>::= <RP>
                   <END-POINTS>
                   [<LSP>]              <--- New Object
                   [<LSPA>]
                   [<BANDWIDTH>]
                   [<metric-list>]
                   [<RRO>[<BANDWIDTH>]]
                   [<IRO>]
                   [<LOAD-BALANCING>]
```

Where:

```
      <metric-list>::=<METRIC>[<metric-list>]
```

3.4. MPLS-TE specific encoding for the PCRep Message for stateful PCE

A PCE MAY include the LSP object defined in [I-D.ietf-pce-stateful-pce] in the PCRep message if the stateful PCE capability has been negotiated on a PCEP session between the PCC and the PCE and the LSP object was included in the corresponding PCReq message from the PCC. The definition of the PCRep message (see [RFC5440], Section 6.5) is then extended as follows

```
      <PCRep Message> ::= <Common Header>
                          <response-list>
```

Where:

```
      <response-list>::=<response>[<response-list>]

      <response>::=<RP>
                  [<LSP>]                <--- New Object
                  [<NO-PATH>]
                  [<attribute-list>]
                  [<path-list>]

      <path-list>::=<path>[<path-list>]

      <path>::= <ERO><attribute-list>
```

Where:

```
    <attribute-list>::=[<LSPA>]
                       [<BANDWIDTH>]
                       [<metric-list>]
                       [<IRO>]

    <metric-list>::=<METRIC>[<metric-list>]
```

4. IANA Considerations

This document requests IANA actions to allocate code points for the protocol elements defined in this document. Values shown here are suggested for use by IANA.

4.1. PCEP-Error Object

This document defines new Error-Type and Error-Value for the following new error conditions:

Error-Type	Meaning
6	Mandatory Object missing
	Error-value=9: ERO Object missing for a path in an LSP Update Request where TE-LSP setup is requested
	Error-value=10: BANDWIDTH Object missing for a path in an LSP Update Request where TE-LSP setup is requested

Error-value=11: LSPA Object missing for a path in an LSP Update Request where TE-LSP setup is requested

5. Security Considerations

The security considerations listed in [I-D.ietf-pce-stateful-pce] apply to this document as well.

6. Acknowledgements

We would like to thank Adrian Farrel, Cyril Margaria and Ramon Casellas for their contributions to this document.

We would like to thank Shane Amante,Julien Meuric, Kohei Shiomoto, Paul Schultz and Raveendra Torvi for their comments and suggestions. Thanks also to Dhruv Dhoddy, Oscar Gonzales de Dios, Tomas Janciga, Stefan Kobza and Kexin Tang for helpful discussions.

7. References

7.1. Normative References

[I-D.ietf-pce-stateful-pce] Crabbe, E., Medved, J., Minei, I., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-03 (work in progress), March 2013.

[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.

[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.

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

[RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang,

"OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.

[RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

7.2. Informative References

[RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3346] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", RFC 3346, August 2002.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394,

December 2008.

[RFC5557] Lee, Y., Le Roux, JL., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", RFC 5557, July 2009.

Authors' Addresses

Edward Crabbe
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: edc@google.com

Jan Medved
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
US

Email: jmedved@cisco.com

Ina Minei
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: ina@juniper.net

Robert Varga
Pantheon Technologies SRO
Mlynske Nivy 56
Bratislava 821 05
Slovakia

Email: robert.varga@pantheon.sk