Network Working Group
Internet Draft
Intended status: Standards Track
Expires: January 4, 2008

S. Sivabalan, Ed.
J. Parker
S. Boutros
Cisco Systems, Inc.

K. Kumaki
KDDI Corporation

July 8, 2007

Diff-Serv Aware Class Type Object for Path Computation Element Communication Protocol
draft-sivabalan-pce-dste-01.txt

Status of this Memo

This Internet-Draft will expire on August 17, 2007.

Copyright Notice

Abstract

This document specifies a CLASSTYPE object to support Diff-Serve Aware Traffic Engineering (DS-TE) where path computation is performed with an aid of Path Computation Element (PCE).

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.

Table of Contents

1. Introduction

The Internet Draft [PCEP-ID] specifies the Path Computation Element communication Protocol (PCEP) for communications between a Path Computation Client(PCC) and a Path Computation Element (PCE), or between two PCEs, in compliance with [RFC4657].

Differentiated Service aware MPLS Traffic Engineering (DS-TE) addresses the fundamental requirement to be able to enforce different bandwidth constraints for different classes of traffic and describes mechanisms to achieve per-class traffic engineering, rather than on an aggregate basis across all classes by enforcing Bandwidth

Constraints (BCs) on different classes. Requirements for DS-TE and the associated protocol extensions are specified in [RFC3564] and [RFC4124] respectively.

As per [RFC4657], PCEP must support traffic class-type as an MPLS TE specific constraint. However, in the present form, PCEP [PCEP-ID] does not have the capability to specify the class-type in the path computation request.

In this document, we define a new PCEP object called CLASSTYPE which carries the class-type of the TE LSP in the path computation request. During path computation, a PCE uses the class-type to identify the bandwidth constraint of the TE-LSP.

2. Terminology

CT: Class type: A set of Traffic Trunks governed by a set of bandwidth constraints. Used for the purpose of link bandwidth allocation, constraint based routing and admission control. A given Traffic Trunk belongs to the same CT on all links.

DS-TE: Diff-Serv Aware Traffic Engineering.

LSR: Label Switching Router.

LSP: Label Switched Path.

PCC: Path Computation Client: any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element: an entity (component, application or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP Peer: an element involved in a PCEP session (i.e. a PCC or the PCE).

TE-Class: A pair consisting of a class-type and a preemption priority allowed for that class type. An LSP transporting a Traffic Trunk from that class type can use that preemption priority as the setup priority, the holding priority, or both.

TE LSP: Traffic Engineering Label Switched Path.

Traffic Trunk: An aggregation of traffic flows of the same class (i.e. treated equivalently from the DS-TE perspective) which is placed inside a TE LSP.

3. CLASSTYPE object

The CLASSTYPE object is optional and is used to specify the class-type of a TE LSP. This object is meaningful only within the path computation request, and is ignored in the path reply message. If the TE LSP for which path is to be computed belongs to Class 0, the path computation request MUST not contain the CLASSTYPE object. This allows backward compatibility with PCE that does not support CLASSTYPE object.

3.1. Object definition

The CLASSTYPE object contains a 32-bit word PCEP common object header defined in [PCEP-ID] followed by another 32-bit word object body as shown in Figure 1.

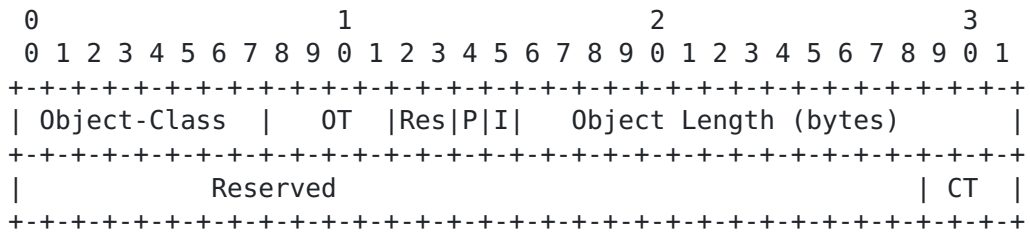

Figure 1 CLASSTYPE object format.

The fields in the object common header are processed as specified in [PCEP-ID]. We explain these fields again for completion. For more details, please refer to [PCEP-ID].

Object-Class is to be assigned by IANA (recommended value=16).

Object-Type (OT) is to be assigned by IANA (recommended value=1).

Res flags (2 bits). Reserved field. This field MUST be set to zero on transmission and MUST be ignored on receipt.

P flag (1 bit): When the P flag is set, the CLASSTYPE object MUST be taken into account by the PCE. Conversely, when the P flag is cleared, the object is optional and the PCE is free to ignore it if not supported.

I flag (1 bit): The PCE MAY include the ignored optional object in its reply and set the I flag to indicate that the optional object was ignored during path computation.

Object Length (16 bits). Specifies the total object length including the header, in bytes. The Object Length field MUST always be a multiple of 4, and at least 4. The maximum object content length is 65528 bytes.

The CLASSTYPE object body contains the following fields:

CT: 3-bit field that indicates the class-type. Values allowed are 1, 2, ... , 7. Value of 0 is Reserved.

Reserved: 29-bit reserved field. It MUST be set to zero on transmission and MUST be ignored on receipt.

3.2. Path Computation Request Message with CLASSTYPE object

The draft [PCEP-ID] specifies the object orders in which objects must be inserted in the PCEP messages. This document specifies that the CLASSTYPE object be inserted after the END-POINT objects as shown below:

The format of a PCReq message is 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>
                [<CLASSTYPE>]
                [<LSPA>]
                [<BANDWIDTH>]
                [<metric-list>]
                [<RRO>]
                [<IRO>]
                [<LOAD-BALANCING>]
where:
<metric-list>::=<METRIC>[<metric-list>]
```

3.3. Handling of the CLASSTYPE object

If the LSP is associated with Class-Type N (1 <= N <= 7), the PCC originating the path computation request MUST include the CLASSTYPE object in the Path computation request message with the Class-Type (CT) field set to N.

If a path computation request contains multiple CLASSTYPE objects, only the first one is meaningful; subsequent CLASSTYPE object(s) MUST be ignored and MUST NOT be forwarded.

If the CLASSTYPE object is not present in the path computation request message, the LSR MUST associate the Class-Type 0 to the LSP.

Path computation reply message MUST NOT include a CLASSTYPE object. If a PCE needs to forward a path computation request containing the CLASSTYPE object to another PCE, it MUST store the class-type of the TE LSP in order to complete the path computation when the path computation reply arrives.

A PCE receiving a path computation request message with the CLASSTYPE object with P flag set that does not recognize the CLASSTYPE object MUST reject the entire PCEP message and MUST send a PCE error message with Error-Type="Unknown Object" or "Not supported Object" defined in [PCEP-ID].

A PCE receiving a path computation request message with the CLASSTYPE object that recognizes the CLASSTYPE object, but does not support the particular Class-Type, MUST send a PCE error message towards the sender with the error type "Diff-Serv aware TE Error" and an error value of "Unsupported Class-Type" (new error code provided below).

A PCE receiving a path computation request message with the CLASSTYPE object that recognizes the CLASSTYPE object, but determines that the Class-Type value is not valid (i.e., Class Type value 0), MUST send a PCE error towards the sender with the error type "Diff-Serve aware TE Error" and an error value of "Invalid Class-Type value" (new error code provided below).

3.4. Determination of Traffic Engineering Class (TE-Class)

As specified in RFC4124, a CT and a Preemption priority map to a Traffic Engineering Class (TE-Class), and there can be up to 8 TE-classes. The TE-class value is used to determine the unreserved bandwidth on the links during path computation. In the case of a PCE, the CT value carried in the CLASSTYPE object and the setup priority in the LSP Attribute (LSPA) object are used to determine the TE-class corresponding to the path computation request. If LSPA object is absent, the setup priority is assumed to be 0.

3.5. Significance of Class-type and TE-Class

To ensure coherent DS-TE operation, a PCE and a PCC should have a common understanding of a particular DS-TE classtype and TE-Class.

If a path computation request crosses an AS boundary, these should have global significance in all domains. Enforcement of this global significance is outside the scope of this document.

3.6. Error Codes for CLASSTYPE object

This document defines the following error type and values:

Error-Type	Meaning
11	Diff-Serve aware TE Error Error-value=1: unsupported class-type. Error-value=2: invalid class-type. Error-value=3: class-type and setup priority does not form a configured TE class.

4. Security Considerations

This document does not introduce new security issues. The security considerations pertaining to PCEP [PCEP-ID] remain relevant.

5. IANA Considerations

IANA assigns value to PCEP parameters. Each PCEP object has an Object-Class and an Object-Type. For the CLASSTYPE object, the suggested values for Object-Class and Object-Type are 16 and 1 respectively.

6. Acknowledgments

The authors would like to thank Jean Philippe Vasseur for his valuable comments.

7. References

7.1. Normative References

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

[RFC4124]Le Faucheur, F. and W. Lai, "Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering", RFC 4124, June 2005.

[PCEP-ID] Path Computation Element (PCE) communication Protocol (PCEP)", draft-ietf-pce-pcep-07.txt (work in progress), February 2007.

7.2. Informative References

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

[RFC3564] Le Faucheur, F. and W. Lai, "Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering", RFC 3564, July 2003.

Author's Addresses

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada

Email: msiva@cisco.com

Jon Parker
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada

Email: jdparker@cisco.com

Sami Boutros
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada

Email: sboutros@cisco.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower Iidabashi, Chiyoda-ku
Tokyo, 102-8460
Japan

Email: ke-kumaki@kddi.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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