

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: July 2, 2017

H. Chen
Huawei Technologies
X. Liu
Ericsson
M. Toy
Verizon
V. Liu
China Mobile
L. Liu
Fujitsu
K. Pithewan
Infinera
December 29, 2016

**Extensions to PCEP for Temporal LSP
draft-chen-tts-05.txt**

Abstract

For existing MPLS LSP tunnel services, it is hard for LSP tunnels to be booked in advance. In addition, once an LSP tunnel is set up, it is assumed to consume a certain amount of resources such as link bandwidth forever.

Temporal LSP tunnel services (TTS) provides an easy way for us to book temporal LSP tunnels in advance. More importantly, a temporal LSP is an LSP with one or more time intervals and it is assumed to consume the resources and carry traffic only in these time intervals.

This document specifies extensions to PCEP for computing a path for a temporal LSP, initiating and maintaining a temporal LSP with a sequence of time intervals, during each of which the LSP carries traffic.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference

material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 2, 2017.

Copyright Notice

Copyright (c) 2016 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	4
2.	Terminology	4
3.	Conventions Used in This Document	5
4.	Operations Overview	5
4.1.	Simple Time Interval	5
4.2.	Recurrent Time Interval	5
4.3.	Elastic Time Interval	6
4.4.	Changes to Time Interval	6
4.5.	Graceful Periods	7
5.	Extensions to PCEP	8
5.1.	Capability TLV in Existing PCE Discovery Protocol	8
5.2.	Open Message Extension	10
5.3.	RP Object Extension	10
5.4.	TIME INTERVAL Object	11
5.4.1.	Absolute Time Interval TLV	11
5.4.2.	Relative Time Interval TLV	13
5.4.3.	Recurrent Absolute Time Interval TLV	14
5.4.4.	Recurrent Relative Time Interval TLV	16
5.5.	Messages for Temporal LSP	17
5.5.1.	Messages between PCE and PCC on Ingress	18
5.5.2.	Messages between two PCEs	19
6.	Procedures	20
6.1.	Creating a Temporal LSP	21
6.2.	Deleting a Temporal LSP	21
7.	Considerations on TED	22
7.1.	TE Representation in Absolute Time	23
7.2.	TE Representation in Relative Time	23
8.	Security Considerations	24
9.	IANA Considerations	24
10.	Acknowledgement	24
11.	References	24
11.1.	Normative References	24
11.2.	Informative References	25
	Authors' Addresses	25

1. Introduction

Once an existing multiprotocol label switching (MPLS) traffic engineering (TE) label switched path (LSP) is set up, it is assumed to carry traffic forever until it is down. When an MPLS TE LSP tunnel is up, it is assumed that the LSP consumes its reserved network resources forever even though the LSP may only use network resources during some period of time. As a result, the network resources are not used efficiently. Moreover, a tunnel service can not be reserved or booked in advance for a period of time or a sequence of time periods.

Temporal LSP tunnel services (TTS) provides an easy way for us to book temporal LSP tunnels in advance. More importantly, a temporal LSP is an LSP with one or more time intervals and it is assumed to consume the resources and carry traffic only in each of these time intervals.

This document specifies extensions to PCEP for computing a path for a temporal LSP, initiating and maintaining a temporal LSP with a period of time called a time interval or a sequence of time intervals. It is assumed that the LSP carries traffic during this time interval or each of these time intervals. Thus the network resources are efficiently used. More importantly, some new services can be provided. For example, a consumer can book a tunnel service in advance for a given time interval or a sequence of time intervals. Tunnel services may be scheduled.

2. Terminology

A Time Interval: a time period from time T_a to time T_b .

LSP: Label Switched Path. An LSP is a P2P (point-to-point) LSP or a P2MP (point-to-multipoint) LSP.

LSP with a time interval: LSP that carries traffic in the time interval.

LSP with a sequence of time intervals: LSP that carries traffic in each of the time intervals.

Temporal LSP: LSP with a time interval or LSP with a sequence of time intervals.

TED: Traffic Engineering Database.

CSPF: Constrained Shortest Path First.

LER: Label Edge Router.

This document uses terminologies defined in [RFC5440](#).

3. 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 [RFC2119](#).

4. Operations Overview

This section briefly describes some operations on a temporal LSP.

4.1. Simple Time Interval

For a temporal LSP, a user configures it with a time interval or a sequence of time intervals. A simple time interval is a time period from time T_a to time T_b , which may be represented as $[T_a, T_b]$.

When an LSP is configured with time interval $[T_a, T_b]$, a path satisfying the constraints for the LSP in the time interval is computed and the LSP along the path is set up to carry traffic from time T_a to time T_b .

In addition to simple time intervals, there are recurrent time intervals and elastic time intervals. Sometimes a simple time interval is called a time interval.

4.2. Recurrent Time Interval

A recurrent time interval represents a series of repeated simple time intervals. It has a simple time interval such as $[T_a, T_b]$, a number of repeats such as 10 (repeats 10 times), and a repeat cycle/time such as a week (repeats every week). The recurrent time interval: " $[T_a, T_b]$ repeats n times with repeat cycle C " represents $n+1$ simple time intervals as follows:

$[T_a, T_b], [T_a+C, T_b+C], [T_a+2C, T_b+2C], \dots, [T_a+nC, T_b+nC]$

When an LSP is configured with a recurrent time interval such as " $[T_a, T_b]$ repeats 10 times with a repeat cycle a week" (representing 11 simple time intervals), a path satisfying the constraints for the LSP in each of the simple time intervals represented by the recurrent time interval is computed and the LSP along the path is set up to

carry traffic in each of the simple time intervals.

4.3. Elastic Time Interval

An elastic time interval is a time interval with an elastic range, which is represented as within $-P$ and Q , where P and Q is an amount of time such as 300 seconds. P is called elastic range lower bound and Q is called elastic range upper bound.

For a simple time interval such as $[T_a, T_b]$ with an elastic range, elastic time interval: " $[T_a, T_b]$ within $-P$ and Q " means a time period from (T_a+X) to (T_b+X) , where $-P \leq X \leq Q$. Note that both T_a and T_b may be shifted the same X .

When an LSP is configured with elastic time interval " $[T_a, T_b]$ within $-P$ and Q ", a path is computed such that the path satisfies the constraints for the LSP in the time period from (T_a+X) to (T_b+X) and $|X|$ is the minimum value from 0 to $\max(P, Q)$. That is that $[T_a+X, T_b+X]$ is the time interval closest to time interval $[T_a, T_b]$ within the elastic range. The LSP along the path is set up to carry traffic in the time period from (T_a+X) to (T_b+X) .

Similarly, for a recurrent time interval with an elastic range, elastic time interval: " $[T_a, T_b]$ repeats n times with repeat cycle C within $-P$ and Q " represents $n+1$ simple elastic time intervals as follows:

$$[T_a+X_0, T_b+X_0], [T_a+C+X_1, T_b+C+X_1], \dots, [T_a+nC+X_n, T_b+nC+X_n]$$

where $-P \leq X_i \leq Q$, $i = 0, 1, 2, \dots, n$.

If a user wants to keep the same repeat cycle between any two adjacent time intervals, elastic time interval: " $[T_a, T_b]$ repeats n times with repeat cycle C within $-P$ and Q SYNC" may be used, which represents $n+1$ simple elastic time intervals as follows:

$$[T_a+X, T_b+X], [T_a+C+X, T_b+C+X], \dots, [T_a+nC+X, T_b+nC+X]$$

where $-P \leq X \leq Q$.

4.4. Changes to Time Interval

After a temporal LSP is configured, a user may change its parameters including some of the time intervals configured. A new time interval may be added, an existing time interval may be removed or changed.

When a new time interval is added to an existing LSP, a path satisfying the constraints for the LSP in the time interval is computed and the LSP along the path is set up to carry traffic in the time interval.

When an existing time interval is removed from an existing LSP, the time interval is deleted from the lifetime of the LSP. If the lifetime is over, the LSP is deleted.

A change to an existing time interval may generate some of four possible results:

1. The existing time interval is extended for a time period EA after the existing time period;
2. The existing time interval is extended for a time period EB before the existing time period;
3. The existing time interval is shrunk for a time period SA from the end of the existing time period; and
4. The existing time interval is shrunk for a time period SB from the beginning of the existing time period.

When an existing time interval for an LSP is extended, a path satisfying the constraints for the LSP in the extended time interval is computed and the LSP along the path is set up to carry traffic in the extended time interval. If the LSP is already up to carry traffic in the existing time interval, the lifetime of the LSP is extended for time period EA following the existing time interval.

When an existing time interval for an LSP is shrunk, the shrunk time periods are removed from the lifetime of the LSP.

4.5. Graceful Periods

For a temporal LSP, a user may want to have some graceful periods for each or some of the time intervals for the LSP. Two graceful periods may be configured for a time interval. One is the graceful period before the time interval, called grace-before, which extends the lifetime of the LSP for grace-before (such as 30 seconds) before the time interval. The other is the one after the time interval, called grace-after, which extends the lifetime of the LSP for grace-after (such as 60 seconds) after the time interval.

When an LSP is configured with a simple time interval such as [Ta, Tb] with graceful periods such as grace-before GB and grace-after GA, a path is computed such that the path satisfies the constraints for

the LSP in the time period from Ta to Tb. The LSP along the path is set up to carry traffic in the time period from (Ta-GB) to (Tb+GA). During graceful periods from (Ta-GB) to Ta and from Tb to (Tb+GA), the LSP is up to carry traffic (maybe in best effort).

5. Extensions to PCEP

This section describes the extensions to PCEP for computing paths for temporal LSP, initiating and maintaining temporal LSPs.

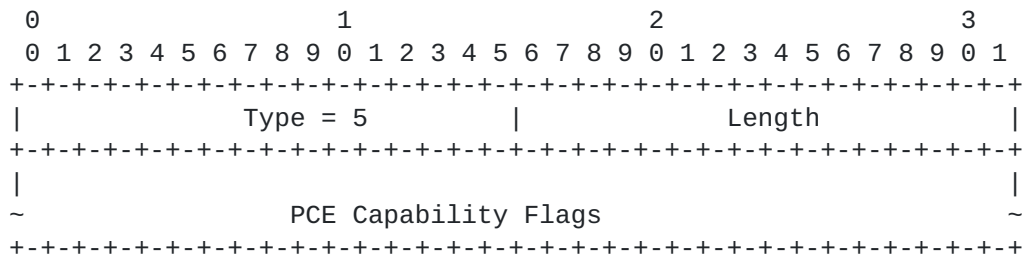
5.1. Capability TLV in Existing PCE Discovery Protocol

There are a couple of options for advertising a PCE capability for computing paths for temporal LSP, initiating and maintaining temporal LSPs.

The first option is to define a new flag in the OSPF and ISIS PCE Capability Flags to indicate the capability that a PCE is capable to compute paths for temporal LSPs, initiate and maintain temporal LSPs. This includes the capability of computing both a path for a temporal P2MP LSP and a path for a temporal P2P LSP.

The second option is to define three new flags. The first new flag in the OSPF and ISIS PCE Capability Flags indicates the capability that a PCE is capable to compute a path for a temporal P2MP LSP; the second new flag indicates the capability that a PCE is capable to compute a path for a temporal P2P LSP; and the third new flag indicates the capability that a PCE is capable to initiate and maintain a temporal LSP.

The format of the PCE-CAP-FLAGS sub-TLV is as follows:



Type: 5
Length: Multiple of 4 octets
Value: This contains an array of units of 32-bit flags numbered from the most significant as bit zero, where each bit represents one PCE capability.

The following capability bits have been assigned by IANA:

Bit	Capabilities
0	Path computation with GMPLS link constraints
1	Bidirectional path computation
2	Diverse path computation
3	Load-balanced path computation
4	Synchronized path computation
5	Support for multiple objective functions
6	Support for additive path constraints (max hop count, etc.)
7	Support for request prioritization
8	Support for multiple requests per message
9	Global Concurrent Optimization (GCO)
10	P2MP path computation
...	

Reserved bits SHOULD be set to zero on transmission and MUST be ignored on receipt.

For the first option, one bit such as bit 16 may be assigned to indicate that a PCE is capable to compute paths for temporal LSPs, initiate and maintain temporal LSPs.

Bit	Capabilities
16	Path computation for temporal LSPs, initiation and maintenance of temporal LSPs
17-31	Reserved for future assignments by IANA.

For the second option, one bit such as bit 16 may be assigned to indicate that a PCE is capable to compute a path for a temporal P2MP LSP; another bit such as bit 17 may be assigned to indicate that a PCE is capable to compute a path for a temporal P2P LSP; and yet another bit such as bit 18 may be assigned to indicate that a PCE is capable to initiate and maintain temporal LSPs.

Bit	Capabilities
16	Path computation for temporal P2MP LSP
17	Path computation for temporal P2P LSP
18	Initiation and maintenance of temporal LSP
19-31	Reserved for future assignments by IANA.

5.2. Open Message Extension

If a PCE does not advertise its capability related to computation of paths for a temporal LSP, initiation and maintenance of a temporal LSP during discovery, PCEP should be used to allow a PCC to discover, during the Open Message Exchange, which PCEs are capable of supporting computation of a path for a temporal LSP, initiation and maintenance of a temporal LSP.

To achieve this, one option is to extend the PCEP OPEN object by defining new flag bits in the value field of an existing capability TLV such as stateful PCE capability TLV in the same way as the PCE Capability Flags described in the previous section. Another option is to extend the PCEP OPEN object by defining a new optional TLV to indicate the PCE's capability to compute paths for a temporal LSP, initiate and maintain a temporal LSP.

For the second option, we need to request IANA to allocate a value such as 10 from the "PCEP TLV Type Indicators" subregistry, as documented in Section below ("Temporal LSP Capability TLV"). The description is "temporal LSP capable", and the length value is 2 bytes. The value field is set to indicate the capability of a PCE for computation of paths for a temporal LSP, initiation and maintenance of a temporal LSP in details. We can use flag bits in the value field in the same way as the PCE Capability Flags described in the previous section.

The inclusion of this TLV in an OPEN object indicates that the sender can compute paths for a temporal LSP, initiate and maintain a temporal LSP.

The capability TLV is meaningful only for a PCE, so it will typically appear only in one of the two Open messages during PCE session establishment. However, in case of PCE cooperation (e.g., inter-domain), when a PCE behaving as a PCC initiates a PCE session it SHOULD also indicate its capabilities.

5.3. RP Object Extension

The following flags are added into the RP Object:

A T bit is added in the flag bits field of the RP object to tell a receiver of a message that the message is for (computing paths for a temporal LSP) a temporal LSP.

o T (Temporal LSP bit - 1 bit):

0: This indicates that this is not a message for a temporal LSP.

1: This indicates that this is a message for a temporal LSP.

The IANA request is referenced in Section below (Request Parameter Bit Flags) of this document.

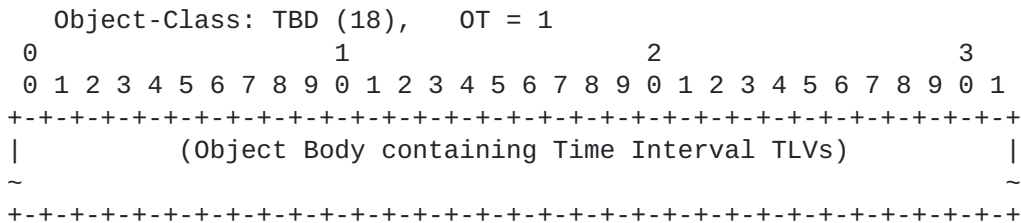
This T bit with the N bit defined in RFC6006 can indicate whether the message is for a temporal P2P LSP or P2MP LSP.

o T = 1 and N = 0: This indicates that this is a message for a temporal P2P LSP

o T = 1 and N = 1: This indicates that this is a message for a temporal P2MP LSP

5.4. TIME INTERVAL Object

For a TIME-INTERVAL object, its Class is to be assigned by IANA, here we use 18, which may be changed late. Its OT is 1, exact number to be assigned by IANA. The format of a TIME-INTERVAL object body is illustrated below, which comprises a number of time interval TLVs.



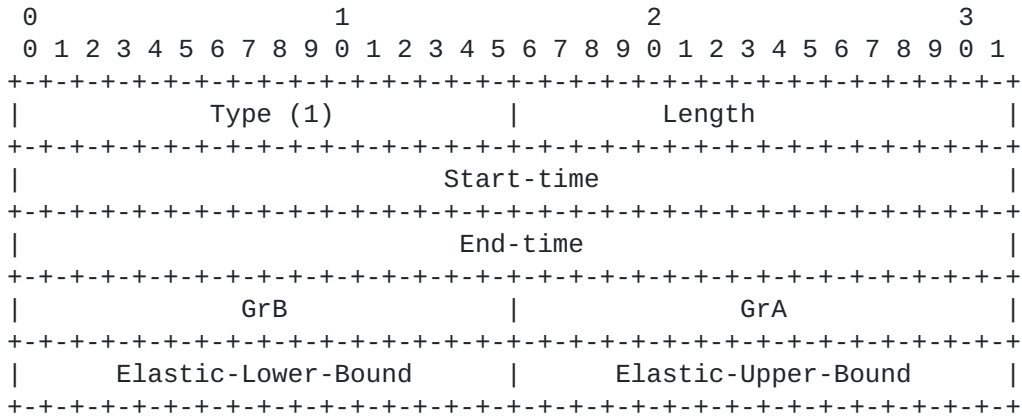
A time interval TLV may be a relative time interval TLV or an absolute time interval TLV, which are two different representations of a time interval. Their advantages and disadvantages are discussed below.

5.4.1. Absolute Time Interval TLV

The format of an absolute time interval TLV (Type = 1) for an LSP is illustrated below. It mainly contains a Start-time and an End-time, representing time interval [Start-time, End-time]. Both of these two

times are the times that are synchronized among all the elements involved. Thus the clocks on all the elements MUST be synchronized if an absolute time interval TLV is used. The time period represented in an absolute time interval TLV is more accurate.

In addition, it contains a non zero grace-before and grace-after if graceful periods are configured. It includes a non zero elastic range lower bound and upper bound if there is an elastic range configured.



- o Start-time: The time LSP starts to carry traffic.
- o End-time: The time LSP stops carrying traffic.
- o GrB (Grace-Before): The graceful period time length in seconds before time interval [Start-time, End-time].
- o GrA (Grace-After): The graceful period time length in seconds after time interval [Start-time, End-time].
- o Elastic-Lower-Bound: The maximum amount of time in seconds that time interval [Start-time, End-time] can shift to lower/left.
- o Elastic-Upper-Bound: The maximum amount of time in seconds that time interval [Start-time, End-time] can shift to upper/right.

Discussions: Optionally, we may define three TLVs below:

1. an absolute time interval TLV containing only a Start-time and an End-time;
2. an elastic range TLV containing just an elastic range lower bound and upper bound; and

3. a graceful period TLV containing only a grace-before and a grace-after.

If a time interval is with an elastic range, an absolute time interval TLV followed by an elastic range TLV is used. If a time interval is with graceful periods, an absolute time interval TLV followed by a graceful period TLV is used.

5.4.2. Relative Time Interval TLV

The format of a relative time interval TLV (Type = 2) for an LSP is shown below. It mainly contains a Start-time-length and an End-time-length, representing the time interval below for the LSP:

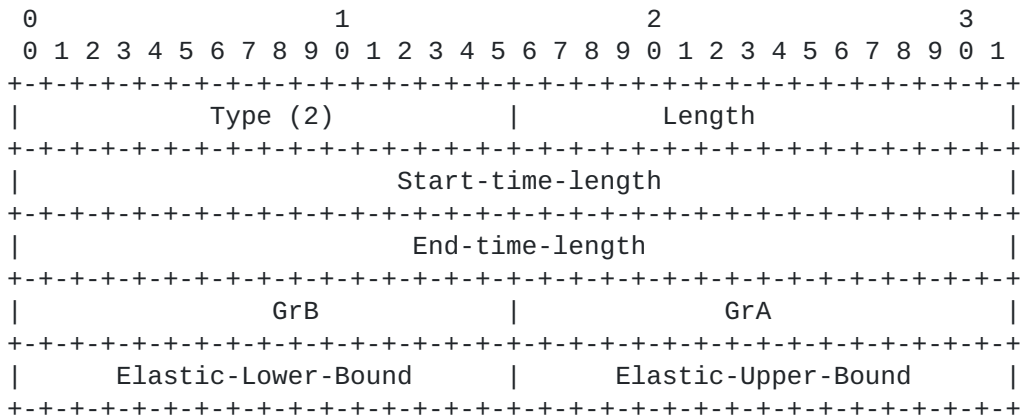
```
[current-time + Start-time-length, current-time + End-time-length]
```

where current-time is a current local time. When a time interval from time Ta to time Tb is configured on a node/element, these two time lengths are the time lengths that are computed on the node using a current local time as follows.

```
Start-time-length = Ta - current-time;  
End-time-length   = Tb - current-time;
```

For a relative time interval TLV, the clocks/times on all the elements involved can be different. But the time period represented in a relative time interval TLV on one element/node may be shifted a little bit from another element's point of view since transmitting the TLV from one element to another takes a little time, which is hard to be considered accurately.

The TLV also includes a non zero grace-before and grace-after if graceful periods are configured. It contains a non zero elastic range lower bound and upper bound if there is an elastic range configured.



- o Start-time-length: The time length in seconds from a current local time to the time LSP starts to carry traffic.
- o End-time-length: The time length in seconds from a current local time to the time LSP ends carrying traffic.
- o GrB (Grace-Before): The graceful period time length in seconds before the time interval above for the LSP.
- o GrA (Grace-After): The graceful period time length in seconds after the time interval above for the LSP.
- o Elastic-Lower-Bound: The maximum amount of time in seconds that the time interval above for the LSP can shift to lower/left.
- o Elastic-Upper-Bound: The maximum amount of time in seconds that the time interval above can shift to upper/right.

5.4.3. Recurrent Absolute Time Interval TLV

The format of a recurrent absolute time interval TLV (Type = 3) for an LSP is illustrated below. It mainly contains a Start-time, an End-time, a Repeat-time-length, a Options field and a Number-repeats.

The Start-time and End-time represents time interval [Start-time, End-time]. The Repeat-time-length represents a repeat cycle/time, which is valid if the Options field is set to indicate the way to repeat is "repeat every Repeat-time-length". The Options field indicates a way to repeat. The Number-repeats indicates the number of repeats of time interval [Start-time, End-time].

In addition, the TLV includes an non zero grace-before and grace-after if graceful periods are configured. It contains an non zero

elastic range lower bound and upper bound if there is an elastic range configured.

0									1									2									3								
0	1	2	3	4	5	6	7	8	0	1	2	3	4	5	6	7	8	0	1	2	3	4	5	6	7	8	0	1	2	3	4	5	6	7	8
Type (3)									Length																										
Start-time																																			
End-time																																			
Repeat-time-length																																			
Options			Number-repeats									Reserved (0)																							
GrB									GrA																										
Elastic-Lower-Bound									Elastic-Upper-Bound																										

- o Start-time: The time LSP starts to carry traffic.
- o End-time: The time LSP stops carrying traffic.
- o Repeat-time-length: The time length in seconds after which LSP starts to carry traffic again for (End-time - Start-time).
- o Options: Indicates a way to repeat.
 - Options = 1: repeat every day;
 - Options = 2: repeat every week;
 - Options = 3: repeat every month;
 - Options = 4: repeat every year;
 - Options = 5: repeat every Repeat-time-length.
- o Number-repeats: The number of repeats. In each of repeats, LSP carries traffic.
- o GrB (Grace-Before): The graceful period time length in seconds before each of the time intervals represented by the recurrent time interval.

- o GrA (Grace-After): The graceful period time length in seconds after each of the time intervals.
- o Elastic-Lower-Bound: The maximum amount of time in seconds that each of the time intervals can shift to lower/left.
- o Elastic-Upper-Bound: The maximum amount of time in seconds that each of the time intervals can shift to upper/right.

5.4.4. Recurrent Relative Time Interval TLV

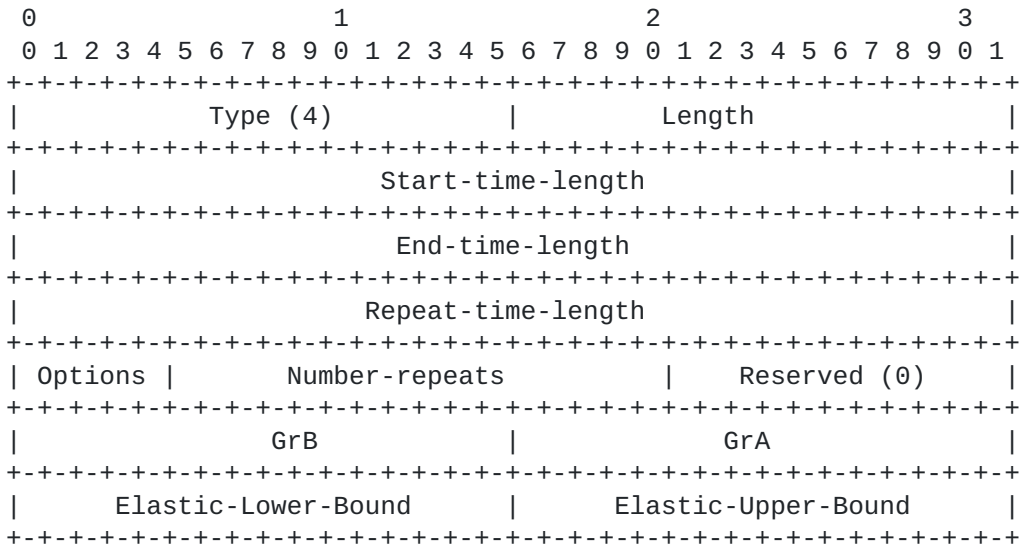
The format of a recurrent relative time interval TLV (Type = 4) is shown below. It mainly contains a Start-time-length, an End-time-length, a Repeat-time-length, a Options field and a Number-repeats.

The Start-time-length and End-time-length represents time interval

[current-time + Start-time-length, current-time + End-time-length]

where current-time is a current local time. The Repeat-time-length represents a repeat cycle/time, which is valid if the Options field is set to indicate the way to repeat is "repeat every Repeat-time-length". The Options field indicates a way to repeat. The Number-repeats indicates the number of repeats of the time interval above.

In addition, the TLV includes an non zero grace-before and grace-after if graceful periods are configured. It contains an non zero elastic range lower bound and upper bound if there is an elastic range configured.



- o Start-time-length: The time length in seconds from a current local time to the time LSP starts to carry traffic.
- o End-time-length: The time length in seconds from a current local time to the time LSP stops carrying traffic.
- o Repeat-time-length: The time length in seconds after which LSP starts to carry traffic again for (End-time-length - Start-time-length).
- o Options: Indicates a way to repeat.
 - Options = 1: repeat every day;
 - Options = 2: repeat every week;
 - Options = 3: repeat every month;
 - Options = 4: repeat every year;
 - Options = 5: repeat every Repeat-time-length.
- o Number-repeats: The number of repeats. In each of repeats, LSP carries traffic.
- o GrB (Grace-Before): The graceful period time length in seconds before each of the time intervals represented by the recurrent time interval.
- o GrA (Grace-After): The graceful period time length in seconds after each of the time intervals.
- o Elastic-Lower-Bound: The maximum amount of time in seconds that each of the time intervals can shift to lower/left.
- o Elastic-Upper-Bound: The maximum amount of time in seconds that each of the time intervals can shift to upper/right.

5.5. Messages for Temporal LSP

This section presents and discusses two classes of messages. One class is the messages between a PCE and a PCC on the ingress of a temporal LSP for initiating and maintaining the LSP. The other is the messages between two PCEs, one of which acts as a PCC.

5.5.1. Messages between PCE and PCC on Ingress

From function's point of view, there are four groups of messages:

1. LSP creation request messages,
2. LSP deletion request messages,
3. LSP creation response messages, and
4. LSP deletion response messages.

A message for an LSP in the first two groups is sent from a PCE to the PCC on the ingress of the LSP. A message for an LSP in the last two groups is sent from the PCC on the ingress of the LSP to a PCE.

A Path Computation LSP Initiate Request (PCInitiate) message without any extensions can be used for a message in the first two groups. A Path Computation LSP State Report (PCRpt) message without any extensions can be used for a message in the last two groups.

Alternatively, a PCInitiate message with some optional extensions such as TIME-INTERVAL can be used for a message in the first two groups. A PCRpt message with some optional extensions such as TIME-INTERVAL can be used for a message in the last two groups.

For an LSP creation request, a PCInitiate message includes objects: SRP, LSP, END-POINTS, ERO and optional TIME-INTERVAL. SRP (Stateful PCE Request Parameters) object comprises an SRP-ID-number. LSP object comprises PLSP-ID of 0, and SYMBOLIC-PATH-NAME TLV with path name. END-POINTS object comprises the source and destination addresses of the LSP. ERO object comprise the path (i.e., ERO) for the LSP. TIME-INTERVAL object comprises the time intervals for the LSP (the path satisfies constraints for the LSP in each of the time intervals).

For an LSP creation response, a PCRpt message includes objects: SRP, LSP, ERO and optional TIME-INTERVAL. SRP object comprises the SRP-ID-number in the corresponding LSP creation request message. LSP object comprises a PLSP-ID assigned to the LSP by the PCC, SYMBOLIC-PATH-NAME TLV with path name, C Flag set to 1 indicates that this LSP is created by the LSP creation request. ERO object comprise the path (i.e., ERO) for the LSP. TIME-INTERVAL object comprises the time intervals for the LSP.

For an LSP deletion request, a PCInitiate message includes objects: SRP, LSP, and optional TIME-INTERVAL. SRP object comprises an SRP-ID-number and R (remove) flag set to 1. LSP object comprises the

PLSP-ID for the LSP created. TIME-INTERVAL object comprises the time intervals for the LSP.

For an LSP deletion response, a PCRpt message includes objects: SRP, LSP, and optional TIME-INTERVAL. SRP object comprises the SRP-ID-number in the corresponding LSP deletion request message. LSP object comprises R(Remove) flag set to 1 indicating that the LSP has been removed from the PCC, and LSP Identifiers TLV.

Note: The PCC on the ingress of an LSP does not use any time intervals in the TIME-INTERVAL object received for signaling the LSP. For just creating and deleting LSPs, we do not need to include any TIME-INTERVAL object in a message if the PCE creates the LSP with a sequence of time intervals at the beginning of each of the time intervals and deletes the LSP at the end of each of the time intervals.

Discussions: For an LSP having a time interval TLV with graceful periods, we may create the LSP in the time period including the graceful periods and the LSP has the reserved bandwidth during that period (including the graceful periods).

Another option is that we create the LSP in the time period including the graceful periods, but do not reserve any bandwidth for the LSP in the beginning. The desired bandwidth for the LSP is reserved in the time period without graceful periods.

After the graceful period before the time interval, the bandwidth for the LSP is reserved through a update message from the PCE to the PCC on the ingress of the LSP. After the time interval (i.e., just before the graceful period after the time interval), the bandwidth for the LSP is released through another update message from the PCE to the PCC on the ingress of the LSP.

5.5.2. Messages between two PCEs

Figure below illustrates the format of a request message with a optional TIME-INTERVAL object for computing paths for a temporal LSP with a sequence of time intervals:

```

<PCReq Message> ::= <Common Header>
                    [<svec-list>]
                    <request-list>
<request-list> ::= <request> [<request-list>]
<request> ::= <RP> <END-POINTS> [<OF>] [<LSPA>] [<BANDWIDTH>]
               [<metric-list>] [<RRO> [<BANDWIDTH>]] [<IRO>]
               [<LOAD-BALANCING>]
               [<TIME-INTERVAL>]

```

Figure 1: Format for Request Message

Below is the format of a reply message with a optional TIME-INTERVAL object:

```

<PCRep Message> ::= <Common Header>
                    <response-list>
<response-list> ::= <response> [<response-list>]
<response> ::= <RP>
               [<NO-PATH>]
               [<attribute-list>]
               [<path-list>]
<path-list> ::= <path> [<path-list>]
<path> ::= <ERO> <attribute-list> [<TIME-INTERVAL>]

```

Figure 2: Format for Reply Message

6. Procedures

This section focuses on the procedures for creating and deleting a temporal LSP. When a PCE receives a request for an LSP with a sequence of time intervals from a user or application, it computes a path satisfying the constraints for the LSP in each of the time intervals and reserved the bandwidth for the LSP along the path in each of the time intervals. And then it initiates the creation of the LSP in the network to carry traffic in each of the time intervals.

There are a couple of ways for a PCE to create an LSP with a sequence of time intervals. One way is that the PCE initiates the creation of the LSP at the beginning of each of the time intervals. At the end of each of the time intervals or when a deletion request for the LSP received, the PCE initiates the deletion of the LSP.

Another way is that the PCE initiates the creation of the LSP at or before the beginning of the first time interval and the deletion of

the LSP at the end of the last time interval. At the start of each time interval, the PCE initiates the update of the LSP with the reserved resource such as link bandwidth. At the end of the each time interval, the PCE initiates the update of the LSP with zero resource.

We will focus on the first way below.

6.1. Creating a Temporal LSP

A procedure for creating a temporal LSP is as follows:

Step 1: A PCE receives a request for creating a temporal LSP from a user or application and stores the parameters of the LSP into an LSP database (LSPDB) such as LSP State Database. The parameters include a number of time intervals for the LSP.

Step 2: The PCE computes a shortest path satisfying constraints for the LSP in each of the time intervals given. It reserves the resources such as the bandwidth in TED on each of the links the LSP traverses for each of the time intervals and stores the information about the LSP into the LSPDB. The information includes the paths computed for the LSP and the resources such as link bandwidth reserved for the LSP.

Step 3: At the beginning of each of the time intervals, the PCE initiates the setup of the LSP in a network through sending an LSP creation request (e.g., a PCInitiate with LSP object with PLSP-ID=0) with the path for the LSP to the PCC on the ingress of the LSP, which triggers RSVP-TE to signal the LSP along the path in the network (Note that the RSVP-TE is not aware of any time interval for the LSP and just sets up the LSP in a normal way). The PCC sends an LSP creation response (e.g., a PCRpt) to the PCE after the LSP is up.

Step 4: The PCE receives the LSP creation response (e.g., the PCRpt) from the PCC corresponding to the request and updates the status of the LSP in the LSPDB accordingly.

6.2. Deleting a Temporal LSP

Suppose that a temporal LSP has been created to carry traffic in a sequence of time intervals. A procedure for deleting this temporal LSP is as follows:

Step 1: A PCE receives a request for deleting the temporal LSP from an client, or the lifetime for the LSP in a time interval is over and the LSP needs to be deleted.

Step 2: The PCE finds the LSP in the LSPDB and gets the information about the LSP.

Step 3: The PCE initiates the deletion of the LSP in the network through sending an LSP deletion request (e.g., a PCInitiate with R flag set and PLSP ID for the LSP) to the PCC on the ingress of the LSP, which triggers the RSVP-TE to tear down the LSP in the network (Note that the RSVP-TE is not aware of any time interval for the LSP and just tears down the LSP in a normal way). The PCC generates an LSP deletion response (e.g., a PCRpt with R flag set) and sends it to the PCE after the LSP is torn down.

Step 4: The PCE receives the LSP deletion response (e.g., the PCRpt) from the PCC corresponding to the request and updates the status of the LSP in the LSPDB accordingly. For deleting the LSP completely as requested, it releases the resources such as the link bandwidth reserved for the LSP in TED for each of the time intervals and removes the information about the LSP from the LSPDB after the LSP is deleted.

7. Considerations on TED

The existing Traffic Engineering (TE) information in a TED represents an unreserved bandwidth B_i at each of eight priority levels for a link at one point of time, for example, at the current time.



This means that the link has bandwidth B_i at a priority level from now to forever until there is a change to it. Thus, a TE Label Switching Path (LSP) tunnel for a given time interval cannot be set up in advance using the information in the TED and the bandwidth cannot be reserved in advance for the LSP in the time interval given.

TED needs to be enhanced for supporting temporal LSPs. Two options


for enhancing TED are presented below.

7.1. TE Representation in Absolute Time

Suppose that the amount of the unreserved bandwidth at a priority level for a link is B_j in a time interval from time T_j to T_k ($k = j+1$), where $j = 0, 1, 2, \dots$. The unreserved bandwidth for the link can be represented as

$$[T_0, B_0], [T_1, B_1], [T_2, B_2], [T_3, B_3], \dots$$

This is an absolute time representation of bandwidths for a link. Time T_j ($j = 0, 1, 2, \dots$) MUST be a synchronized time among all the elements involved.



If an LSP is completely deleted at time T and uses bandwidth B , then for every time interval/period (after time T) during which bandwidth B is reserved for the LSP on a link, B is added to the link for that interval/period.

If an LSP is to be up at time T and uses bandwidth B , then for every time interval/period (after time T) during which bandwidth B is reserved for the LSP on a link, B is subtracted from the link for that interval/period.

7.2. TE Representation in Relative Time

Alternatively, a relative time representation of bandwidths for a link can be used. For example, the amount of the unreserved bandwidth at a priority level for a link is B_j during a series of time intervals/periods can be expressed as

$$[P_0, B_0], [P_1, B_1], [P_2, B_2], [P_3, B_3], \dots, \text{ where } P_j = T_k - T_j, k = (j+1) \text{ and } j = 0, 1, 2, 3, \dots$$

In this representation, every time T_j ($j = 0, 1, 2, \dots$) can be a local time. A timer may expire for every unit of time (e.g., every second) and may trigger $--P_0$, which decrements P_0 . When $P_0 = 0$, P_1 becomes P_0 , P_2 becomes P_1 , and so on.

If an LSP is completely deleted at time T and uses bandwidth B , then for every time interval/period (after time T) during which bandwidth B is reserved for the LSP on a link, B is added to the link for that interval/period.

If an LSP is to be up at time T and uses bandwidth B , then for every time interval/period (after time T) during which bandwidth B is reserved for the LSP on a link, B is subtracted from the link for that interval/period.

An advantage of using relative time representation is that the times or clocks on all the elements involved can be different.

8. Security Considerations

The mechanism described in this document does not raise any new security issues for the PCEP protocols.

9. IANA Considerations

This section specifies requests for IANA allocation.

10. Acknowledgement

The authors would like to thank everyone who give his/her valuable comments on this draft.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009,

<<http://www.rfc-editor.org/info/rfc5440>>.

[RFC6006] Zhao, Q., Ed., King, D., Ed., Verhaeghe, F., Takeda, T., Ali, Z., and J. Meuric, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", [RFC 6006](#), DOI 10.17487/RFC6006, September 2010, <<http://www.rfc-editor.org/info/rfc6006>>.

[I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", [draft-ietf-pce-stateful-pce-18](#) (work in progress), December 2016.

[I-D.ietf-pce-pce-initiated-lsp]
Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", [draft-ietf-pce-pce-initiated-lsp-07](#) (work in progress), July 2016.

11.2. Informative References

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.

[RFC5862] Yasukawa, S. and A. Farrel, "Path Computation Clients (PCC) - Path Computation Element (PCE) Requirements for Point-to-Multipoint MPLS-TE", [RFC 5862](#), DOI 10.17487/RFC5862, June 2010, <<http://www.rfc-editor.org/info/rfc5862>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
US

Email: huaimo.chen@huawei.com

Xufeng Liu
Ericsson
USA

Email: xliu@kuatrotech.com

Mehmet Toy
Verizon
USA

Email: mehmet.toy@verizon.com

Vic Liu
China Mobile
No.32 Xuanwumen West Street, Xicheng District
Beijing, 100053
China

Email: liu.cmri@gmail.com

Lei Liu
Fujitsu
USA

Email: lliu@us.fujitsu.com

Khuzema Pithewan
Infinera

Email: kpithewan@infinera.com