

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: April 18, 2016

H. Chen
Huawei Technologies
M. Toy
Comcast
L. Liu
Fijitsu
October 16, 2015

Framework for Temporal Tunnel Services
draft-chen-teas-frmwk-tts-00.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 a framework for temporal LSP tunnel services and provides a few of reference models along with logical components required to design a solution for TTS.

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 April 18, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the

Internet-Draft

Framework for TTS

October 2015

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
2.	Terminology	3
3.	Operations Overview	4
3.1.	Simple Time Interval	4
3.2.	Recurrent Time Interval	4
3.3.	Changes to Time Interval	4
4.	Reference Models	5
4.1.	Building Blocks	5
4.1.1.	Temporal TED	6
4.1.2.	Temporal CSPF	7
4.1.3.	Temporal Label Database	7
4.1.4.	Temporal LSP Tunnel Manager	8
4.1.5.	Temporal LSP Database	8
4.1.6.	Temporal PCE	9
4.2.	Centralized Model	9
4.2.1.	Centralized Model for Single Domain	9
4.2.2.	Centralized Model for Multiple Domains	12
4.3.	Hybrid Model	13
4.3.1.	Hybrid Model for Single Domain	14
4.3.2.	Hybrid Model for Multiple Domains	16
4.3.3.	Temporal Stateful PCE	17
4.4.	Distributed Model	18
5.	Security Considerations	20
6.	Acknowledgement	20
7.	References	20
7.1.	Normative References	20
7.2.	Informative References	21
	Authors' Addresses	22

[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 to consume the reserved network resources for it forever even though the LSP may only use the network resources during some period of time. As a result, the entire 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/intervals.

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/periods and it is assumed to consume the resources and carry traffic only in these time intervals. Thus the entire network resources are efficiently used. Moreover, some new services can be provided easily. For example, a user can book a tunnel service in advance for a given time interval or a sequence of given time intervals. Tunnel services may be easily scheduled.

This document specifies a framework for temporal LSP tunnel services and provides a few of reference models along with logical components required to design a solution for TTS.

[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.

PCE: Path Computation Element.

PCEP: Path Computation Element Communication Protocol.

[3.](#) Operations Overview

This section briefly describes some operations on a temporal LSP.

[3.1.](#) 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.

[3.2.](#) 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: "[Ta, Tb] repeats n times with repeat cycle C" represents n+1 simple time intervals as follows:

[Ta, Tb], [Ta+C, Tb+C], [Ta+2C, Tb+2C], ..., [Ta+nC, Tb+nC]

When an LSP is configured with a recurrent time interval such as "[Ta, Tb] 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.

[3.3.](#) 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.](#) Reference Models


This section presents a few of reference models for TTS after introducing some of building blocks.

[4.1.](#) Building Blocks

This section briefly describes some of the components that may be used to build a solution for creating and maintaining temporal LSP tunnels.

[4.1.1.](#) Temporal TED

The Traffic Engineering (TE) information in a normal TE Database (TED) represents a unreserved bandwidth B_i at each of eight priority levels for a link at one point of time, i.e., 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.

The TED needs to be enhanced for supporting temporal LSPs. The enhanced TED is called Temporal TED (T-TED for short). It maintains the TE information such as bandwidth for every link with a series of time intervals.

For example, 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

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



The unreserved bandwidth information above for the link will be stored in the T-TED.

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 time 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 time interval/period.

[4.1.2.](#) Temporal CSPF

An existing constrained shortest path first (CSPF) (or say a normal CSPF) computes a path for a normal LSP that satisfies a set of given constraints using a traffic engineering database (TED).

A temporal CSPF (T-CSPF for short) computes a path for a temporal LSP (i.e., an LSP with a number of time intervals) that satisfies a set of given constraints in each of the time intervals through using a temporal TED (T-TED).

[4.1.3.](#) Temporal Label Database

In a centralized controller, a normal label database (LDB) records and maintains the status of every label for every node (and/or interface) in a network, which the controller controls. The status of a label indicates whether the label is assigned to an LSP.

A temporal label database (T-LDB) in a centralized controller records and maintains the status of every label in a series of time intervals for every node (and/or interface) in a network, on which the controller controls. The status of a label in a time interval indicates whether the label is assigned to an LSP in the time interval.

If there are enough labels anytime, we do not need any temporal label database and we can just use a normal label database. For example, if we can make sure that at any time the number of LSPs going through any node in the network is less than the number of labels on the node, then there are enough labels anytime. Thus, we can just use a normal label database.

[4.1.4.](#) Temporal LSP Tunnel Manager

An existing LSP tunnel manager (or say a normal LSP tunnel manager) receives a request for an operation on an MPLS TE LSP from a user or an application. The operation may be a creation of a new MPLS TE LSP tunnel, a deletion of an existing MPLS TE LSP tunnel, or a change to an existing LSP tunnel.

A temporal LSP tunnel manager (T-LSP Manager for short) receives a request for an operation on a temporal LSP from a user or an application. The operation may be a creation of a new temporal LSP tunnel, a deletion of an existing temporal LSP tunnel, or a change to an existing temporal LSP tunnel.

When receiving a request for creating a new temporal LSP (i.e., an LSP with a sequence of time intervals), the T-LSP manager asks the T-CSPF to compute a path for the LSP that satisfies the constraints given for the LSP in each of the time intervals.

After obtaining the path for the LSP from the T-CSPF, the T-LSP manager requests the T-LDB to assign labels along the path for the LSP and asks the T-TED to reserve the resources such as link bandwidth along the path for the LSP in each of the time intervals if it is used in a centralized controller.

The T-LSP manager in a centralized controller really sets up the LSP along the path in the network by writing a forwarding entry on every node along the path through the API to the network in each of the time intervals. The T-LSP manager in a distributed environment initiates the RSVP signaling to set up the LSP along the path.

The T-LSP manager records the related information for the LSP into the temporal LSP database (T-LSPDB). The information includes the time intervals for the LSP, the path computed for the LSP, the resources such as bandwidth reserved along the path in each of the time intervals for the LSP (for centralized controller), the labels assigned along the path for the LSP (for centralized controller), and the status of the LSP.

[4.1.5.](#) Temporal LSP Database

A temporal LSP database (T-LSPDB for short) in a centralized controller stores the related information for every temporal LSP. For each LSP, the following information will be stored in the T-LSPDB:

- o the time intervals for the LSP,
- o the paths computed for the LSP,
- o the resources such as bandwidth reserved along the path in each of the time intervals for the LSP,
- o the labels assigned along the path for the LSP, and
- o the status of the LSP.

In a distributed environment, a T-LSPDB on a label edge router (LER) stores the following information for every temporal LSP originating from the LER (i.e., the LER is the ingress of the LSP):

- o the time intervals for the LSP,
- o the paths computed for the LSP, and
- o the status of the LSP.

[4.1.6.](#) Temporal PCE

A temporal PCE (T-PCE for short) is an enhanced version of the existing PCE. It receives a request for computing a path for a temporal LSP crossing multiple domains, computes the path for the LSP and replies to the request with the path computed. For the LSP with a number of time intervals and some constraints, the path computed satisfies the constraints in each of the time intervals.

[4.2.](#) Centralized Model

This section presents two centralized reference models. one model is for a single domain, the other for multiple domains.

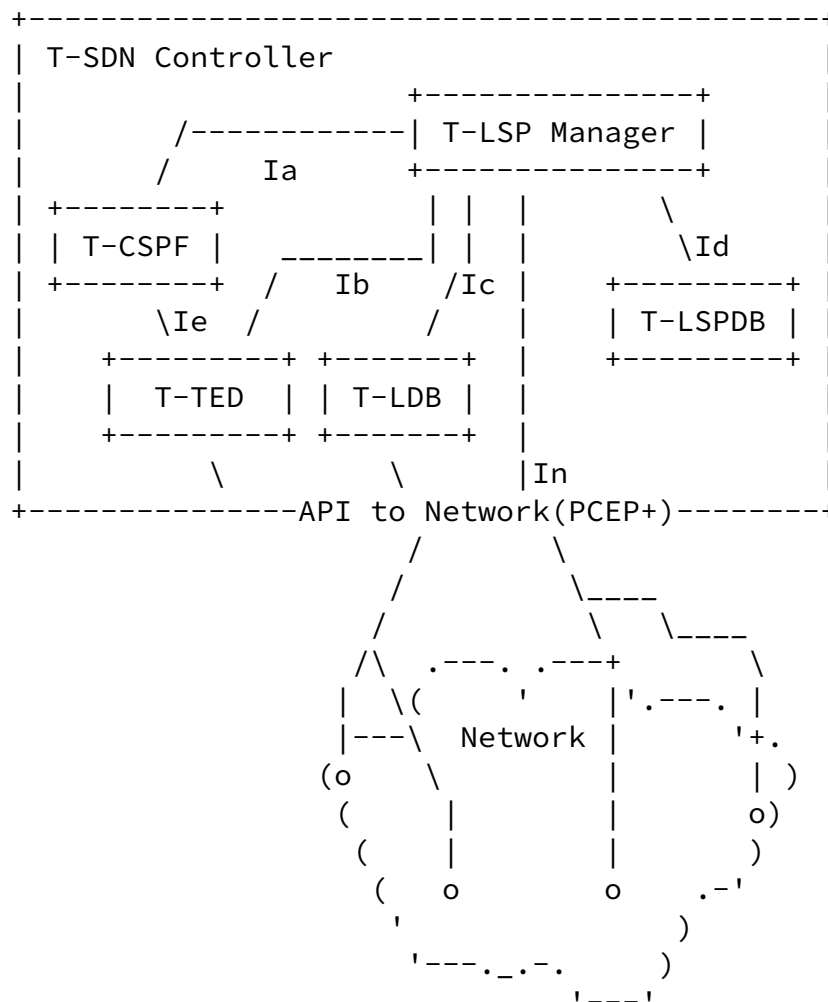
[4.2.1.](#) Centralized Model for Single Domain

Figure below illustrates a centralized SDN controller reference model for temporal LSP tunnel services for a network (i.e., a single domain).

Internet-Draft

Framework for TTS

October 2015



The temporal SDN (T-SDN) controller in the reference model controls a network through an API to the network such as PCEP+ (extensions to PCEP for central controller). The T-SDN controller is responsible for creating and maintaining every temporal LSP in the network.

The T-SDN controller comprises a number of modules, including a T-LSP manager, a T-CSPF, a T-TED, a T-LDB and a T-LSPDB. The interfaces among these modules are listed as follows:

- o Interface Ia between the T-LSP manager and the T-CSPF. Through

this interface, the T-LSP manager requests the T-CSPF to compute a path for a temporal LSP with a number of time intervals and a set of constraints, and the T-CSPF responses the T-LSP manager with the path computed that satisfies the constraints in each of the time intervals.

- o Interface Ib between the T-LSP manager and the T-TED. When a temporal LSP with a number of time intervals is to be created, through this interface, the T-LSP manager reserves in the T-TED

the TE resources such as link bandwidths on every link in each of the time intervals along the path computed for the LSP. When a temporal LSP with a number of time intervals is deleted, the T-LSP manager releases the TE resources such as link bandwidths on every link in each of the time intervals along the path for the LSP.

- o Interface Ic between the T-LSP manager and the T-LDB. When a temporal LSP with a number of time intervals is to be created, through this interface, the T-LSP manager reserves in the T-LDB a label for every link in each of the time intervals along the path computed for the LSP. When a temporal LSP with a number of time intervals is deleted, the T-LSP manager releases the label for every link in each of the time intervals along the path for the LSP.
- o Interface Id between the T-LSP manager and the T-LSPDB. the T-LSP manager updates the information for every LSP in the T-LSPDB through this interface.
- o Interface Ie between the T-CSPF and the T-TED. Through this interface, the T-CSPF accesses the traffic engineering information such as link bandwidths when it computes a path for a temporal LSP with a number of time intervals.

There is an interface In between the T-SDN controller and the network. In fact, there is a control channel (or interface) between the T-SDN controller and every node in the network.

Initially, the T-TED obtains the original traffic engineering (TE) information such as link bandwidths from the network through the interface In (i.e., API to network) for every link in the network. The T-LDB gets the original label resources from the network through

the interface In for every node and link in the network.

Then the TE information in the T-TED is updated mostly by the following events.

- o When a temporal LSP with a number of time intervals is to be created, the T-LSP manager reserves in the T-TED bandwidths on every link in each of the time intervals along the path for the LSP.
- o When a temporal LSP with a number of time intervals is deleted, the T-LSP manager releases bandwidths on every link in each of the time intervals along the path for the LSP.
- o When a link in the network is down, the TE information about the link is removed from the T-TED.

- o When a link in the network is up, the TE information about the link is added into the T-TED.

The label resources in the T-LDB may be updated as follows:

- o When a temporal LSP with a number of time intervals is to be created, the T-LSP manager reserves in the T-LDB a label for every link in each of the time intervals along the path for the LSP. For a node specific label space, a label on the downstream node is assigned for the link. For a link specific label space, a label on the link is assigned for the link.
- o When a temporal LSP with a number of time intervals is deleted, the T-LSP manager releases the label for every link in each of the time intervals along the path for the LSP.
- o When a node in the network is down, the label resources on the node is removed from the T-LDB if a node specific label space is used. When a link in the network is down, the label resources on the link is removed from the T-LDB if a link specific label space is used.
- o When a node in the network is up, the label resources on the node is added into the T-LDB if a node specific label space is used. When a link in the network is up, the label resources on the link

is added into the T-LDB if a link specific label space is used.

There are a couple of ways in which the T-SDN controller sets up a temporal LSP with a number of time intervals in the network.

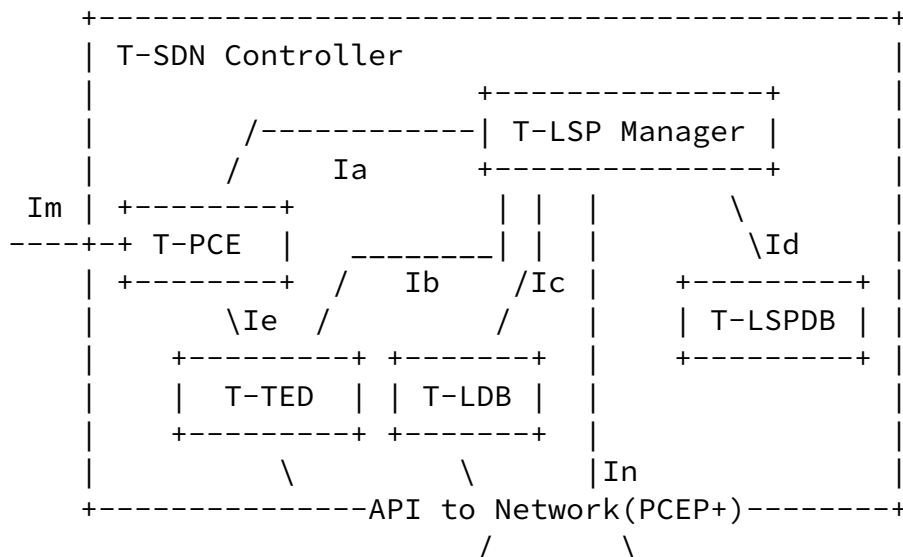
One way is to set up the LSP in the network along the path computed for the LSP at the start of each time interval and to delete the LSP at the end of each time interval.


Another way is to set up the LSP in the network along the path computed for the LSP before or at the start of the first time interval, to update the parameters such as bandwidth for each time interval, and to delete the LSP at the end of last time interval.

[4.2.2.](#) Centralized Model for Multiple Domains

The centralized model described in the previous section is for temporal LSPs within a single domain, which is called single-domain centralized model. It can be easily extended to support temporal LSPs crossing multiple domains. The extended model is called multi-domain centralized model. Basically, through replacing the T-CSPF module with a T-PCE module in the single-domain centralized model, we obtain a multi-domain centralized model.

Figure below illustrates a centralized SDN controller reference model for temporal LSPs crossing multiple domains.





The T-PCE may be outside of the T-SDN controller. When receiving a request for creating a new temporal LSP with a number of time intervals and some constraints, the T-SDN controller (or say the T-LSP manager) asks the T-PCE to compute a path for the LSP.


For computing a path for a temporal LSP crossing multiple domains, the T-PCE communicates with other T-PCEs through interface I_m to get an end to end path for the LSP crossing domains. For computing a path for a temporal LSP within the network (one domain), the T-PCE uses a T-CSPF inside it to obtain a path for the LSP.

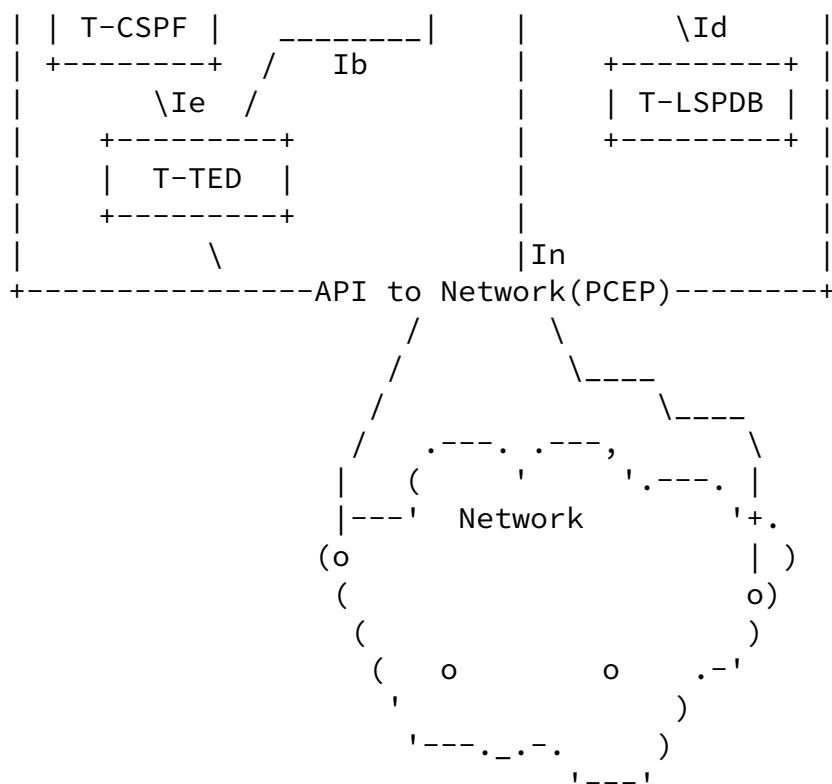
4.3. Hybrid Model

This section presents a couple of hybrid reference models. one model is for a single domain, the other for multiple domains.

4.3.1. Hybrid Model for Single Domain

Figure below illustrates a hybrid SDN controller reference model for temporal LSP tunnel services within a single domain.





The temporal SDN (T-SDN) controller in this hybrid reference model manages some parts of the resources in the network it controls. For example, it may just manage the link bandwidth for every link in the network. The label resources in the network is not managed by the T-SDN controller. It may still be managed by each node in the network.

The T-SDN controller controls the network through an API to the network such as PCEP. There is a control channel between the T-SDN controller and each of the LERs in the network. The T-SDN controller is responsible for creating and maintaining every temporal LSP in the network through the control channel to the ingress node of the LSP.

The T-SDN controller comprises a T-LSP manager, a T-CSPF, a T-TED and

a T-LSPDB. The interfaces among these modules are listed as follows:

- o Interface Ia between the T-LSP manager and the T-CSPF. This interface is the same as the one between the T-LSP manager and the

T-CSPF in the centralized model.

- o Interface Ib between the T-LSP manager and the T-TED. This interface is the same as the one between the T-LSP manager and the T-TED in the centralized model.
- o Interface Id between the T-LSP manager and the T-LSPDB. This interface is similar to the one between the T-LSP manager and the T-LSPDB in the centralized model. Most of the information for a temporal LSP stored in the T-LSPDB in the hybrid model is the same as that stored in the T-LSPDB in the centralized model. For example, the time intervals associated with the LSP and the link bandwidths reserved for the LSP in each of the time intervals are the same in both models. The labels assigned to the LSP is stored in the T-LSPDB in the centralized model, but there is not any label information for the LSP stored in the T-LSPDB in the hybrid model.
- o Interface Ie between the T-CSPF and the T-TED. This interface is the same as the one between the T-CSPF and the T-TED in the centralized model.

The TE information in the T-TED in the hybrid model is updated in the same way as that in the T-TED in the centralized model. But the way in which the T-TED in one model obtains the original TE information from the network may be different from the one in another model.

For example, the T-TED in the centralized model may obtain the original TE information from the network through polling every node in the network. The T-TED in the hybrid model may get the original TE information from the network through an OSPF or ISIS adjacency between the T-SDN controller and one of the nodes in the network.

There are a couple of ways in which the T-SDN controller sets up a temporal LSP with a number of time intervals in the network.

One way is that the T-SDN controller asks the ingress of the LSP to signal the LSP in the network along the path computed for the LSP at the start of each time interval and to tear down the LSP at the end of each time interval.

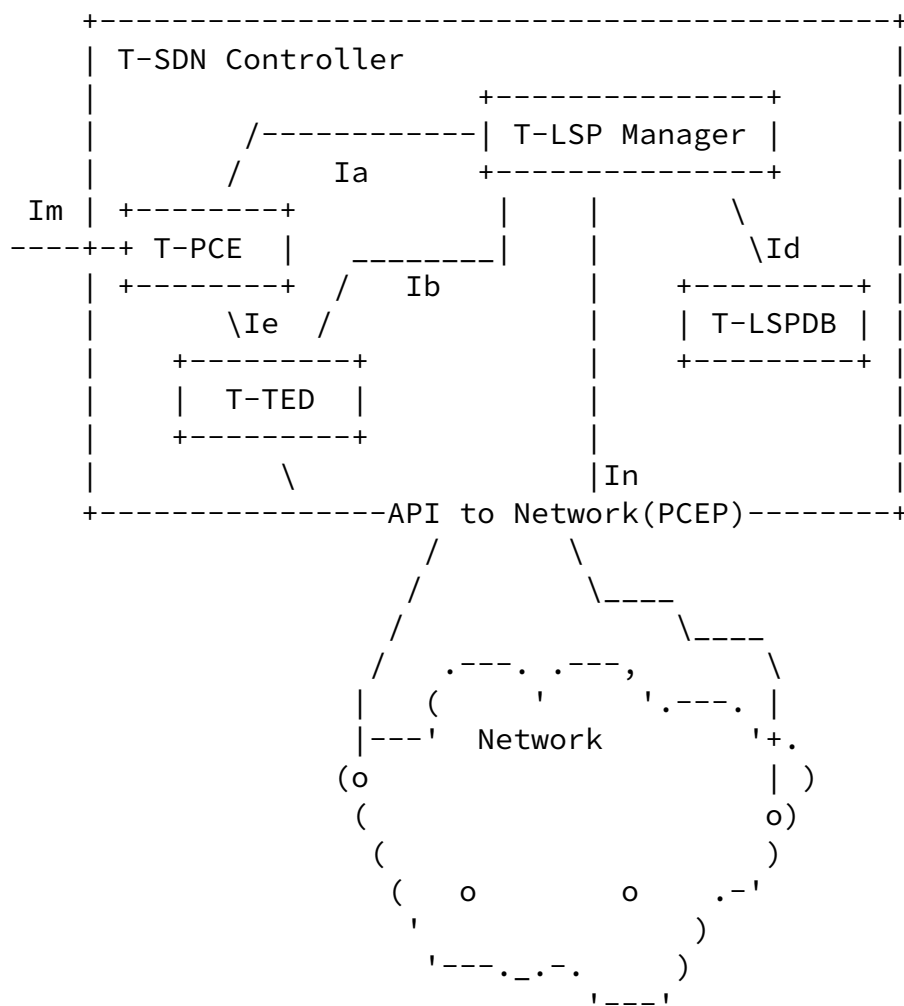
Another way is that the T-SDN controller asks the ingress of the LSP to signal the LSP in the network along the path computed for the LSP before or at the start of the first time interval, to update the

parameters such as bandwidth for each time interval, and to tear down the LSP at the end of the last time interval.

[4.3.2.](#) Hybrid Model for Multiple Domains

The hybrid model described in the previous section is for temporal LSPs within a single domain, which is called single-domain hybrid model. It can be easily extended to support temporal LSPs crossing multiple domains. The extended model is called multi-domain hybrid model. Basically, through replacing the T-CSPF module with a T-PCE module in the single-domain hybrid model, we obtain a multi-domain hybrid model.

Figure below illustrates a hybrid SDN controller reference model for temporal LSPs crossing multiple domains.



The T-PCE may be outside of the T-SDN controller. When receiving a request for creating a new temporal LSP with a number of time

intervals and some constraints, the T-SDN controller (or say the

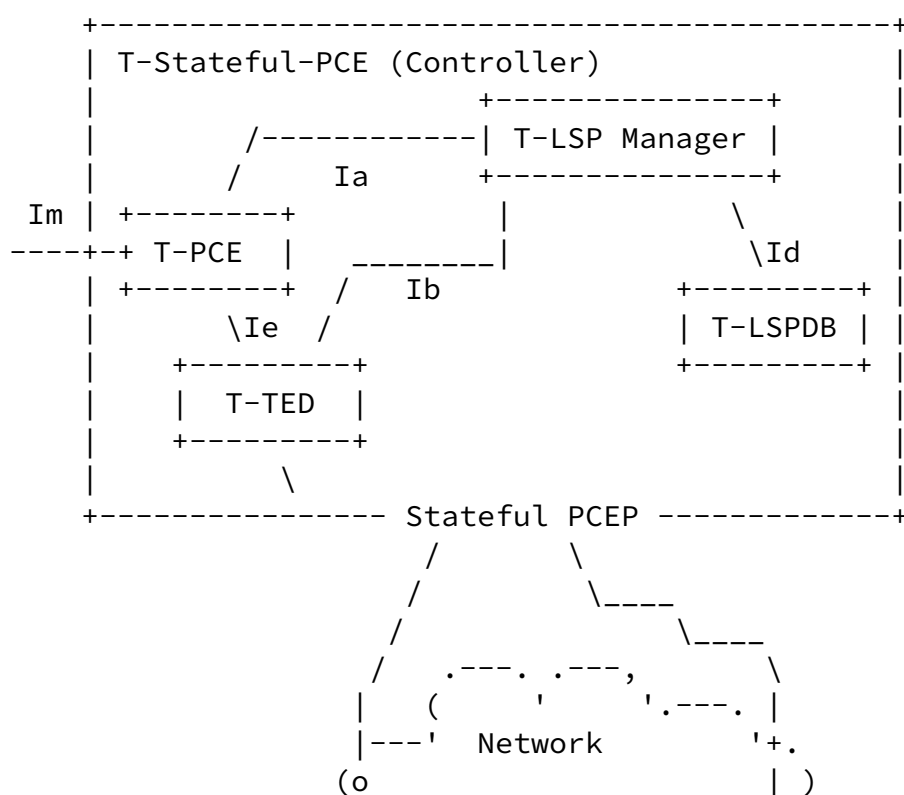
T-LSP manager) asks the T-PCE to compute a path for the LSP.

For computing a path for a temporal LSP crossing multiple domains, the T-PCE communicates with other T-PCEs through interface Im to get an end to end path for the LSP crossing domains. For computing a path for a temporal LSP within the network (one domain), the T-PCE uses a T-CSPF inside it to obtain a path for the LSP.

4.3.3. Temporal Stateful PCE

From the multi-domain hybrid model described in the previous section, we can get a temporal stateful PCE (controller) if we use the stateful PCEP as the interface between the temporal stateful PCE (T-Stateful-PCE for short) controller and the network on which the PCE controls.

Figure below illustrates a temporal stateful PCE controller reference model.



The T-PCE may be outside of the T-Stateful PCE controller. When receiving a request for creating a new temporal LSP with a number of time intervals and some constraints, the T-Stateful PCE (controller) asks the T-PCE to compute a path for the LSP.

For computing a path for a temporal LSP crossing multiple domains, the T-PCE communicates with other T-PCEs through interface *Im* to get an end to end path for the LSP crossing domains. For computing a path for a temporal LSP within the network (one domain), the T-PCE uses a T-CSPF inside it to obtain a path for the LSP.

After obtaining the path for the LSP, the T-Stateful PCE (controller) reserves in the T-TED the TE resources such as link bandwidths for the LSP along the path in each of the time intervals, updates in the T-LSPDB the information about the LSP, initiates the creation of the LSP at the start of each time interval through sending a Path Computation LSP Initiate Request (PCInitiate) message to the ingress of the LSP, and deletes the LSP at the end of each time interval through sending another PCInitiate message with R (remove) flag set to 1.

The T-Stateful PCE (controller) updates the information about the LSP in the T-LSPDB accordingly after receiving a Path Computation LSP State Report (PCRpt) message from the ingress of the LSP.

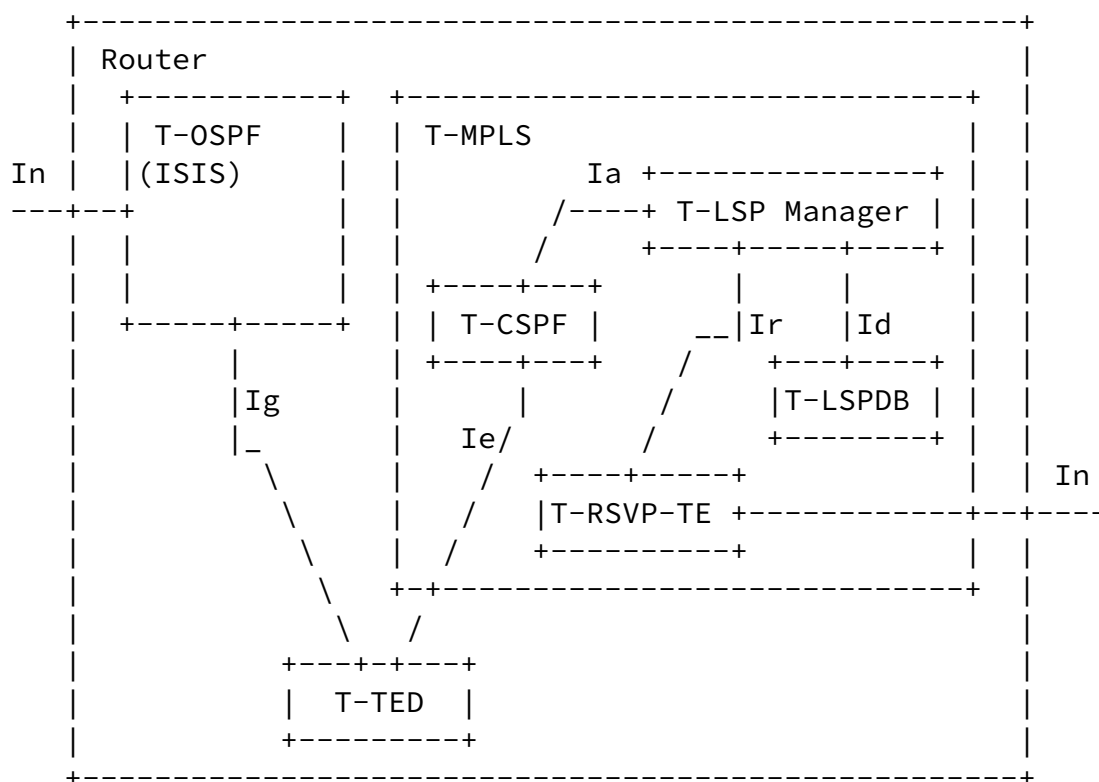
[4.4.](#) Distributed Model

Figure below illustrates a distributed reference model for temporal LSP tunnel services.

Internet-Draft

Framework for TTS

October 2015



In an existing distributed network, the existing MPLS and OSPF/ISIS running on every node in the network need to be enhanced to support temporal LSP tunnel services.

The enhanced OSPF is represented by T-OSPF in the distributed model. The T-OSPF running on a router distributes the traffic engineering (TE) information such as bandwidth about every link connected to the router in a series of time intervals. It also receives the TE information about a link in a number of time intervals from a router in the network. It updates the TE information about every link in the network in the T-TED. The T-OSPF distributes and receives the TE information about a link through interface In connecting to another router.

The T-TED stores the TE information about every link in the network. It is updated by the T-OSPF through interface Ig when the T-OSPF receives the TE information about a link that is changed or the TE information about a link connected to the router is changed.

The T-CSPF in the distributed model has the same function as the one in the other two models. It computes a path for a temporal LSP using the T-TED. It accesses the T-TED through interface Ie.

The enhanced RSVP-TE for supporting temporal LSP tunnel services is represented by T-RSVP-TE. The T-RSVP-TE running on every node along

the path for a temporal LSP signals and maintains the LSP with time intervals. The signaling messages for the LSP is sent and received through interface In connecting to another router.

The T-LSP manager on an LER receives a request to create, delete or update a temporal LSP with a number of time intervals.

When receiving a request for creating a new temporal LSP with a sequence of time intervals and constraints, the T-LSP manager asks the T-CSPF to compute a path for the LSP that satisfies the constraints for the LSP in each of the time intervals.

After obtaining the path for the LSP from the T-CSPF, the T-LSP manager requests T-RSVP-TE to signal the LSP along the path with the time intervals.

The T-LSP manager records the related information for the LSP into the temporal LSP database (T-LSPDB). The information includes the time intervals for the LSP, the path computed for the LSP, and the

status of the LSP.

5. Security Considerations

The mechanism described in this document does not raise any new security issues.

6. Acknowledgement

The authors would like to thank every one who gives his valuable comments on this draft.

7. References

7.1. 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>>.
- [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>>.

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), DOI 10.17487/[RFC3031](#), January 2001, <<http://www.rfc-editor.org/info/rfc3031>>.
- [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>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/[RFC2328](#), April 1998,

<<http://www.rfc-editor.org/info/rfc2328>>.

- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [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-11](#) (work in progress) , April 2015.
- [I-D.ietf-pce-pce-initiated-lsp]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", [draft-ietf-pce-pce-initiated-lsp-04](#) (work in progress) , April 2015.

[7.2](#). 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>>.

- [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>>.

- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol – Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 4875](#), DOI 10.17487/RFC4875, May 2007, <<http://www.rfc-editor.org/info/rfc4875>>.

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
US

Email: huaimo.chen@huawei.com

Mehmet Toy
Comcast
1800 Bishops Gate Blvd.
Mount Laurel, NJ 08054
USA

Email: mehmet_toy@cable.comcast.com

Lei Liu
Fujitsu
USA

Email: lliu@us.fujitsu.com