

Internet Engineering Task Force (IETF)
Request for Comments: 8934
Category: Standards Track
ISSN: 2070-1721

H. Chen, Ed.
Futurewei
Y. Zhuang, Ed.
Q. Wu
Huawei
D. Ceccarelli
Ericsson
October 2020

PCE Communication Protocol (PCEP) Extensions for Label Switched Path (LSP) Scheduling with Stateful PCE

Abstract

This document defines a set of extensions to the stateful PCE Communication Protocol (PCEP) to enable Label Switched Path (LSP) path computation, activation, setup, and deletion based on scheduled time intervals for the LSP and the actual network resource usage in a centralized network environment, as stated in RFC 8413.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8934>.

Copyright Notice

Copyright (c) 2020 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 (<https://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
2. Conventions Used in This Document
 - 2.1. Terminology
3. Motivation and Objectives
4. Procedures and Mechanisms
 - 4.1. LSP Scheduling Overview
 - 4.2. Support of LSP Scheduling
 - 4.2.1. LSP Scheduling
 - 4.2.2. Periodical LSP Scheduling
 - 4.3. Scheduled LSP Creation

- 4.4. Scheduled LSP Modifications
- 4.5. Scheduled LSP Activation and Deletion
- 5. PCEP Objects and TLVs
 - 5.1. Stateful PCE Capability TLV
 - 5.2. LSP Object
 - 5.2.1. SCHED-LSP-ATTRIBUTE TLV
 - 5.2.2. SCHED-PD-LSP-ATTRIBUTE TLV
- 6. The PCEP Messages
 - 6.1. The PCRpt Message
 - 6.2. The PCUpd Message
 - 6.3. The PCInitiate Message
 - 6.4. The PCReq message
 - 6.5. The PCRep Message
 - 6.6. The PCErr Message
- 7. Security Considerations
- 8. Manageability Consideration
 - 8.1. Control of Function and Policy
 - 8.2. Information and Data Models
 - 8.3. Liveness Detection and Monitoring
 - 8.4. Verify Correct Operations
 - 8.5. Requirements on Other Protocols
 - 8.6. Impact on Network Operations
- 9. IANA Considerations
 - 9.1. PCEP TLV Type Indicators
 - 9.1.1. SCHED-PD-LSP-ATTRIBUTE TLV Opt Field
 - 9.1.2. Schedule TLVs Flag Field
 - 9.2. STATEFUL-PCE-CAPABILITY TLV Flag Field
 - 9.3. PCEP-ERROR Object Error Types and Values
- 10. References
 - 10.1. Normative References
 - 10.2. Informative References

Acknowledgments

Contributors

Authors' Addresses

1. Introduction

The PCE Communication Protocol (PCEP) defined in [RFC5440] is used between a Path Computation Element (PCE) and a Path Computation Client (PCC) (or other PCE) to enable path computation of Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs).

[RFC8231] describes a set of extensions to PCEP to provide stateful control. A stateful PCE has access to not only the information carried by the network's Interior Gateway Protocol (IGP) but also the set of active paths and their reserved resources for its computations. The additional state allows the PCE to compute constrained paths while considering individual LSPs and their interactions.

Traditionally, the usage and allocation of network resources, especially bandwidth, can be supported by a Network Management System (NMS) operation such as path pre-establishment. However, this does not provide efficient usage of network resources. The established paths reserve the resources forever, so they cannot be used by other services even when they are not used for transporting any service. [RFC8413] then provides a framework that describes and discusses the problem and defines an appropriate architecture for the scheduled reservation of TE resources.

The scheduled reservation of TE resources allows network operators to reserve resources in advance according to the agreements with their customers and allows them to transmit data about scheduling, such as a specified start time and duration (for example, for a scheduled bulk data replication between data centers). It enables the

activation of bandwidth usage at the time the service is really being used while letting other services use the bandwidth when it is not being used by this service. The requirement of scheduled LSP provisioning is mentioned in [RFC8231] and [RFC7399]. Also, for deterministic networks [RFC8655], the scheduled LSP or temporal LSP can provide better network resource usage for guaranteed links. This idea can also be applied in segment routing [RFC8402] to schedule the network resources over the whole network in a centralized manner.

With this in mind, this document defines a set of needed extensions to PCEP used for stateful PCEs so as to enable LSP scheduling for path computation and LSP setup/deletion based on the actual network resource usage duration of a traffic service. A scheduled LSP is characterized by a start time and a duration. When the end of the LSP life is reached, it is deleted to free up the resources for other LSPs (scheduled or not).

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.1. Terminology

The following terminology is reused from existing PCE documents.

- * Active Stateful PCE [RFC8051]
- * Delegation [RFC8051]
- * PCE-initiated LSP [RFC8281]
- * PCC [RFC5440]
- * PCE [RFC5440]
- * TE LSP [RFC5440]
- * TED (Traffic Engineering Database) [RFC5440]
- * LSP-DB (LSP State Database) [RFC8051]

In addition, this document defines the following terminologies.

Scheduled TE LSP (or Scheduled LSP for short): An LSP with scheduling attributes that carries traffic flow demand at a start time and lasts for a certain duration (or from a start time to an end time, where the end time is the start time plus the duration). A scheduled LSP is also called a "temporal LSP". The PCE operates path computation per LSP availability for the required time and duration.

Scheduled LSP-DB (SLSP-DB): A database of scheduled LSPs.

Scheduled TED: Traffic engineering database with the awareness of scheduled resources for TE. This database is generated by the PCE from the information in the TED and scheduled LSP-DB; it allows knowing, at any time, the expected amount of available resources (discounting the possibility of failures in the future).

Start time (Start-Time): This value indicates when the scheduled LSP is used and the corresponding LSP must be set up and active. At other times (i.e., before the start time or after the start time

plus duration), the LSP can be inactive to include the possibility of the resources being used by other services.

Duration: This value indicates the length of time that the LSP carries a traffic flow and the corresponding LSP must be set up and active. At the end of the duration, the LSP is torn down and removed from the database.

3. Motivation and Objectives

A stateful PCE [RFC8231] can support better efficiency by using LSP scheduling described in the use case of [RFC8051]. This requires the PCE to maintain the scheduled LSPs and their associated resource usage (e.g., bandwidth for packet-switched network) as well as have the ability to trigger signaling for the LSP setup/tear-down at the correct time.

Note that existing configuration tools can be used for LSP scheduling, but as highlighted in Section 3.1.3 of [RFC8231] as well as discussions in [RFC8413], doing this as a part of PCEP in a centralized manner has obvious advantages.

This document provides a set of extensions to PCEP to enable LSP scheduling for LSP creation/deletion under the stateful control of a PCE and according to traffic service requests from customers, so as to improve the usage of network resources.

4. Procedures and Mechanisms

4.1. LSP Scheduling Overview

LSP scheduling allows PCEs and PCCs to provide scheduled LSP for customers' traffic services at its actual usage time, so as to improve the network resource utilization efficiency.

For stateful PCE supporting LSP scheduling, there are two types of LSP databases used in this document. One is the LSP-DB defined in PCEP [RFC8231], while the other is the scheduled LSP database (SLSP-DB). The SLSP-DB records scheduled LSPs and is used in conjunction with the TED and LSP-DB. Note that the two types of LSP databases can be implemented in one physical database or two different databases. This is an implementation matter, and this document does not state any preference.

Furthermore, a scheduled TED can be generated from the scheduled LSP-DB, LSP-DB, and TED to indicate the network links and nodes with resource availability information for now and the future. The scheduled TED MUST be maintained by all PCEs within the network environment.

In the case of implementing PCC-initiated scheduled LSPs, when delegating a scheduled LSP, a PCC MUST include that LSP's scheduling parameters (see Section 5.2.1), including the start time and duration, using a Path Computation State Report (PCRpt) message. Since the LSP is not yet signaled, at the time of delegation, the LSP would be in down state. Upon receiving the delegation of the scheduled LSP, a stateful PCE MUST check whether the parameters are valid. If they are valid, it SHALL check the scheduled TED for the network resource availability on network nodes, compute a path for the LSP with the scheduling information, and update to the PCC as per the active stateful PCE techniques [RFC8231].

Note that the active stateful PCE can update to the PCC with the path for the scheduled LSP at any time. However, the PCC should not signal the LSP over the path after receiving these messages since the path is not active yet; the PCC signals the LSP at the start time.

In the case of multiple PCEs within a single domain, the PCE would need to synchronize their scheduling information with other PCEs within the domain. This could be achieved by proprietary database-synchronization techniques or via a possible PCEP extension (see [PCE-STATE-SYNC]). The technique used to synchronize an SLSP-DB is out of scope for this document. When the scheduling information is out of synchronization among some PCEs, some scheduled LSPs may not be set up successfully.

The scheduled LSP can also be initiated by a PCE itself. In the case of implementing a PCE-initiated scheduled LSP, the stateful PCE SHALL check the network resource availability for the traffic, compute a path for the scheduled LSP, and initiate a scheduled LSP at the PCC and synchronize the scheduled LSP to other PCEs. Note that the PCC could be notified immediately or at the start time of the scheduled LSP, based on the local policy. In the former case, the SCHED-LSP-ATTRIBUTE TLV (see Section 5.2.1) MUST be included in the message, whereas for the latter, the SCHED-LSP-ATTRIBUTE TLV SHOULD NOT be included. Either way, the synchronization to other PCEs MUST be done when the scheduled LSP is created.

In both modes, for activation of scheduled LSPs, the PCC MUST initiate the setup of the scheduled LSP at the start time. Similarly, on the scheduled usage expiry, the PCC MUST initiate the removal of the LSP based on the flag set in the SCHED-LSP-ATTRIBUTE TLV.

4.2. Support of LSP Scheduling

4.2.1. LSP Scheduling

For a scheduled LSP, a user configures it with an arbitrary scheduling period from time T_a to time T_b , which may be represented as $[T_a, T_b]$.

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

4.2.2. Periodical LSP Scheduling

In addition to LSP scheduling at an arbitrary time period, there is also periodical LSP scheduling.

Periodical LSP scheduling means an LSP has multiple time intervals and the LSP is set up to carry traffic in every time interval. It has a scheduling period such as $[T_a, T_b]$, a number of repeats such as 10 (repeats 10 times), and a repeat cycle/time interval such as a week (repeats every week). The scheduling interval " $[T_a, T_b]$ repeats n times with repeat cycle C " represents $n+1$ scheduling 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 scheduling interval such as " $[T_a, T_b]$ repeats 10 times with a repeat cycle of a week" (representing 11 scheduling intervals), a path satisfying the constraints for the LSP in every interval represented by the periodical scheduling interval is computed once. Note that the path computed for one recurrence may be different from the path for another recurrence. And then the LSP along the path is set up to carry traffic in each of the scheduling intervals. If there is no path satisfying the constraints for some of the intervals, the LSP MUST NOT be set up at all. It MUST generate a PCEP Error (PCErr) with Error-Type = 29 (Path computation

failure) and Error-value = 5 (Constraints could not be met for some intervals).

4.2.2.1. Elastic Time LSP Scheduling

In addition to the basic LSP scheduling at an arbitrary time period, another option is elastic time intervals, which is represented as within $-P$ and Q , where P and Q are amounts of time such as 300 seconds. P is called the elastic range lower bound, and Q is called the elastic range upper bound.

For a simple time interval such as $[Ta, Tb]$ with an elastic range, elastic time interval " $[Ta, Tb]$ within $-P$ and Q " means a time period from $(Ta+X)$ to $(Tb+X)$, where $-P \leq X \leq Q$. Note that both Ta and Tb are shifted by the same X . This elastic time interval is suitable for the case where a user wants to have a scheduled LSP up to carry the traffic in time interval $[Ta, Tb]$ and has some flexibility on shifting the time interval a little bit, such as up to P seconds earlier/left or some time such as up to Q seconds later/right.

When an LSP is configured with elastic time interval " $[Ta, Tb]$ within $-P$ and Q ", a path is computed such that the path satisfies the constraints for the LSP in the time period from $(Ta+Xv)$ to $(Tb+Xv)$, and an optimization is performed on Xv from $-P$ to Q . The optimization makes $[Ta+Xv, Tb+Xv]$ the time interval closest to time interval $[Ta, Tb]$ within the elastic range. The LSP along the path is set up to carry traffic in the time period from $(Ta+Xv)$ to $(Tb+Xv)$.

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

$[Ta+X_0, Tb+X_0], [Ta+C+X_1, Tb+C+X_1], \dots, [Ta+nC+X_n, Tb+nC+X_n]$,
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 " $[Ta, Tb]$ 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:

$[Ta+X, Tb+X], [Ta+C+X, Tb+C+X], \dots, [Ta+nC+X, Tb+nC+X]$, where $-P \leq X \leq Q$.

4.2.2.2. Grace Periods

Besides the stated time scheduling, a user may want to have some grace periods (short for "graceful time periods") for each or some of the time intervals for the LSP. Two grace periods may be configured for a time interval. One is the grace period before the time interval, called "Grace-Before", which extends the lifetime of the LSP by an amount of time (such as 30 seconds) before the time interval. The other grace period is after the time interval and is called "Grace-After"; it extends the lifetime of the LSP by an amount of time (such as 60 seconds) after the time interval. Note that no network resources such as link bandwidth will be reserved for the LSP during the grace periods.

When an LSP is configured with a simple time interval such as $[Ta, Tb]$ with grace periods such as Grace-Before GrB and Grace-After GrA , 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 $(Ta-GrB)$ to $(Tb+GrA)$. During grace periods from $(Ta-GrB)$ to Ta and from Tb to $(Tb+GrA)$, the LSP is up to carry traffic in best effort.

4.3. Scheduled LSP Creation

In order to realize PCC-initiated scheduled LSPs in a centralized network environment, a PCC MUST separate the setup of an LSP into two steps. The first step is to request/delegate and get an LSP but not signal it over the network. The second step is to signal the scheduled LSP over the Label Switching Routers (LSRs) at its start time.

For PCC-initiated scheduled LSPs, a PCC MUST delegate the scheduled LSP by sending a PCRpt message by including its demanded resources with the scheduling information to a stateful PCE. Note that the PCC MAY use Path Computation Request (PCReq) and Path Computation Reply (PCRep) messages with scheduling information before delegating.

Upon receiving the delegation via PCRpt message, the stateful PCE MUST compute a path for the scheduled LSP per its start time and duration based on the network resource availability stored in the scheduled TED (see Section 4.1).

The stateful PCE will send a Path Computation Update Request (PCUpd) message with the scheduled path information and the scheduled resource information for the scheduled LSP to the PCC. The stateful PCE MUST update its local scheduled LSP-DB and scheduled TED with the scheduled LSP and would need to synchronize the scheduling information with other PCEs in the domain.

For a PCE-initiated scheduled LSP, the stateful PCE MUST automatically compute a path for the scheduled LSP per requests from network management systems, based on the network resource availability in the scheduled TED, and send an LSP Initiate Request (PCInitiate) message with the path information to the PCC. Based on the local policy, the PCInitiate message could be sent immediately to ask the PCC to create a scheduled LSP (as per this document), or the PCInitiate message could be sent at the start time to the PCC to create a normal LSP (as per [RFC8281]).

For both PCC-initiated and PCE-initiated Scheduled LSPs:

- * The stateful PCE MUST update its local scheduled LSP-DB and scheduled TED with the scheduled LSP.
- * Upon receiving the PCUpd message or PCInitiate message for the scheduled LSP from PCEs with a found path, the PCC determines that it is a scheduled path for the LSP by the SCHED-LSP-ATTRIBUTE TLV (see Section 5.2.1) or SCHED-PD-LSP-ATTRIBUTE TLV (see Section 5.2.2) in the message and does not trigger signaling for the LSP setup on LSRs immediately.
- * The stateful PCE MUST update the scheduled LSP parameters on any network events using the PCUpd message to the PCC. These changes are also synchronized to other PCEs.
- * When it is time for the LSP to be set up (i.e., at the start time), based on the value of the C flag for the scheduled TLV, either the PCC MUST trigger the LSP to be signaled or the delegated PCE MUST send a PCUpd message to the head-end LSR providing the updated path to be signaled (with the A flag set to indicate LSP activation).

4.4. Scheduled LSP Modifications

After a scheduled LSP is configured, a user may change its parameters, including the requested time and the bandwidth. For a periodic-scheduled LSP, its unused recurrences can be modified or

canceled. For a scheduled LSP that is currently active, its duration (the lifetime) can be reduced.

In the PCC-initiated case, the PCC MUST send the PCE a PCRpt message for the scheduled LSP with updated parameters, as well as scheduled information included in the SCHED-LSP-ATTRIBUTE TLV (see Section 5.2.1) or SCHED-PD-LSP-ATTRIBUTE TLV (see Section 5.2.2) carried in the LSP object. The PCE SHOULD take the updated resources and schedule into consideration, and update the new path for the scheduled LSP to the PCC, and synchronize to other PCEs in the network. If the path cannot be set based on new requirements, the previous LSP will not be impacted, and this MUST be conveyed by the use of an empty Explicit Route Object (ERO) in the PCEP messages.

In the PCE-initiated case, the stateful PCE would recompute the path based on updated parameters and scheduled information. If it has already conveyed this information to the PCC by sending a PCInitiate message, it SHOULD update the path and other scheduling and resource information by sending a PCUpd message.

4.5. Scheduled LSP Activation and Deletion

In the PCC-initiated case, when it is time for the LSP to be set up (i.e., at the start time), based on the value of the C flag for the scheduled TLV, either the PCC MUST trigger the LSP to be signaled, or the delegated PCE MUST send a PCUpd message to the head-end LSR providing the updated path to be signaled (with the A flag set to indicate LSP activation). The PCC MUST report the status of the active LSP as per the procedures in [RFC8231], and at this time, the LSP MUST be considered part of the LSP-DB. The A flag MUST be set in the scheduled TLV to indicate that the LSP is active now. After the scheduled duration expires, based on the C flag, the PCC MUST trigger the LSP deletion on itself, or the delegated PCE MUST send a PCUpd message to the PCC to delete the LSP as per the procedures in [RFC8231].

In the PCE-initiated case, based on the local policy, if the scheduled LSP is already conveyed to the PCC at the time of creation, the handling of LSP activation and deletion is handled in the same way as the PCC-initiated case, as per the setting of the C flag. Otherwise, the PCE MUST send the PCInitiate message to the PCC at the start time to create a normal LSP without the scheduled TLV and remove the LSP after the duration expires, as per [RFC8281].

5. PCEP Objects and TLVs

5.1. Stateful PCE Capability TLV

A PCC and a PCE indicate their ability to support LSP scheduling during their PCEP session establishment phase. For an environment with multiple PCEs, the PCEs SHOULD also establish a PCEP session and indicate its ability to support LSP scheduling among PCEP peers. The OPEN object in the Open message contains the STATEFUL-PCE-CAPABILITY TLV. Note that the STATEFUL-PCE-CAPABILITY TLV is defined in [RFC8231] and updated in [RFC8281] and [RFC8232]. In this document, we define a new flag bit B (LSP-SCHEDULING-CAPABILITY) in the Flags field of the STATEFUL-PCE-CAPABILITY TLV to indicate the support of LSP scheduling. We also define another flag bit PD (PD-LSP-CAPABILITY) to indicate the support of LSP periodical scheduling.

B (LSP-SCHEDULING-CAPABILITY) - 1 bit (Bit Position 22): If set to 1 by a PCC, the B flag indicates that the PCC allows LSP scheduling; if set to 1 by a PCE, the B flag indicates that the PCE is capable of LSP scheduling. The B bit MUST be set by both PCEP peers in order to support LSP scheduling for path computation.

PD (PD-LSP-CAPABILITY) - 1 bit (Bit Position - 21): If set to 1 by a PCC, the PD flag indicates that the PCC allows LSP scheduling periodically; if set to 1 by a PCE, the PD flag indicates that the PCE is capable of periodical LSP scheduling. Both the PD bit and the B bit MUST be set to 1 by both PCEP peers in order to support periodical LSP scheduling for path computation. If the PD bit or B bit is 0, then the periodical LSP scheduling capability MUST be ignored.

5.2. LSP Object

The LSP object is defined in [RFC8231]. This document adds an optional SCHED-LSP-ATTRIBUTE TLV for normal LSP scheduling and an optional SCHED-PD-LSP-ATTRIBUTE TLV for periodical LSP scheduling. The LSP object for a scheduled LSP MUST NOT include these two TLVs. Only one scheduling, either normal or periodical, is allowed for a scheduled LSP.

The presence of the SCHED-LSP-ATTRIBUTE TLV in the LSP object indicates that this LSP is normal scheduling while the SCHED-PD-LSP-ATTRIBUTE TLV indicates that this scheduled LSP is periodical. The SCHED-LSP-ATTRIBUTE TLV MUST be present in the LSP object for each normal-scheduled LSP carried in the PCEP messages. The SCHED-PD-LSP-ATTRIBUTE TLV MUST be used in the LSP object for each periodic-scheduled LSP carried in the PCEP messages.

Only one SCHED-LSP-ATTRIBUTE TLV SHOULD be present in the LSP object. If more than one SCHED-LSP-ATTRIBUTE TLV is found, the first instance is processed and others ignored. The SCHED-PD-LSP-ATTRIBUTE TLV is the same as the SCHED-LSP-ATTRIBUTE TLV with regard to its presence in the LSP object.

5.2.1. SCHED-LSP-ATTRIBUTE TLV

The SCHED-LSP-ATTRIBUTE TLV MAY be included as an optional TLV within the LSP object for LSP scheduling for the requesting traffic service.

This TLV MUST NOT be included unless both PCEP peers have set the B (LSP-SCHEDULING-CAPABILITY) bit in the STATEFUL-PCE-CAPABILITY TLV carried in the Open message to one. If the TLV is received by a peer when both peers didn't set the B bit to one, the peer MUST generate a PCEP Error (PCerr) with a PCEP-ERROR object having Error-Type = 19 (Invalid Operation) and Error-value = 15 (Attempted LSP scheduling while the scheduling capability was not advertised).

The format of the SCHED-LSP-ATTRIBUTE TLV is shown in Figure 1.

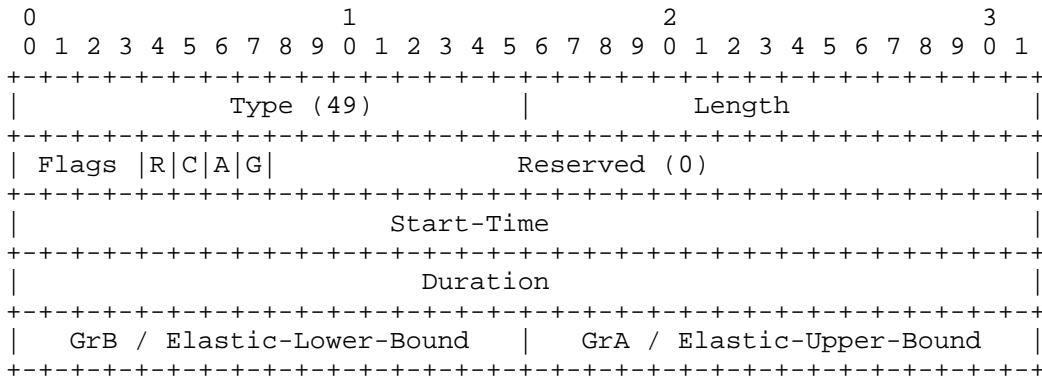


Figure 1: SCHED-LSP-ATTRIBUTE TLV

The type of the TLV is 49, and the TLV has a fixed length of 16 octets.

The fields in the format are:

Flags (8 bits): The following flags are defined in this document.

R (1 bit): Set to 1 to indicate that the Start-Time is a relative time, which is the number of seconds from the current time. The PCEs and PCCs MUST synchronize their clocks when relative time is used. It is RECOMMENDED that the Network Time Protocol [RFC5905] be used to synchronize clocks among them. When the transmission delay from a PCE or PCC to another PCE or PCC is too big (such as greater than 1 second), the scheduling interval represented is not accurate if the delay is not considered. Set to 0 to indicate that the 32-bit Start-Time is an absolute time, which is the number of seconds since the epoch. The epoch is 1 January 1970 at 00:00 UTC. It wraps around every 2^{32} seconds, which is roughly 136 years. The next wraparound will occur in the year 2106. The received Start-Time is considered after the wraparound if the resulting value is less than the current time. In that case, the value of the 32-bit Start-Time is considered to be the number of seconds from the time of wraparound (because the Start-Time is always a future time).

C (1 bit): Set to 1 to indicate that the PCC is responsible to set up and remove the scheduled LSP based on the Start-Time and Duration. The PCE holds these responsibilities when the bit is set to zero.

A (1 bit): Set to 1 to indicate that the scheduled LSP has been activated.

G (1 bit): Set to 1 to indicate that the grace period is included in the fields GrB/Elastic-Lower-Bound and GrA/Elastic-Upper-Bound; set to 0 to indicate that the elastic range is included in the fields.

Reserved (24 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

Start-Time (32 bits): This value, in seconds, indicates when the scheduled LSP is used to carry traffic and the corresponding LSP MUST be set up and activated. Note that the transmission delay SHOULD be considered when R=1 and the value of Start-Time is small.

Duration (32 bits): This value, in seconds, indicates the duration that the LSP carries a traffic flow and the corresponding LSP MUST be up to carry traffic. At the expiry of this duration, the LSP MUST be torn down and deleted. A value of 0 MUST NOT be used in Duration since it does not make any sense. The value of Duration SHOULD be greater than a constant MINIMUM-DURATION seconds, where MINIMUM-DURATION is 5.

Start-Time indicates a time at or before which the scheduled LSP MUST be set up. When the R bit is set to 0, the value of Start-Time represents the number of seconds since the epoch. When the R bit is set to 1, the value of Start-Time represents the number of seconds from the current time.

In addition, the SCHED-LSP-ATTRIBUTE TLV contains the G flag set to 1 and a nonzero Grace-Before and Grace-After in the fields GrB/Elastic-Lower-Bound and GrA/Elastic-Upper-Bound if grace periods are configured. It includes the G flag set to 0 and a nonzero elastic range lower bound and upper bound in the fields if there is an elastic range configured. A TLV can configure a nonzero grace period or elastic range, but it MUST NOT provide both for an LSP.

GrB (Grace-Before, 16 bits): The grace period time length, in seconds, before the start time.

GrA (Grace-After, 16 bits): The grace period time length, in seconds, after time interval [start time, start time + duration].

Elastic-Lower-Bound (16 bits): The maximum amount of time, in seconds, that the time interval can shift lower/left.

Elastic-Upper-Bound (16 bits): The maximum amount of time, in seconds, that the time interval can shift higher/right.

5.2.2. SCHED-PD-LSP-ATTRIBUTE TLV

The periodical LSP is a special case of LSP scheduling. The traffic service happens in a series of repeated time intervals. The SCHED-PD-LSP-ATTRIBUTE TLV can be included as an optional TLV within the LSP object for this periodical LSP scheduling.

This TLV MUST NOT be included unless both PCEP peers have set the B (LSP-SCHEDULING-CAPABILITY) bit and PD (PD-LSP-CAPABILITY) bit in STATEFUL-PCE-CAPABILITY TLV carried in Open message to one. If the TLV is received by a peer when either bit is zero (or both bits are zero), the peer MUST generate a PCEP Error (PCErr) with a PCEP-ERROR object having Error-Type = 19 (Invalid Operation) and Error-value = 15 (Attempted LSP scheduling while the scheduling capability was not advertised).

The format of the SCHED-PD-LSP-ATTRIBUTE TLV is shown in Figure 2.

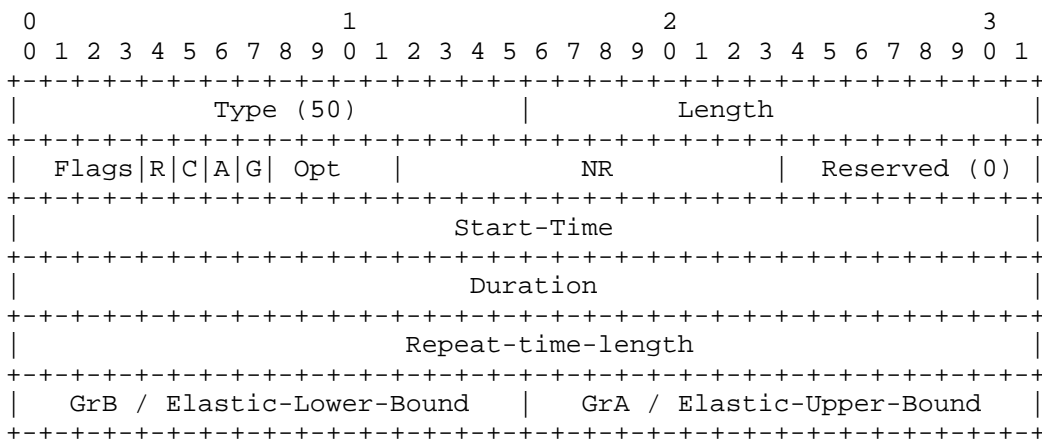


Figure 2: SCHED-PD-LSP-ATTRIBUTE TLV

The type of the TLV is 50, and the TLV has a fixed length of 20 octets. The description, format, and meaning of the flags (R, C, A, and G bits), Start-Time, Duration, GrB, GrA, Elastic-Lower-Bound, and Elastic-Upper-Bound fields remain the same as in the SCHED-LSP-ATTRIBUTE TLV.

The following fields are new:

Opt (4 bits): Indicates options to repeat. When a PCE receives a

TLV with an unknown Opt value, it does not compute any path for the LSP. It MUST generate a PCEP Error (PCErr) with a PCEP-ERROR object having Error-Type = 4 (Not supported object) and Error-value = 4 (Unsupported parameter).

Opt = 1: repeat every month

Opt = 2: repeat every year

Opt = 3: repeat every Repeat-time-length

A user may configure a Repeat-time-length in time units weeks, days, hours, minutes, and/or seconds. The value represented by these units is converted to the number of seconds in the TLV. For example, repeat every 2 weeks is equivalent to repeat every Repeat-time-length = $2 * 7 * 86,400$ (seconds), where 86,400 is the number of seconds per day.

NR (12 bits): The number of repeats. During each repetition, LSP carries traffic.

Reserved (8 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

Repeat-time-length (32 bits): The time in seconds between the Start-Time of one repetition and the Start-Time of the next repetition.

6. The PCEP Messages

6.1. The PCRpt Message

The Path Computation State Report (PCRpt) message is a PCEP message sent by a PCC to a PCE to report the status of one or more LSPs, as per [RFC8231]. Each LSP State Report in a PCRpt message contains the actual LSP's path, bandwidth, operational and administrative status, etc. An LSP Status Report carried in a PCRpt message is also used in delegation or revocation of control of an LSP to/from a PCE. In the case of a scheduled LSP, a scheduled TLV MUST be carried in the LSP object, and the ERO conveys the intended path for the scheduled LSP. The scheduled LSP MUST be delegated to a PCE.

6.2. The PCUpd Message

The Path Computation Update Request (PCUpd) message is a PCEP message sent by a PCE to a PCC to update LSP parameters on one or more LSPs, as per [RFC8231]. Each LSP Update Request in a PCUpd message contains all LSP parameters that a PCE wishes to be set for a given LSP. In the case of a scheduled LSP, a scheduled TLV MUST be carried in the LSP object, and the ERO conveys the intended path for the scheduled LSP. If no path can be found, an empty ERO is used. The A bit is used in the PCUpd message to indicate the activation of the scheduled LSP if the PCE is responsible for the activation (as per the C bit).

6.3. The PCInitiate Message

The LSP Initiate Request (PCInitiate) message is a PCEP message sent by a PCE to a PCC to trigger LSP instantiation or deletion, as per [RFC8281]. In the case of a scheduled LSP, based on the local policy, the PCE MAY convey the scheduled LSP to the PCC by including a scheduled TLV in the LSP object. Alternatively, the PCE would initiate the LSP only at the start time of the scheduled LSP, as per [RFC8281], without the use of scheduled TLVs.

6.4. The PCReq message

The Path Computation Request (PCReq) message is a PCEP message sent by a PCC to a PCE to request a path computation [RFC5440], and it may contain the LSP object [RFC8231] to identify the LSP for which the path computation is requested. In the case of a scheduled LSP, a scheduled TLV MUST be carried in the LSP object in the PCReq message to request the path computation based on the scheduled TED and LSP-DB. A PCC MAY use the PCReq message to obtain the scheduled path before delegating the LSP. The parameters of the LSP may be changed (refer to Section 4.4).

6.5. The PCRep Message

The Path Computation Reply (PCRep) message is a PCEP message sent by a PCE to a PCC in reply to a path computation request [RFC5440], and it may contain the LSP object [RFC8231] to identify the LSP for which the path is computed. A PCRep message can contain either a set of computed paths if the request can be satisfied or a negative reply if not. A negative reply may indicate the reason why no path could be found. In the case of a scheduled LSP, a scheduled TLV MUST be carried in the LSP object in PCRep message to indicate the path computation based on the scheduled TED and LSP-DB. A PCC and PCE MAY use PCReq and PCRep messages to obtain the scheduled path before delegating the LSP.

6.6. The PCErr Message

The PCEP Error (PCErr) message is a PCEP message, as described in [RFC5440], for error reporting. This document defines new error values for several error types to cover failures specific to scheduling and reuses the applicable error types and error values of [RFC5440] and [RFC8231] wherever appropriate.

The PCEP extensions for scheduling MUST NOT be used if one or both of the PCEP speakers have not set the corresponding bits in the STATEFUL-PCE-CAPABILITY TLV in their respective Open messages to one. If the PCEP speaker supports the extensions of this specification but did not advertise this capability, then upon receipt of LSP object with the scheduled TLV, it MUST generate a PCEP Error (PCErr) with Error-Type = 19 (Invalid Operation) and Error-value = 15 (Attempted LSP scheduling while the scheduling capability was not advertised), and it SHOULD ignore the TLV. As per Section 7.1 of [RFC5440], a legacy PCEP implementation that does not understand this specification would consider a scheduled TLV unknown and ignore it.

If the PCC decides that the scheduling parameters proposed in the PCUpd/PCInitiate message are unacceptable, it MUST report this error by including the LSP-ERROR-CODE TLV (Section 7.3.3 of [RFC8231]) with LSP Error-value = 4 (Unacceptable parameters) in the LSP object (with the scheduled TLV) in the PCRep message to the PCE.

The scheduled TLV MUST be included in the LSP object for the scheduled LSPs. If the TLV is missing, the receiving PCEP speaker MUST send a PCErr message with Error-Type = 6 (Mandatory Object missing) and Error-value = 16 (Scheduled TLV missing).

7. Security Considerations

This document defines the LSP-SCHEDULING-CAPABILITY TLV and SCHED-LSP-ATTRIBUTE TLV; the security considerations discussed in [RFC5440], [RFC8231], and [RFC8281] continue to apply. In some deployments, the scheduling information could provide details about the network operations that could be deemed extra sensitive. Additionally, snooping of PCEP messages with such data or using PCEP messages for network reconnaissance may give an attacker sensitive information about the operations of the network. A single PCEP message can now instruct a PCC to set up and tear down an LSP every

second for a number of times. That single message could have a significant effect on the network. Thus, such deployments SHOULD employ suitable PCEP security mechanisms like TCP Authentication Option (TCP-AO), which is discussed in [RFC5925] and [RFC8253]. Note that [RFC8253] is considered a security enhancement and thus is much better suited for sensitive information. PCCs may also need to apply some form of rate limit to the processing of scheduled LSPs.

8. Manageability Consideration

8.1. Control of Function and Policy

The LSP scheduling feature MUST be controlled per tunnel by the active stateful PCE. The values for parameters like start time and duration SHOULD be configurable by customer applications and based on the local policy at PCE. The suggested default values for start time and duration are one day (in seconds) from the current time and one year (in seconds), respectively. One day has 86,400 seconds. One year has 31,536,000 seconds.

When configuring the parameters for time, a user SHOULD consider leap years and leap seconds. If a scheduled LSP has a time interval containing a leap year, the duration of the LSP is 366 days plus the rest of the interval.

8.2. Information and Data Models

An implementation SHOULD allow the operator to view the information about each scheduled LSP defined in this document. To serve this purpose, the PCEP YANG module [PCE-PCEP-YANG] could be extended.

8.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

8.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440]. An implementation SHOULD allow a user to view information, including the status of a scheduled LSP, through a Command Line Interface (CLI) tool. In addition, it SHOULD check and handle the cases where there is a significant time correction or a clock skew between PCC and PCE.

8.5. Requirements on Other Protocols

Mechanisms defined in this document do not imply any new requirements on other protocols.

8.6. Impact on Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [RFC5440].

9. IANA Considerations

9.1. PCEP TLV Type Indicators

IANA maintains the "PCEP TLV Type Indicators" subregistry within the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA has made the following allocations in this subregistry for the new PCEP TLVs defined in this document.

Value	Description	Reference
49	SCHED-LSP-ATTRIBUTE	This document
50	SCHED-PD-LSP-ATTRIBUTE	This document

Table 1: Additions to PCEP TLV Type Indicators Subregistry

9.1.1. SCHED-PD-LSP-ATTRIBUTE TLV Opt Field

IANA has created and will maintain a new subregistry named "SCHED-PD-LSP-ATTRIBUTE TLV Opt Field" within the "Path Computation Element Protocol (PCEP) Numbers" registry. Initial values for the subregistry are given below. New values are assigned by Standards Action [RFC8126].

Value	Description	Reference
0	Reserved	
1	REPEAT-EVERY-MONTH	This document
2	REPEAT-EVERY-YEAR	This document
3	REPEAT-EVERY-REPEAT-TIME-LENGTH	This document
4-14	Unassigned	
15	Reserved	

Table 2: New SCHED-PD-LSP-ATTRIBUTE TLV Opt Field Subregistry

9.1.2. Schedule TLVs Flag Field

IANA has created a new subregistry named "Schedule TLVs Flag Field" within the "Path Computation Element Protocol (PCEP) Numbers" registry. New values are assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (counting from bit 0 as the most significant bit)
- * Capability description
- * Defining RFC

The following values are defined in this document:

Bit	Description	Reference
0-3	Unassigned	
4	Relative Time (R-bit)	This document
5	PCC Responsible (C-bit)	This document
6	LSP Activated (A-bit)	This document
7	Grace Period Included (G-bit)	This document

Table 3: New Schedule TLVs Flag Field Subregistry

9.2. STATEFUL-PCE-CAPABILITY TLV Flag Field

This document defines new bits in the Flags field in the STATEFUL-PCE-CAPABILITY TLV in the OPEN object. IANA maintains the "STATEFUL-PCE-CAPABILITY TLV Flag Field" subregistry within the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA has made the following allocations in this subregistry.

Bit	Description	Reference
22	LSP-SCHEDULING-CAPABILITY (B-bit)	This document
21	PD-LSP-CAPABILITY (PD-bit)	This document

Table 4: Additions to STATEFUL-PCE-CAPABILITY TLV Flag Field Subregistry

9.3. PCEP-ERROR Object Error Types and Values

IANA has allocated the following new error types to the existing error values within the "PCEP-ERROR Object Error Types and Values" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry:

Error-Type	Meaning	Error-value
6	Mandatory Object missing	16: Scheduled TLV missing
19	Invalid Operation	15: Attempted LSP scheduling while the scheduling capability was not advertised
29	Path computation failure	5: Constraints could not be met for some intervals

Table 5: Additions to PCEP-ERROR Object Error Types and Values Subregistry

10. References

10.1. 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, <<https://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, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.

- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", RFC 8232, DOI 10.17487/RFC8232, September 2017, <<https://www.rfc-editor.org/info/rfc8232>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.
- [RFC8413] Zhuang, Y., Wu, Q., Chen, H., and A. Farrel, "Framework for Scheduled Use of Resources", RFC 8413, DOI 10.17487/RFC8413, July 2018, <<https://www.rfc-editor.org/info/rfc8413>>.

10.2. Informative References

- [PCE-PCEP-YANG]
Dhody, D., Hardwick, J., Beeram, V. P., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-15, 31 October 2020, <<https://tools.ietf.org/html/draft-ietf-pce-pcep-yang-15>>.
- [PCE-STATE-SYNC]
Litkowski, S., Sivabalan, S., Li, C., and H. Zheng, "Inter Stateful Path Computation Element (PCE) Communication Procedures.", Work in Progress, Internet-Draft, draft-litkowski-pce-state-sync-08, 12 July 2020, <<https://tools.ietf.org/html/draft-litkowski-pce-state-sync-08>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<https://www.rfc-editor.org/info/rfc5925>>.
- [RFC7399] Farrel, A. and D. King, "Unanswered Questions in the Path Computation Element Architecture", RFC 7399, DOI 10.17487/RFC7399, October 2014, <<https://www.rfc-editor.org/info/rfc7399>>.
- [RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", RFC 8051, DOI 10.17487/RFC8051, January 2017, <<https://www.rfc-editor.org/info/rfc8051>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the

Path Computation Element Communication Protocol (PCEP)",
RFC 8253, DOI 10.17487/RFC8253, October 2017,
<<https://www.rfc-editor.org/info/rfc8253>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L.,
Decraene, B., Litkowski, S., and R. Shakir, "Segment
Routing Architecture", RFC 8402, DOI 10.17487/RFC8402,
July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", RFC 8655,
DOI 10.17487/RFC8655, October 2019,
<<https://www.rfc-editor.org/info/rfc8655>>.

Acknowledgments

The authors of this document would also like to thank Rafal Szarecki,
Adrian Farrel, and Cyril Margaria for the review and comments.

Contributors

Dhruv Dhody
Huawei
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: dhruv.ietf@gmail.com

Xufeng Liu
Ericsson
United States of America

Email: xliu@kuatrotech.com

Mehmet Toy
Verizon
United States of America

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
United States of America

Email: lliu@us.fujitsu.com

Khuzema Pithewan
Infinera

Email: kpithewan@infinera.com

Zitao Wang
Huawei
101 Software Avenue, Yuhua District
Nanjing
Jiangsu, 210012
China

Email: wangzitao@huawei.com

Xian Zhang
Huawei Technologies
Research Area F3-1B
Huawei Industrial Base
Shenzhen
518129
China

Email: zhang.xian@huawei.com

Authors' Addresses

Huaimo Chen (editor)
Futurewei
Boston, MA
United States of America

Email: huaimo.chen@futurewei.com

Yan Zhuang (editor)
Huawei
Yuhua District
101 Software Avenue
Nanjing
Jiangsu, 210012
China

Email: zhuangyan.zhuang@huawei.com

Qin Wu
Huawei
Yuhua District
101 Software Avenue
Nanjing
Jiangsu, 210012
China

Email: bill.wu@huawei.com

Daniele Ceccarelli
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: daniele.ceccarelli@ericsson.com