Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2017

Z. Li
D. Dhody
Huawei Technologies
C. Margaria
C. Barth
Juniper
X. Chen
S. Zhuang
Huawei Technologies

July 8, 2016

PCEP Extension for Flow Specification
draft-li-pce-pcep-flowspec-01

Abstract

Dissemination of the traffic flow specifications was first introduced in the BGP protocol via RFC 5575. In order to distribute the flow specifications from PCE controller to network device without BGP protocol it is desirable to extend PCEP with flow specification information.

This document specifies a set of extensions to PCEP to support dissemination of flow specifications. The extensions include the instantiation, updation and deletion of flow specifications.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 January 9, 2017.

Table of Contents

1. Introduction

Dissemination of the traffic flow specifications was first introduced in the BGP protocol [RFC5575]. The traffic flow specification is comprised of traffic filtering rules and actions. The routers which received the flow specification can take advantage of the ACL (Access Control List) or firewall capabilities in the router's forwarding path. The routers can classify the packets according to the traffic filtering rules and shape, rate limit, filter, or redirect packets based on the actions. The flow specification carried by BGP can be used to automate inter-domain coordination of traffic filtering to mitigate (distributed) denial-of-service attacks and can also be used to provide traffic filtering in the context of a BGP/MPLS VPN service.

[RFC5575] also defines that a flow specification received from an external autonomous system will need to be validated against unicast routing before being accepted. [I-D.ietf-idr-bgp-flowspec-oid] describes a modification to the validation procedure defined in [I-D.ietf-idr-bgp-flowspec-oid] for the dissemination of BGP flow specifications. The modification proposed enables flow specifications to be originated from a centralized BGP route controller.

[I-D.ietf-ospf-flowspec-extensions] defines the extensions to OSPF to distribute flow specifications in the networks that only deploy an IGP (Interior Gateway Protocol) (e.g., OSPF). It also defines the validation procedures for imposing the filtering information on the routers.

[RFC5440] describes the Path Computation Element Protocol (PCEP). PCEP defines the communication between a Path Computation Client (PCC) and a Path Control Element (PCE), or between PCE and PCE, enabling computation of Multiprotocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP) characteristics.

Stateful pce [I-D.ietf-pce-stateful-pce] specifies a set of extensions to PCEP to enable stateful control of TE LSPs between and across PCEP sessions in compliance with [RFC4657]. It includes mechanisms to effect LSP state synchronization between PCCs and PCEs, delegation of control of LSPs to PCEs, and PCE control of timing and sequence of path computations within and across PCEP sessions and focuses on a model where LSPs are configured on the PCC and control over them is delegated to the PCE. [I-D.ietf-pce-pce-initiated-lsp] describes the setup, maintenance and teardown of PCE- initiated LSPs under the stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network that is centrally controlled and deployed.

In case PCE is used to initiate tunnels via PCEP, it is desirable to use the same protocol to also distribute the flow specifications to describe what data flows on those tunnels. Thus, in order to distribute the flow specifications from PCE controller to network device, PCEP is extended with flow specification information in this document.

[I-D.zhao-teas-pce-control-function] introduces the architecture for PCE as a central controller and describes how PCE can be viewed as a component that perfor computation to place 'flows' within the network and decide how these flows are routed.

This document specifies a set of extensions to PCEP to support dissemination of flow specifications. The flow specifications can be disseminated between PCEP peers such as from PCE to PCC or between PCEs . The extensions include the creation, updation and withdrawal of flow specifications via PCEP.

The values of flow filtering rules and actions mainly refer to the BGP flow specification and IGP specification. This document extends new actions which are redirecting to LSP (refered by Symbolic Path Name, IPv4 LSP, or IPv6 LSP).

2. Terminology

This document uses the terms defined in [RFC5440] and [RFC5575].

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

The following term is from [RFC5575]. It is used frequently throughout this document:

Flow Specification (FlowSpec): A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic, including filters and actions. Each FlowSpec consists of a set of filters and a set of actions.

3. Procedures for Dissemination of FlowSpec

3.1. Overview of Procedures

A PCC or PCE indicates its ability to support PCE FlowSpec during the PCEP Initialization Phase via "PCE FlowSpec Capability" TLV (see details in Section 5.1.1).

This section introduces the procedure to support PCE FlowSpec as follows:

Firstly both the PCE and PCC advertise the PCE FlowSpec Capability during the PCE session initiation phase.

On the PCEP session with PCE FlowSpec Capability PCE communicates with PCC to create, update and withdraw PCE FlowSpec.

[Editor's Note - The procedure about PCE FlowSpec synchronization, the session failure process, etc. will be specified in the future version.]

3.2. Capability Advertisement

During PCEP session establishment, both the PCC and the PCE must announce their support of PCEP extensions for FlowSpec defined in this document.

A PCEP Speaker (PCE or PCC) includes the "PCE FlowSpec Capability" TLV, described in Section 5.1.1, in the OPEN Object to advertise its support for PCEP extensions for PCE FlowSpec Capability.

The presence of the PCE FlowSpec Capability TLV in PCE's OPEN message indicates that the PCE can support distribute the FlowSpec to PCC.

The presence of such Capability TLV in PCC's OPEN Object indicates that the PCC can be in support of Flowspec functionality to instantiate the FlowSpec according to the PCE's indication and can apply the FlowSpec to the incoming packets.

If PCE has such capability TLV and PCC has no such capability TLV PCE MUST NOT send the PCE messages with FlowSpec information. And if PCC receives such messages it should send PCErr message to PCE.

[Editor's Note - PCE discovery via IGP should also be extended for this.]

3.3. Operations

To instantiate a FlowSpec which is comprised of a set of FlowSpec filter rules and actions, the PCE sends a new PCEP message (called FlowSpec message) to the PCC. The FlowSpec message MUST include the SRP object[I-D.ietf-pce-stateful-pce], a new FLOW object (see details in Section 5.2) and a new ACTION object (see details in Section 5.3). FLOW object carries a set of FlowSpec filter rules. A list of ACTION objects specify a set of FlowSpec actions.

To update the FlowSpec actions of a specified FlowSpec which has been created, the same PCEP message "FlowSpec" is used. The PCE sends a

FlowSpec message to the PCC. The FlowSpec message MUST include the SRP object, FLOW object and ACTION object.

To delete the specified FlowSpec which has been created, the PCE sends a FlowSpec message to the PCC with a flag indicating the removal action. The FlowSpec message MUST include the SRP object (with R flag set) and FLOW object.

3.4. Flowspec Synronization

[I-D.kuppani-pce-pcep-flowspec-sync] specify the flow specification synchronization mechanism for managing of flow specification (FLOWSPEC-DB) at node (PCC) aligning with FLOWSPEC-DB at PCE on initial session UP or session flap and specifies the required Path Computation Element Communication Protocol (PCEP) extensions. This includes full synchronization as well as optimizations such as synchronization avoidance and incremental synchronization.

4. PCEP Messages

As defined in [RFC5440], a PCEP message consists of a common header followed by a variable-length body made of a set of objects that can be either mandatory or optional. An object is said to be mandatory in a PCEP message when the object must be included for the message to be considered valid. For each PCEP message type, a set of rules is defined that specify the set of objects that the message can carry. An implementation MUST form the PCEP messages using the object ordering specified in this document.

To support the PCEP FlowSpec functionality one new PCEP messages is introduced.

4.1. PCEP FlowSpec Message

A FlowSpec message which is also referred to as FlowSpec message is a PCEP message sent by a PCE to a PCC to trigger creation, modification or deletion of a FlowSpec.

The Message-Type field of the PCEP common header for the FlowSpec message is TBD17 (to be assigned by IANA). The FlowSpec message MUST include the SRP and the FLOW objects.

If FlowSpec message is used to create or update the FlowSpec, it MUST include the ACTION objects too.

If FlowSpec message is used to delete the FlowSpec the ACTION objects SHOULD NOT be carried and the SRP object is set with the R flag.

A FlowSpec is identified by a PCEP specific identifier FS-ID.

The format of a FlowSpec message for creation or deletion of FlowSpec is as follows:

```
   <FlowSpec Message> ::= <Common Header>
                          <flowspec-list>
Where:
   <flowspec-list> ::= <flowspec-request>[<flowspec-list>]

   <flowspec-request>::= (<flowspec-create-or-update>|
                          <flowspec-delete>)

   <flowspec-create-or-update> ::= <SRP>
                                   <FLOW>
                                   <action-list>

   <flowspec-delete> ::= <SRP>
                          <FLOW>

Where:
    <action-list>::=<ACTION>[<action-list>]
```

The SRP object defined in [I-D.ietf-pce-stateful-pce] can be used in this document to correlate FlowSpec requests sent by the PCE with the error reports sent by the PCC.

Every FlowSpec requests from the PCE sends a new SRP-ID-NUMBER as described in [I-D.ietf-pce-stateful-pce]. This number is unique per PCEP session and is incremented each time an FlowSpec operation (creation, update, deletion etc) is requested from the PCE. The value of the SRP-ID-NUMBER MAY be echoed back by the PCC in PCErr messages to allow for correlation between requests made by the PCE and errors generated by the PCC. Procedure of dissemination of FlowSpec from PCE share the same number space of the SRP-ID-NUMBER with procedure of stateful PCE.

The FLOW and ACTION objects are new objects introduced in this document.

5. Objects and TLVs

The PCEP objects defined in this document are compliant with the PCEP object format defined in [RFC5440].

New TLVs about FlowSpec filtering rules are defined. The value portion of the new TLVs can reuse the structure defined in [RFC5575] and [I-D.ietf-idr-flow-spec-v6]. New TLVs about FlowSpec actions are

also defined. The value portion of the new TLVs can reuse the structure defined in [I-D.ietf-ospf-flowspec-extensions]. This document also defines two new actions: Redirect to IPv4 LSP and Redirect to IPv6 LSP.

5.1. OPEN Object

5.1.1. PCE FlowSpec Capability TLV

The PCE-FLOWSPEC-CAPABILITY TLV is an optional TLV associated with the OPEN Object [RFC5440] to exchange PCE FlowSpec capability of PCEP speakers.

Its format is shown in the following figure:

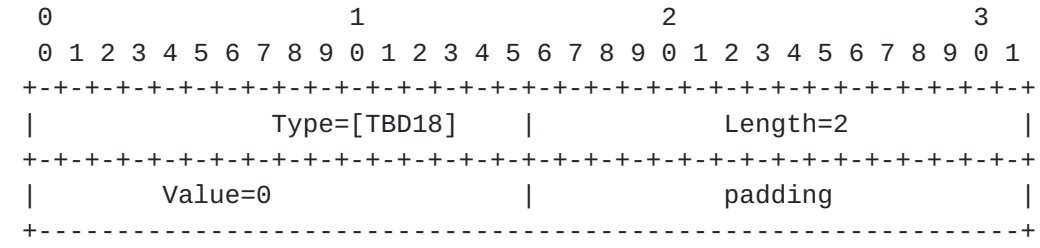

Figure 1: PCE-FLOWSPEC-CAPABILITY TLV format

The type of the TLV is TBD18 (to be assigned by IANA) and it has a fixed length of 2 octets. The value field is set to default value 0.

The inclusion of this TLV in an OPEN object indicate that the sender can perform FlowSpec handling in PCEP.

5.2. FLOW Object

The FLOW object MUST be present within FlowSpec messages. The FLOW object carries a set of FlowSpec filter rules.

FLOW Object-Class is TBD19 (to be assigned by IANA).

FLOW Object-Type is 1.

The format of the FLOW object is as follows:

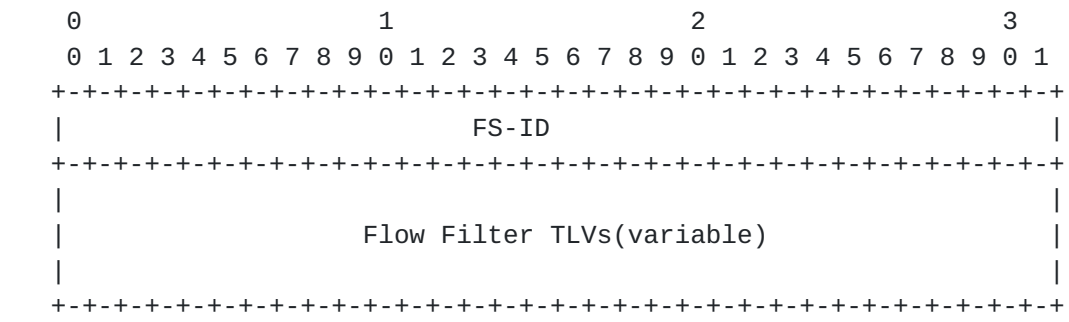

Figure 2: FLOW Object Body Format

FS-ID(32-bit): A PCEP-specific identifier for the FlowSpec information. A PCE creates an unique FS-ID for each FlowSpec that is constant for the lifetime of a PCEP session. All subsequent PCEP messages then address the FlowSpec by the FS-ID. The values of 0 and 0xFFFFFFFF are reserved.

Flow Filter TLVs(variable): The FLOW object body has a variable length and may contain one or more additional TLVs.

The following flow filter types are supported:

Type	Description	Ref TLV	Value defined in
TBD1	Destination IPv4 Prefix	1	RFC5575
TBD2	Source IPv4 Prefix	2	RFC5575
TBD3	IP Protocol	3	RFC5575
TBD4	Port	4	RFC5575
TBD5	Destination port	5	RFC5575
TBD6	Source port	6	RFC5575
TBD7	ICMP type	7	RFC5575
TBD8	ICMP code	8	RFC5575
TBD9	TCP flags	9	RFC5575
TBD10	Packet length	10	RFC5575
TBD11	DSCP	11	RFC5575
TBD12	Fragment	12	RFC5575
TBD13	Flow Label	13	I-D.ietf-idr-flow-spec-v6
TBD14	Destination IPv6 Prefix	1	I-D.ietf-idr-flow-spec-v6
TBD15	Source IPv6 Prefix	2	I-D.ietf-idr-flow-spec-v6
TBD16	Next Header	3	I-D.ietf-idr-flow-spec-v6
*	ROUTE-DISTINGUISHER	-	I-D.dhodylee-pce-pcep-ls

(*) - TLV is defined in another PCEP document.

Figure 3: Table of Flow Filter Types

[RFC5575] and [I-D.ietf-idr-flow-spec-v6] specify the above types for BGP. The encoding for "Destination Prefix" is described in [RFC5575] as -

Encoding: <type (1 octet), prefix length (1 octet), prefix>

In PCEP, the type of flow filter is identified by the type field in the TLV header, TBD1 in case of Destination Prefix. The length field in the TLV header (as per [RFC5440]) is the length of the value portion in octets without padding. The value portion for "Destination IPv4 Prefix" is made up of 1 octet of prefix length followed by the prefix, padded to 4-byte alignment for the TLV.

Similarly for all encoding defined in [RFC5575] and [I-D.ietf-idr-flow-spec-v6], the value portion of the PCEP TLV uses the BGP encoding but without the type octet and pad it to 4-byte alignment.

[I-D.dhodylee-pce-pcep-ls] allow identificatiion of a VPN information in PCEP via a Route Distinguisher (RD) [RFC4364] and encoded in ROUTE-DISTINGUISHER TLV. This TLV MAY be included in the FLOW object to identify the flow filter infomration, say a IPv4 destination prefix, is a VPNv4 destination prefix belonging to the VPN identified by the RD.

5.3. ACTION Object

The ACTION object MUST be present within FlowSpec messages when creating or updating the FlowSpec. The ACTION object carries a set of FlowSpec actions.

ACTION Object-Class is TBD20 (to be assigned by IANA).

ACTION Object-Type is 1.

The format of the ACTION object body is:

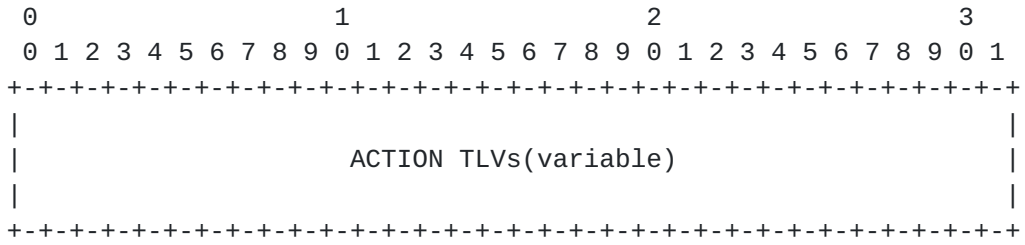

Figure 4: ACTION Object Body Format

The ACTION object body has a variable length and may contain one or more additional TLVs.

The following FlowSpec action types are supported:

Type	Description	Defined in
18(*)	IPV4-LSP-IDENTIFIERS	I-D.ietf-pce-stateful-pce
19(*)	IPV6-LSP-IDENTIFIERS	I-D.ietf-pce-stateful-pce
17(*)	Symbolic-Path-Name	I-D.ietf-pce-stateful-pce

(*) The type is defined in [I-D.ietf-pce-stateful-pce]

Figure 5: Flow Action Types

6. IANA Considerations

IANA maintains the "Path Computation Element Protocol (PCEP) Numbers" registry at <http://www.iana.org/assignments/pcep>. This document requests IANA actions to allocate code points for the protocol elements defined in this document.

6.1. PCEP Messages

IANA maintains a subregistry for PCEP messages called "PCEP Messages". Each PCEP message has a message type value. This document defines a new PCEP message type value.

Value	Meaning	Reference
TBD17	FlowSpec	[This I-D]

6.2. PCEP Objects

Each PCEP object has an Object-Class and an Object-Type. IANA maintains a subregistry called "PCEP Objects". This document defines the following new PCEP Object-classes and Object-values:

Object-Class Value	Name	Object-Type	Reference
TBD19	FLOW	1	[This I-D]
TBD20	ACTION	1	[This I-D]

6.3. PCEP TLV Type Indicators

IANA maintains a subregistry called "PCEP TLV Type Indicators". This document defines the following new PCEP TLVs.

Value	Meaning	Reference
TBD18	PCE-FLOWSPEC-CAPABILITY TLV	[This I-D]
TBD1	Destination IPv4 Prefix	[This I-D]
TBD2	Source IPv4 Prefix	[This I-D]
TBD3	IP Protocol	[This I-D]
TBD4	Port	[This I-D]
TBD5	Destination port	[This I-D]
TBD6	Source port	[This I-D]
TBD7	ICMP type	[This I-D]
TBD8	ICMP code	[This I-D]
TBD9	TCP flags	[This I-D]
TBD10	Packet length	[This I-D]
TBD11	DSCP	[This I-D]
TBD12	Fragment	[This I-D]
TBD13	Flow Label	[This I-D]
TBD14	Destination IPv6 Prefix	[This I-D]
TBD15	Source IPv6 Prefix	[This I-D]
TBD16	Next Header	[This I-D]

7. Security Considerations

TBD.

8. Acknowledgements

TBD.

9. References

9.1. 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>.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <http://www.rfc-editor.org/info/rfc5575>.

[I-D.ietf-idr-flow-spec-v6]
McPherson, D., Raszuk, R., Pithawala, B., akarch@cisco.com, a., and S. Hares, "Dissemination of Flow Specification Rules for IPv6", draft-ietf-idr-flow-spec-v6-07 (work in progress), March 2016.

[I-D.dhodylee-pce-pcep-ls]
Dhody, D., Lee, Y., and D. Ceccarelli, "PCEP Extension for Distribution of Link-State and TE Information.", draft-dhodylee-pce-pcep-ls-04 (work in progress), July 2016.

9.2. Informative References

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[RFC4657] Ash, J., Ed. and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <http://www.rfc-editor.org/info/rfc4657>.

[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-14 (work in progress), March 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-06 (work in progress), July 2016.

[I-D.ietf-idr-bgp-flowspec-oid]
Uttaro, J., Filsfils, C., Smith, D., Alcaide, J., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", draft-ietf-idr-bgp-flowspec-oid-03 (work in progress), March 2016.

[I-D.ietf-ospf-flowspec-extensions]
liangqiandeng, l., You, J., Wu, N., Fan, P., Patel, K., and A. Lindem, "OSPF Extensions for Flow Specification", draft-ietf-ospf-flowspec-extensions-01 (work in progress), April 2016.

[I-D.kuppani-pce-pcep-flowspec-sync]
Kuppani, S. and A. Sinha, "PCEP Flowspec Synchronization Procedures.", draft-kuppani-pce-pcep-flowspec-sync-00 (work in progress), May 2016.

[I-D.zhao-teas-pce-control-function]
Farrel, A., Zhao, Q., Li, Z., and C. Zhou, "An Architecture for Use of PCE and PCEP in a Network with Central Control", draft-zhao-teas-pce-control-function-01 (work in progress), May 2016.

Appendix A. Contributor Addresses

Shankara
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India
Email: shankara@huawei.com
Qiandeng Liang
Huawei Technologies
101 Software Avenue, Yuhuatai District
Nanjing 210012
China
Email: liangqiandeng@huawei.com

Appendix B. Example Usage

Once PCE initiate tunnels, it needs to further decide what data needs to flow on the newly created tunnel, a flow specification can be created at the ingress to redirect the flow to the LSP as shown below.

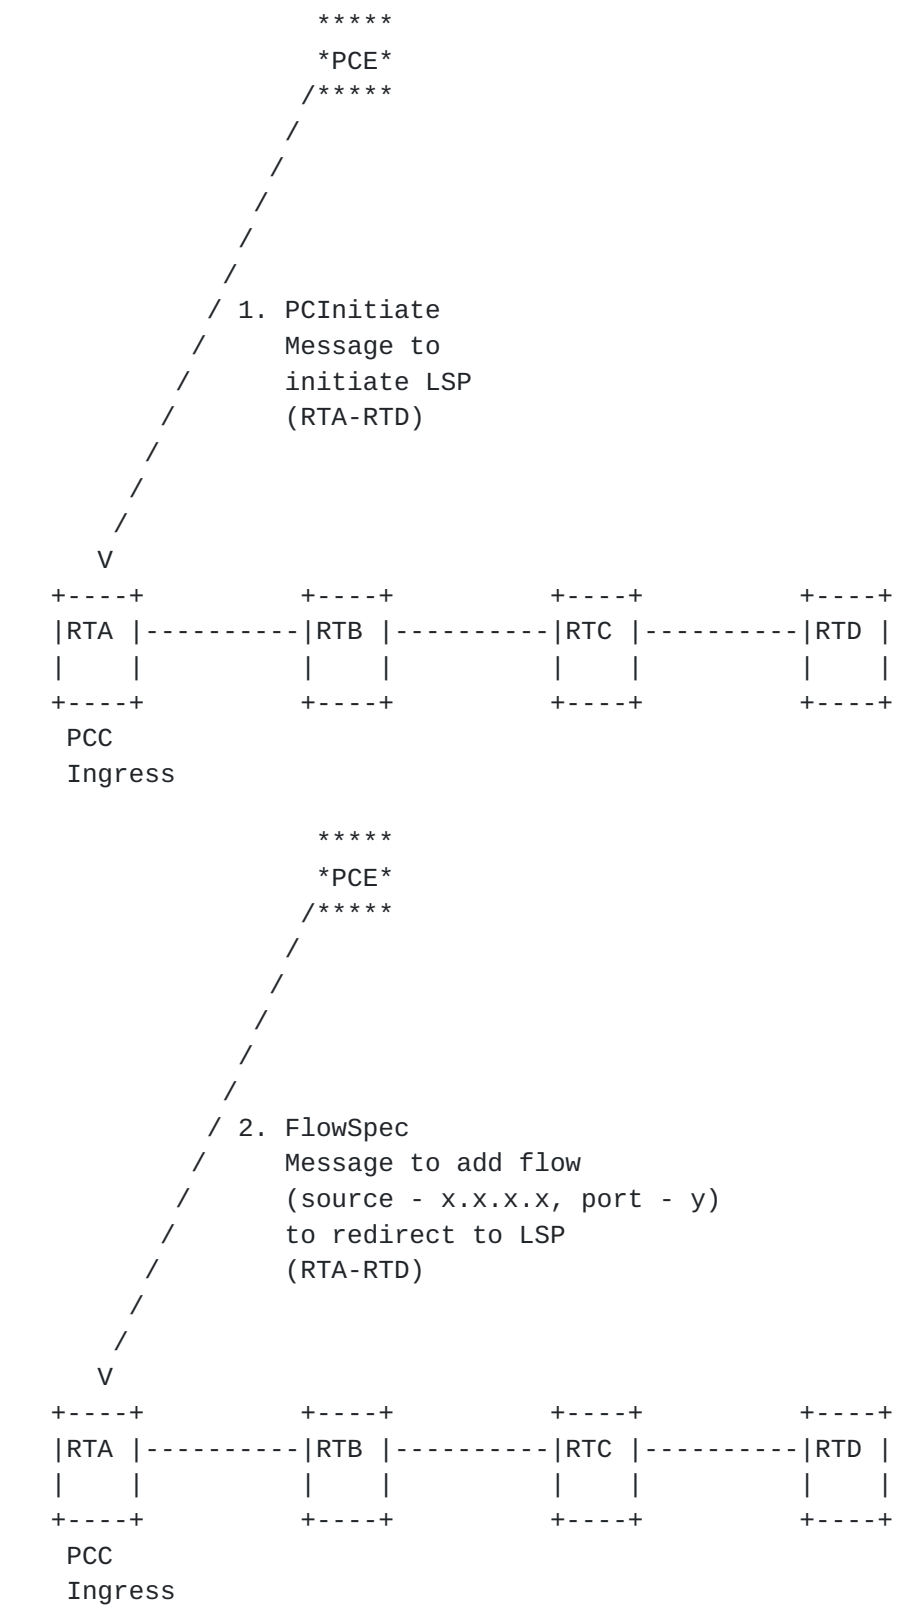
*****
PCE
/*****
1. PCInitiate
Message to
initiate LSP
(RTA-RTD)
V
+----+ +----+ +----+ +----+
|RTA |----------|RTB |----------|RTC |----------|RTD |
| | | | | | | |
+----+ +----+ +----+ +----+
PCC
Ingress
*****
PCE
/*****
2. FlowSpec
Message to add flow
(source - x.x.x.x, port - y)
to redirect to LSP
(RTA-RTD)
V
+----+ +----+ +----+ +----+
|RTA |----------|RTB |----------|RTC |----------|RTD |
| | | | | | | |
+----+ +----+ +----+ +----+
PCC
Ingress

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

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

Email: dhruv.ietf@gmail.com

Cyril Margaria
Juniper
200 Somerset Corporate Boulevard, , Suite 4001
Bridgewater, NJ 08807
USA

Email: cmargaria@juniper.net

Colby Barth
Juniper
200 Somerset Corporate Boulevard, , Suite 4001
Bridgewater, NJ 08807
USA

Email: cbarth@juniper.net

Xia Chen
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: jescia.chenxia@huawei.com

Shunwan Zhuang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: zhuangshunwan@huawei.com