

# IETF 112 Path Computation Element (PCE) WG

Wednesday, Nov 10, 2021 (14:30-15:30 UTC)

Friday, Nov 12, 2021 (16:00-17:00 UTC)

Chairs

Julien Meuric ( [julien.meuric@orange.com](mailto:julien.meuric@orange.com) )

Dhruv Dhody ( [dd@dhruvdhody.com](mailto:dd@dhruvdhody.com) )

Secretary

Hariharan Ananthakrishnan ( [hari@netflix.com](mailto:hari@netflix.com) )

**This session is being recorded**



IETF 112  
hosted by

Internet Engineering Task Force  
© 2021 IETF Trust  
Production by Meetecho



# Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (<https://www.ietf.org/contact/ombudsteam/>) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- [BCP 9](#) (Internet Standards Process)
- [BCP 25](#) (Working Group processes)
- [BCP 25](#) (Anti-Harassment Procedures)
- [BCP 54](#) (Code of Conduct)
- [BCP 78](#) (Copyright)
- [BCP 79](#) (Patents, Participation)
- <https://www.ietf.org/privacy-policy/>(Privacy Policy)



# Reminder for the IETF Guidelines for Conduct

- IETF participants extend respect and courtesy to their colleagues at all times.
- IETF participants have impersonal discussions.
- IETF participants devise solutions for the global Internet that meet the needs of diverse technical and operational environments.
- Individuals are prepared to contribute to the ongoing work of the group
- See BCP 54!

This session is being recorded

# IETF 112 Online Meeting Tips

- Your audio is muted and video is off when you join
  - Unmute only when you are recognized in the queue
  - Mute and remove yourself from the queue after you are done
  - Video is welcome when you are speaking in the queue
  - Use of a headset is strongly recommended
- Session bluesheet is automatically generated based on IETF Datatracker logins
- Chatrooms in Meetecho are connected to the Jabber chatrooms on IETF Datatracker agenda
- More information and assistance:
  - Participant guide  
<https://www.ietf.org/how/meetings/technology/meetecho-guide-participant/>

# Administrivia

- Minute taker(s), jabber scribe(s)
- Meetecho Etiquette
  - Join the queue if you would like to speak/present
    - Do not send audio directly
  - Please state your name before speaking
  - Be mindful of the agenda time
    - Longer discussion on mailing list (or jabber)
- Collaborative minutes
  - <https://notes.ietf.org/notes-ietf-112-pce>

# Usual Reminders

- Please use the mailing list actively!
- Please be more vocal during WG business (WG LC, adoption, etc)!
- Use the WG wiki to track progress -  
<https://trac.ietf.org/trac/pce/wiki/WikiStart>
- Request for early code point allocation when you are planning to interop!

# Agenda Bashing

**Wednesday, November 10, 2021 14:30-15:30 UTC**

## **Introduction**

- 1.1. Administrivia, Agenda Bashing (Chairs, 5 min) [5/60]
- 1.2. WG Status (Chairs, 10 min) [15/60]
- 1.3. State of WG I-Ds and next steps (Chairs, 10 min) [25/60]

## **Segment Routing**

- 2.1. Multipath ERO (Mike Koldychev, 10 mins) [35/60]  
[draft-ietf-pce-multipath-03](#)
- 2.2. SR Policy (Mike Koldychev, 10 min) [45/60]  
[draft-ietf-pce-segment-routing-policy-cp-06](#)
- 2.3. Path MTU (Luc-Fabrice Ndifor, 5 mins) [50/60]  
[draft-li-pce-pcep-pmtu-05](#)
- 2.4. IFIT (Giuseppe Fioccola, 5 mins) [55/60]  
[draft-chen-pce-pcep-ifit-04](#)
- 2.5. Ingress Protection (Huaimo Chen, 5 mins) [60/60]  
[draft-chen-pce-sr-ingress-protection-06](#)

**Friday, November 12, 2021 16:00-17:00 UTC**

## **Update to PCEP**

- 3.1. Relax Object Ordering (Dhruv Dhody, 10 mins) [10/60]  
[draft-dhody-pce-pcep-object-order-00](#)
- 3.2. Topology Filter (Quan Xiong, 10 mins) [20/60]  
[draft-xpbs-pce-topology-filter-01](#)
- 3.3. VTN in PCEP (Minxue Wang, 10 mins) [30/60]  
[draft-dong-pce-pcep-vtn-00](#)
- 3.4. PCEP-LS (Gyan Mishra, 10 mins) [40/60]  
[draft-dhodylee-pce-pcep-ls-22](#)

## **Multicast**

- 4.1. BIER (Huanan Li, 10 mins) [50/60]  
[draft-li-pce-based-bier-02](#)

## **Others**

- 5.1. Color in PCEP (Balaji Rajagopalan, 5 mins) [55/60]  
[draft-rajagopalan-pce-pcep-color-00](#)
- 5.2. VLAN-based Native IP (Yue Wang, 5 mins) [60/60]  
[draft-wang-pce-vlan-based-traffic-forwarding-01](#)

# WG Status



# Beyond the WG

- no new RFCs since IETF 111
- RFC Editor Queue
  - draft-ietf-pce-pcep-flowspec - RFC-EDITOR state
    - -13 posted to strip L2VPN flowspec from this I-D to remove dependency on IDR Flowspec V2 work
    - A separate I-D draft-li-pce-pcep-l2-flowspec-00 is posted
      - A quick WG adoption?
- With the AD
  - draft-ietf-pce-binding-label-sid
    - Changes made based on AD review

# In the WG's Hands

- Errata
  - RFC 8231 - Technical (Rejected)
    - Order of LSP and CLASSTYPE object in PCReq message
    - New I-D proposed (on the agenda)
- Early IANA codepoint allocation
  - draft-ietf-pce-local-protection-enforcement
    - Expires 2022-01-28
  - draft-ietf-pce-segment-routing-policy-cp
    - Expires 2022-03-30
  - draft-ietf-pce-binding-label-sid
    - Expires 2022-03-29

# Status of WG I-Ds & Next Steps



# WG documents “nearing” WG LC

## **draft-ietf-pce-pcep-stateful-pce-gmpls**

- No recent update
- -15 posted on 2021-06-24
- Reorganization done by authors after the merge of 2 I-Ds
- **Ready for WG-LC!**

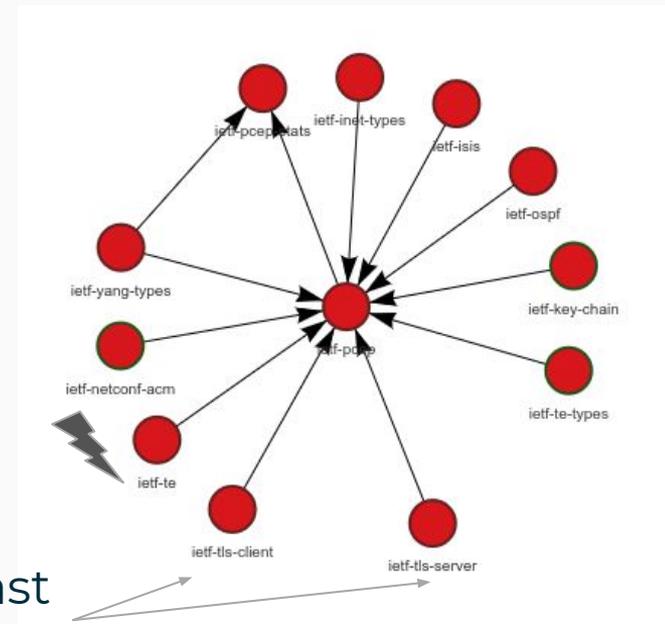
## **draft-ietf-pce-enhanced-errors**

- -10 posted on 2021-08-26
  - No changes!
- Is there still interest in this work?
- **Options:**
  - Progress this work as experimental
    - Would need reviewers to commit
  - Mark it as waiting for implementation
  - Any extension planning to use it

# WG documents “nearing” WG LC

## draft-ietf-pce-pcep-yang

- -17 posted on 2021-10-23
- Comments from Tom Petch are handled
- Very early YANG Doctor review was done
  - ready for another one?
- Or directly WGLC next!



WG Last  
Call

# WG I-Ds

## **draft-ietf-pce-pcep-extension-native-ip**

- -16 posted on 2021-08-15
- Comments from Sue handled
- Agenda time in IDR to get feedback early
  - Cross posting will be done for WG LC

## **draft-ietf-pce-flexible-grid**

- -06 posted on 2021-09-08
  - No major technical change since a long time!
- Is this ready?

# WG I-Ds

## **draft-ietf-pce-segment-routing-ipv6**

- No recent update
- -09 posted on 2021-05-28
- **Ready for WG-LC next?**

## **draft-ietf-pce-vn-association**

- -05 posted on 2021-10-15
  - No major technical change since a long time!
- Authors consider it ready!

# WG I-Ds

## **draft-ietf-pce-sr-path-segment**

- -04 posted on 2021-08-12
- No major technical change
- Are there any open issues?
- Nearing WG LC?

## **draft-ietf-pce-sr-bidir-path**

- -08 posted on 2021-09-09
- Sync with RFC 9059
- Nearing WG LC?

# WG I-Ds

## **draft-ietf-pce-segment-routing-policy-cp**

- -06 posted on 2021-10-22
- On the agenda

## **draft-ietf-pce-local-protection-enforcement**

- -03 posted on 2021-08-05
- No significant change
- Is this ready?

# WG I-Ds

## **draft-ietf-pce-pcep-extension-pce-controller-sr**

- -03 posted on 2021-09-30
- Aligned to the published RFC 9050

## **draft-ietf-pce-stateful-interdomain**

- No recent update
- -02 posted on 2021-07-12

## **draft-ietf-pce-lsp-extended-flags**

- -01 posted on 2021-10-18
- Comments received during adoption call are handled

# WG I-Ds

## **draft-ietf-pce-multipath**

- -03 posted on 2021-10-25
- On agenda

## **draft-ietf-pce-state-sync**

- -01 posted on 2021-10-20
- Comments received during adoption are handled

# Recently adopted WG I-Ds

## **draft-ietf-pce-stateful-pce-optional**

- Adopted on 2021-10-12
- -02 posted on 2021-10-23
- Comments received during adoption are handled

# WG Adoption Poll Queue

- draft-hsd-pce-sr-p2mp-policy (on going)
- ...
- Reshuffling and update to be made at <https://trac.ietf.org/trac/pce/wiki/WikiStart#WGAdoptionCallQueue> after the IETF 112

Thanks!



# Backup!

# Using the Mailing List

- Please use the mailing list actively to discuss all working group business
- Open issues with drafts should be discussed on the list, and conclusions reported to the list
- New drafts should be introduced to the working group first on the mailing list, to gauge interest
- Working group consensus is determined from the mailing list
- Priority in meetings is given to drafts that have been discussed on the list

# Please be Vocal

- During WG Adoption and WG LC calls, response number is low.
- Please be vocal on the list to help us gauge the consensus better.
- The WG mailing lists are looked at by the IESG, IAB, and others (internal and external to IETF) to determine interest/participation level in our standards process.
- Please review ideas from your peers, these are community outputs of the working group as a whole.
- Also help flushing our queues faster
  - we had to extend the calls when response was lacking!

# Using the Wiki

- A way to give you visibility as the document progress through the WG
  - adoption queue
  - WG LC queue
  - balancing work between chairs
  - shepherding responsibilities and opportunities
  - pending actions
  - IPR polls
- Use this wiki
  - make sure this is up to date!
- <https://trac.ietf.org/trac/pce/wiki/WikiStart>

# Early Codepoint Allocation

- If you have an implementation of a WG I-D
  - that requires inter-operation with other implementations
    - Please request for early IANA codepoint allocation
  - Make sure to include an Implementation Status section in your I-D
  - Make sure the IANA section is correct and complete
    - And meets the condition set out in RFC 7120
- Maintained at
  - <https://trac.ietf.org/trac/pce/wiki/WikiStart#CandidateforearlyIANAAllocations>



IETF 112 – Online  
PCE Working Group

# PCEP Extensions for Signaling Multipath Information draft-ietf-pce-multipath

*M. Koldychev – Cisco Systems ([mkoldych@cisco.com](mailto:mkoldych@cisco.com)) – Presenter*  
*M. Sivabalan – Ciena Corporation ([ssivabal@ciena.com](mailto:ssivabal@ciena.com))*  
*T. Saad – Juniper Networks ([tSaad@juniper.net](mailto:tSaad@juniper.net))*  
*V. Beeram – Juniper Networks ([vbeeram@juniper.net](mailto:vbeeram@juniper.net))*  
*H. Bidgoli – Nokia ([Hooman.Bidgoli@Nokia.com](mailto:Hooman.Bidgoli@Nokia.com))*  
*S. Peng – Huawei Technologies ([pengshuping@Huawei.com](mailto:pengshuping@Huawei.com))*  
*B. Yadav – Ciena ([byadav@ciena.com](mailto:byadav@ciena.com))*

# Review

## Abstract

Path computation algorithms are not limited to return a single optimal path. Multiple paths may exist that satisfy the given objectives and constraints. This document defines a mechanism to encode multiple paths for a single set of objectives and constraints. This is a generic PCEP mechanism, not specific to any path setup type or dataplane. The mechanism is applicable to both stateless and stateful PCEP.

## Mechanisms in this draft are as generic as possible:

- Independent of data-plane/setup-type (RSVP-TE/SR-MPLS/SRv6)
- Independent of stateful/stateless PCEP, i.e., it works with PCReq/PCReply

# Forward and Reverse paths

A path computation algorithm that is given a computational problem can output:

- **0 or 1** forward paths (without this draft),
- **N forward paths** (with this draft),
- **N forward paths, plus M reverse paths** (with latest update).

Returning only **0 or 1** forward paths was sufficient for RSVP-TE tunnels.

Returning **N forward** paths is required for SR Policy with multiple Segment Lists (SL) under one Candidate Path.

Returning **N forward paths, plus M reverse paths** is useful for Circuit-Style SR Policy [draft-schmutzer-pce-cs-sr-policy]. It allows the head-end to learn about reverse SL(s) for each forward SL.

# OPPDIR-PATH TLV

New TLV in the PATH-ATTRIB object:

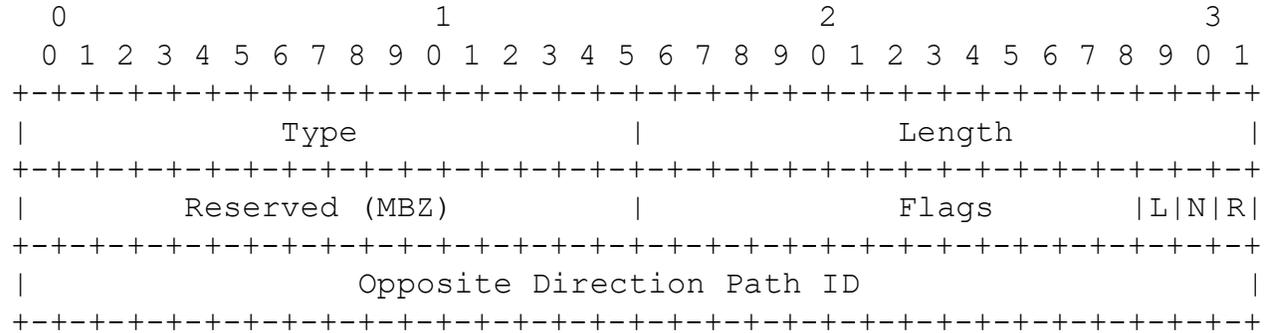


Figure 5: MULTIPATH-OPPDIR-PATH TLV format

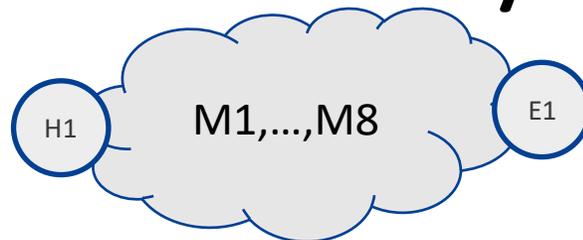
Opposite Direction Path ID points to another PATH-ATTRIB object within the same Tunnel. It signifies that the pointed-to path is the opposite of the current path.

R-flag is set to 1 when the current path (described by the current PATH-ATTRIB object) is going in the reverse direction w.r.t the Tunnel, i.e., from the end-point to the head-end.

Each path can have multiple opposite paths.

Oppositeness property is NOT mutual (eg., Path 2 is an opposite of Path1, but Path 1 is NOT an opposite of Path 2).

# Example: Circuit Style SR Policies



## Router H1:

```
SR policy POL1 <headend = H1, color, endpoint = E1>
Candidate-path CP1
  Preference 200
  Bidirectional Association = A1
  SID-List = <H1,M1,M2,E1>
  SID-List = <H1,M3,M4,E1>
Candidate-path CP2
  Preference 100
  Bidirectional Association = A2
  SID-List = <H1,M5,M6,E1>
  SID-List = <H1,M7,M8,E1>
```

```
<state-report> =
<LSP PLSP ID=100>
<BIDIRECTIONAL ASSOCIATION = A1>
<PATH-ATTRIB PathID=1
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=3>>
<ERO <H1,M1,M2,E1>>
<PATH-ATTRIB PathID=2
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=4>>
<ERO <H1,M3,M4,E1>>
<PATH-ATTRIB PathID=3
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=1>>
<ERO <E1,M2,M1,H1>>
<PATH-ATTRIB PathID=4
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=2>>
<ERO <E1,M4,M3,H1>>
```

```
<state-report> =
<LSP PLSP ID=200>
<BIDIRECTIONAL ASSOCIATION = A2>
<PATH-ATTRIB PathID=1
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=3>>
<ERO <H1,M5,N6,E1>>
<PATH-ATTRIB PathID=2
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=0>>
<ERO <H1,M7,M8,E1>>
<PATH-ATTRIB PathID=3
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=1>>
<ERO <E1,M6,M5,H1>>
```

## Router E1:

```
SR policy POL2 <headend = E1, color, endpoint = H1>
Candidate-path CP1
  Preference 200
  Bidirectional Association = A1
  SID-List = <E1,M2,M1,H1>
  SID-List = <E1,M4,M3,H1>
Candidate-path CP2
  Preference 100
  Bidirectional Association = A2
  SID-List = <E1,M6,M5,H1>
```

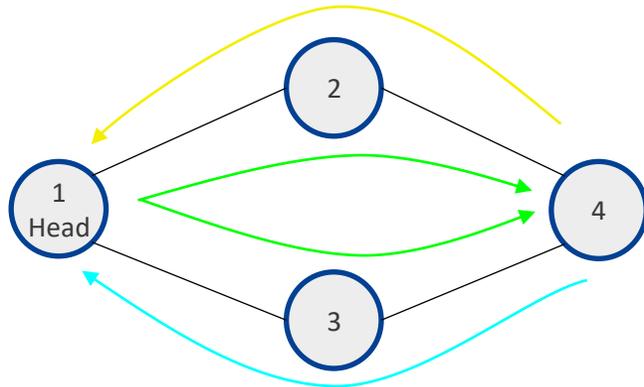
```
<state-report> =
<LSP PLSP ID=100>
<BIDIRECTIONAL ASSOCIATION = A1>
<PATH-ATTRIB PathID=1
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=3>>
<ERO <E1,M2,M1,H1>>
<PATH-ATTRIB PathID=2
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=4>>
<ERO <E1,M4,M3,H1>>
<PATH-ATTRIB PathID=3
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=1>>
<ERO <H1,M1,M2,E1>>
<PATH-ATTRIB PathID=4
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=2>>
<ERO <H1,M3,M4,E1>>
```

```
<state-report> =
<LSP PLSP ID=200>
<BIDIRECTIONAL ASSOCIATION = A2>
<PATH-ATTRIB PathID=1
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=3>>
<ERO <E1,M6,M5,H1>>
<PATH-ATTRIB PathID=2
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=0>>
<ERO <H1,M7,M8,E1>>
<PATH-ATTRIB PathID=3
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=1>>
<ERO <H1,M5,N6,E1>>
```

Data  
model

PCEP  
signaling

# Example: forward path with 2 reverse paths



Node(X) -> node segment of router X  
Adj(X,Y) -> adjacency segment from router X to router Y  
Where X, Y ∈ {1,2,3,4}

In Segment Routing, node segments can send traffic along multiple links. Thus, multiple segment lists may be required to express all the reverse paths.

Multiple instances of the OPPDIR-PATH-TLV encode multiple reverse paths. For example, the following encodes that Segment List <Node(4)> has two reverse Segment Lists: <Adj(42),Adj(21)> and <Adj(43),Adj(31)>.

```
<PATH-ATTRIB PathID=1>  
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=2>  
  <OPPDIR-PATH-TLV R-flag=0 OppositePathID=3>>  
<ERO <Node (4) >>  
<PATH-ATTRIB PathID=2>  
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=0>>  
<ERO <Adj (4,2),Adj (2,1)>>  
<PATH-ATTRIB PathID=3>  
  <OPPDIR-PATH-TLV R-flag=1 OppositePathID=0>>  
<ERO <Adj (4,3),Adj (3,1)>>
```

Note that just because Path 2 is a reverse of Path 1, does NOT mean that Path 1 is a reverse of Path 2.

Value of R-flag among all instances of the OPPDIR-PATH-TLV MUST be the same. We could also put this R-flag into the PATH-ATTRIB object, instead of the OPPDIR-PATH-TLV?

# Next steps

- Get feedback from WG
- Request IANA code point allocation



IETF 112 – Online  
PCE Working Group

# PCEP extension to support Segment Routing Policy Candidate Paths

draft-ietf-pce-segment-routing-policy-cp

*M. Koldychev – Cisco Systems ([mkoldych@cisco.com](mailto:mkoldych@cisco.com)) – Presenter*

*M. Sivabalan – Ciena Corporation ([ssivabal@ciena.com](mailto:ssivabal@ciena.com))*

*C. Barth – Juniper Networks ([cbarth@juniper.net](mailto:cbarth@juniper.net))*

*S. Peng – Huawei Technologies ([pengshuping@Huawei.com](mailto:pengshuping@Huawei.com))*

*H. Bidgoli – Nokia ([Hooman.Bidgoli@Nokia.com](mailto:Hooman.Bidgoli@Nokia.com))*

# Generic mechanisms

Several “tunnel” mechanisms were standardized as part of SR Policy [draft-ietf-spring-segment-routing-policy]:

- Computation Priority
- Explicit Null Label Policy (ENLP)
- Drop-upon-invalid
- Specified-BSID-only

Even though these were originally standardized for SR Policy, they are applicable to other tunnel types. In this draft, we keep them generic, so that they are automatically applicable to RSVP-TE tunnels.



# Specified-BSID-only

When specified-BSID-only is enabled for a particular binding SID, it means that the given binding SID is required to be allocated and programmed for the LSP to be operationally up.

We request a bit in the TE-PATH-BINDING TLV Flags field to signal this behavior.

Note that each LSP can have multiple BSIDs (MPLS, SRv6, etc.) and the value of this flag is per BSID.



# Implementation Status

Cisco IOS-XR and Juniper has proof-of-concept with IANA code-points.

Successful interop testing between Cisco IOS-XR and Juniper has recently been done at EANTC.

## Next steps

Get feedback from WG.

Request IANA code points for the additional TLVs.

# Support for Path MTU (PMTU) in the Path Computation Element (PCE) communication Protocol (PCEP)

draft-li-pce-pcep-pmtu-05

Luc-Fabrice Ndifor [Luc-Fabrice.Ndifor@mtn.com](mailto:Luc-Fabrice.Ndifor@mtn.com)

Shuping Peng Huawei

Cheng Li Huawei

Liuyan Han China Mobile

Luc-Fabrice Ndifor MTN Cameroon

# Motivation

- In traditional MPLS, the Path MTU can be signaled via signaling protocols like RSVP-TE[3209] and LDP[RFC3988].
- However, there is no additional signaling to establish Segment Routing (SR) paths, so the SR tunnel cannot currently support the negotiation mechanism of the Path MTU.
  - SR information is reported by BGP-LS, and the PCE can calculate the SR Paths based on this info.
- When SIDs (Label or IPv6 address) are pushed in a packet, the packet will be dropped (in IPv6) or fragmented in forwarding since the packet size may exceed the Path MTU.
- From Operator:
  - When using leased line over multi-domains, MTU should be learned to avoid dropping packets.
- This draft is to specify the **extensions to PCEP** to carry **Path MTU** in PCEP messages.

# History of this draft

- Presented the version 02 @ IETF108 and received comments and suggestions
- We have updated the draft and addressed all the comments in the version 03
  - ✓ A Terminology session has been added to clarify the often confusing terms including MTU, Link MTU, Path MTU.
  - ✓ A Path MTU Adjustment session has been added to include the case of protection such as TI-LFA.
  - ✓ Some editorial changes.
- Presented the version 03 @ IETF109 and asked for adoption, comments were further addressed in this v04.
  - Text were added to clarify the relationship to MSD
  - Text were added to clarify the support on Multi-path [draft-ietf-pce-multipath]

## Next Step

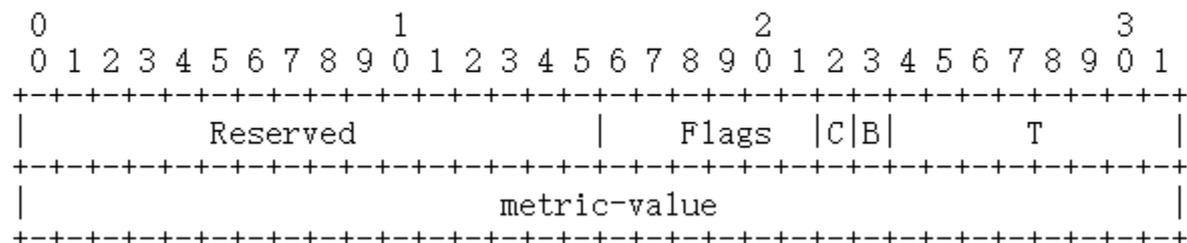
- PMTU is a very important feature to have for Network Operators.
- We would like to again ask for the WG adoption of this draft.

**Thank you for your attention!**

# METRIC Object for Path MTU

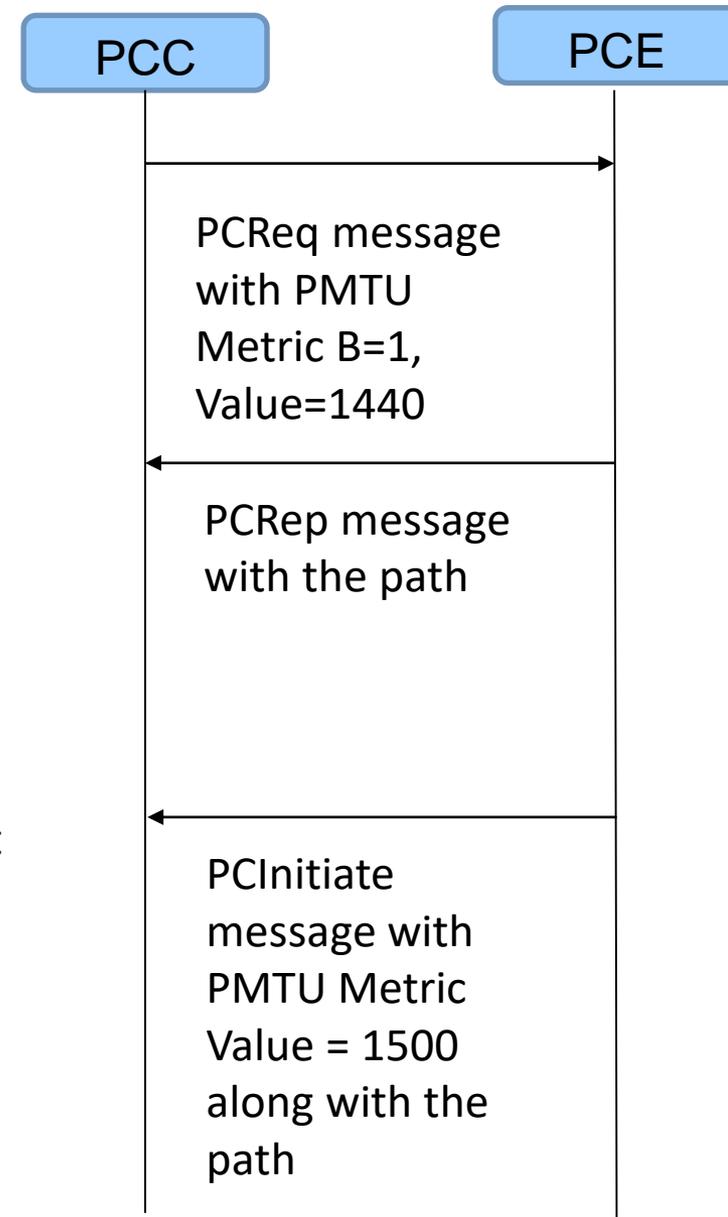
- This document defines a new type for the existing METRIC object for Path MTU.
  - T = TBD by IANA
  - B (Bound - 1 bit): Bound
  - metric-value = PMTU
- The Path MTU metric type of the METRIC object in PCEP represents the minimum of the Link MTU of all the links along the path.

The format of the METRIC object body is as follows:



# PMTU for Segment Routing

- PCE can be used for computing one or more SR-TE paths taking into account various constraints and objective functions.
  - Path MTU could be another metric for PCE to consider
- Once a path is chosen, the PCE can inform an SR-TE path on a PCC using PCEP extensions specified in [RFC8664].
  - PCE could also inform the Path MTU to the PCC
- [I-D.ietf-pce-segment-routing-ipv6] adds the support for IPv6 data plane in SR.
- **The new metric type for path MTU is applicable for the SR-TE path and does not require any additional extensions.**



# **Path Computation Element Communication Protocol (PCEP) Extensions to Enable IFIT**

**draft-chen-pce-pcep-ifit-04**

Online, Nov 2021, IETF 112

Hang Yuan (UnionPay)  
Tianran Zhou (Huawei)  
Weidong Li (Huawei)  
Giuseppe Fioccola (Huawei)  
Yali Wang (Huawei)

# Background and Motivation

- ❑ In-situ Flow Information Telemetry (**IFIT**) refers to dataplane on-path telemetry techniques, including **IOAM** (draft-ietf-ippm-ioam-data) and **Alternate Marking** (RFC8321, RFC8889)
- ❑ The **PCEP extension** defined in this document allows to signal the IFIT capabilities. In this way IFIT methods are automatically activated and running.

The IFIT attributes can be generalized and included as **TLVs** carried inside the **LSPA (LSP Attributes) object** in order to be applied for all path types, as long as they support the relevant data plane telemetry method

# Latest Changes

- Specified the usage scenario of IFIT

IFIT is a solution focusing on specific network domains according to RFC8799.

- For a number of reasons, such as policies, options supported, style of network management and security requirements, it is suggested to limit applications including the emerging IFIT techniques to a controlled domain.

- Improved Security Considerations section

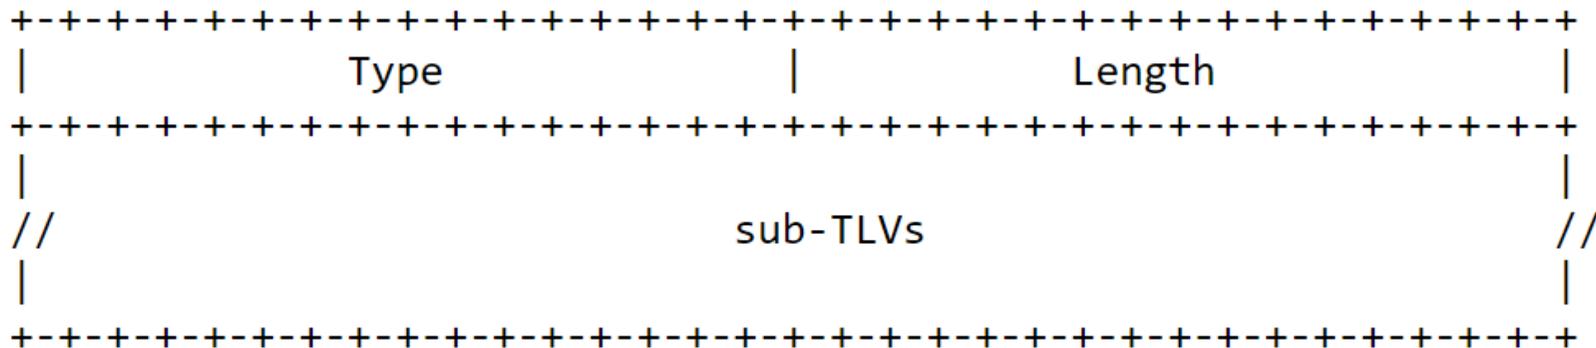
IFIT data **MUST** be propagated in a limited domain to avoid malicious attacks. Solutions to ensure this requirement are respectively discussed in [draft-ietf-ippm-ioam-data](#) and [draft-ietf-6man-ipv6-alt-mark](#).

- A limited administrative domain provides the network administrator with the means to select, monitor and control the access to the network, making it a trusted domain also for the PCEP extensions defined in this document.



# IFIT Attributes TLV

The **IFIT-ATTRIBUTES TLV** provides the configurable knobs of the IFIT feature, and it can be included as an optional TLV in the **LSPA object**



IFIT attribute TLVs, carried inside the LSPA object and applicable to all path types

- IFIT TLVs are optional and can be taken into account by the PCE during path computation and by the PCC during path setup.
- In general, the LSPA object can be carried within a PCInitiate message, a PCUpd message, or a PCRpt message in the stateful PCE model.

# IOAM Sub-TLVs

- IOAM Pre-allocated Trace Option Sub-TLV

Type=1	Length=8	
Namespace ID	Rsvd1	
IOAM Trace Type	Flags	Rsvd2

- IOAM Incremental Trace Option Sub-TLV

Type=2	Length=8	
Namespace ID	Rsvd1	
IOAM Trace Type	Flags	Rsvd2

- IOAM Directly Export Option Sub-TLV

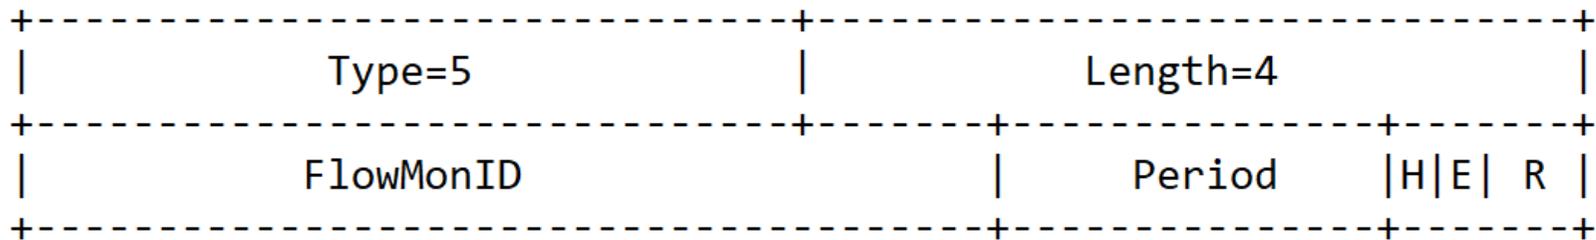
Type=3	Length=12
Namespace ID	Flags
IOAM Trace Type	Rsvd
Flow ID	

- IOAM Edge-to-Edge Option Sub-TLV

Type=4	Length=4
Namespace ID	IOAM E2E Type

# Enhanced Alternate Marking Sub-TLV

- Enhanced Alternate Marking Sub-TLV



H: A flag indicating that the measurement is Hop-By-Hop.

E: A flag indicating that the measurement is end to end.

# Discussion & Next Steps

- Since IFIT methods are becoming mature for SR-MPLS and SRv6, IFIT attributes TLV also complements [draft-ietf-pce-segment-routing-policy-cp](#) to enable SR policy with native IFIT.
- Ask for WG adoption

Welcome questions, comments

Thank you

# PCE for Path Ingress Protections

draft-chen-pce-sr-ingress-protection-06

Huaimo Chen, Mike McBride(Futurewei)  
Mehmet Toy, Gyan S. Mishra (Verizon Inc.)  
Aijun Wang (China Telecom)  
Zhenqiang Li, Yisong Liu (China Mobile)  
Boris Khasanov (Yandex LLC)  
Lei Liu (Fujitsu)  
Xufeng Liu (Volta Networks)

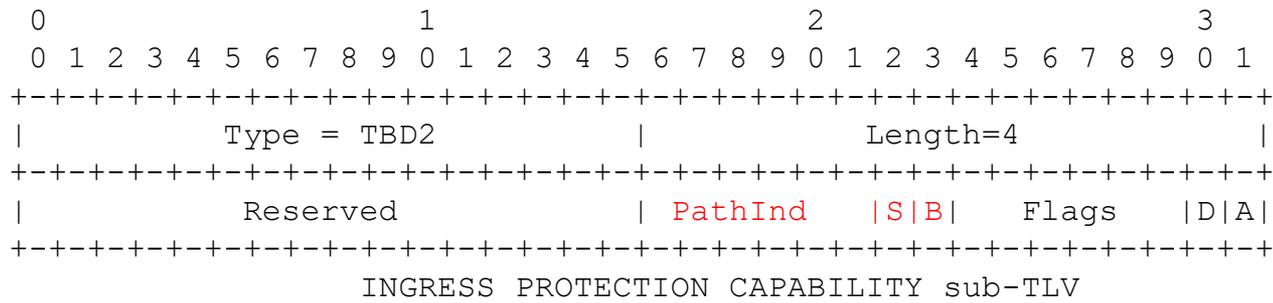
IETF 112

# Overview

- Merged 2 Ingress Protection drafts
  - pce-sr-ingress-protection
  - pce-bier-te-ingress-protect
- PCE for ingress protection of 2 types of paths:
  - SR paths
  - BIER-TE paths
- Foundation for protecting ingress of different types of paths

# Capability for Ingress Protection /w Backup Ingress

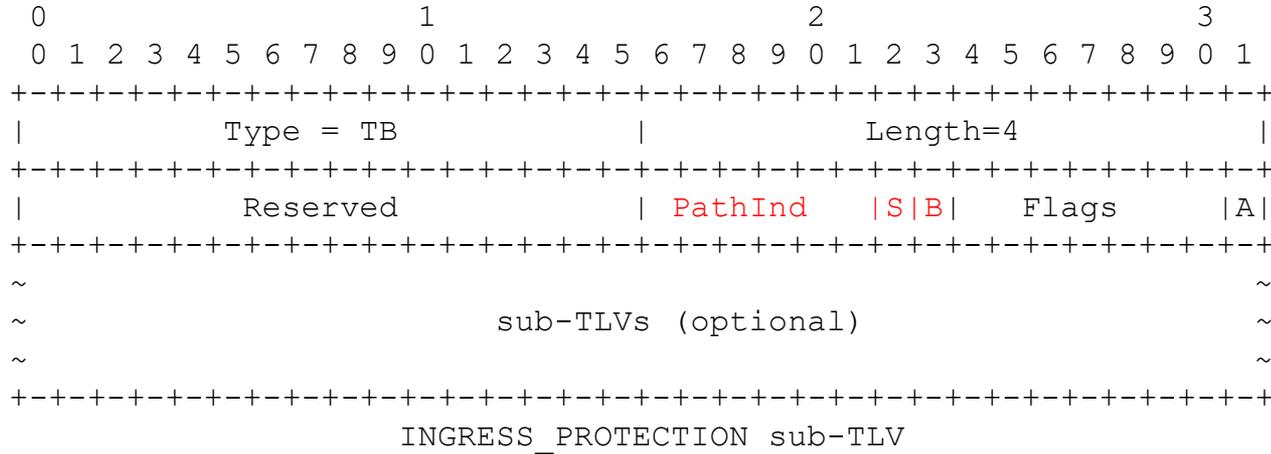
When PCE and PCC running on a backup ingress establish a PCEP session, they exchange their capabilities of protecting ingress for different types of paths.



- PathInd: 1 octet. Indicators for the types of paths whose ingress protections are supported. Two indicators are defined.
- o S : S = 1 indicating that ingress protection of SR path is supported.
  - o B : B = 1 indicating that ingress protection of BIER-TE path is supported.

# Extensions for Backup Ingress

When PCE sends PCC PCInitiate message for initiating a backup path to protect the primary ingress node of a primary path, the message contains:



PathInd: 1 octet. Indicator for the type of path whose ingress is protected. Two indicators are defined.

- o S : S = 1 indicating ingress protection of a SR path.
- o B : B = 1 indicating ingress protection of a BIER-TE path.

# Next Step

- Comments
- Request for adoption