

PCEP Extension for Flow Specification

draft-li-pce-pcep-flowspec-02

Dhruv Dhody : <dhruv.ietf@gmail.com>

Adrian Farrel : <afarrel@juniper.net>

Zhenbin Li : <lizhenbin@huawei.com>

IETF-99, Prague, July 2017

Document History/Status

- Was making good progress, but then stalled
- This revision
 - Injects new momentum
 - Updates the mechanism
 - No new messages
 - Retains the purpose and functionality

Objectives

- When a PCE is in control of an LSP, how does it tell the head end what traffic to put on the LSP?
 - Applies to PCE-initiated LSPs and delegated LSPs
- When a PCE does load balancing TE, how does it know what traffic is on which LSP?
 - Applies to PCC-initiated LSPs
 - Also applies at the moment of delegation
- Don't re-invent wheels
 - Use existing BGP Flowspec encodings
 - **BUT** this is nothing to do with IDR

Mechanism (1/2)

- Capability advertisement in IGP and OPEN message
- New PCEP object
 - Flow Spec Object
 - 0, 1, or more instances on PCReq, PCRep, PCErr, PCInitiate, PCRpt, and PCUpd
 - Each instance describes a traffic flow using TLVs
- New PCEP TLV
 - Flow Filter TLV
 - Exactly one per Flow Spec Object
 - Comprised of sub-TLVs (next slide)
 - This TLV is only present to enable sub-TLV codepoint management

Mechanism (2/2)

- Sub-TLVs
 - Flow Specification TLVs
 - Combined to describe the flow
 - TLV types as follows

Range	
0	Reserved - must not be allocated.
1 .. 255	Per BGP registry defined by [RFC5575]. Not to be allocated in this registry.
256 .. 65535	New PCEP Flow Specs allocated according to the registry defined in this document.

- TLV value fields encoded as BGP specifications or as defined in this or new PCEP specs
- Additional TLVs defined for...
 - VPN Route Distinguisher
 - Multicast flows (for P2MP LSPs)

For Example

- PCE initiates a new P2P LSP and wants it to be used for all traffic destined to 198.51.100.0/24 and all traffic destined to 203.0.113.0/24
- It sends a PCInitiate message for the LSP and includes a Flow Spec Object containing a Flow Filter TLV
- It includes two Flow Specification TLVs
 - Type = 0x0001 (IPv4 destination prefix)
Length = 0x0004
Value = prefix length in bits (1 octet) + prefix
0x18C63364
 - Type = 0x0001 (IPv4 destination prefix)
Length = 0x0004
Value = prefix length in bits (1 octet) + prefix
0x18CB0071

Work Still to be Done

- Push on with implementations
- Include examples
- Simplify / streamline main use cases
 - Consider removing some BGP Flowspecs that are “too complicated”
 - Special consideration of forwarding capabilities
- Consider ordering issues
 - This remains a challenge for Flowspec in BGP
 - It is even an issue for static routes
 - Need clear and unambiguous rules
- Understand from WG if this is:
 - In scope
 - Something they want to work on