

Technical Issues with draft-ietf-mpls-bfd-directed

IETF 99, Praha

C. Pignataro, Cisco

Background

- Loa Andersson's email, July 12th:
 - 3 Unsuccessful WG LCs
 - Current Issues: IPR, Technical, Process
 - We will take half an hour of the allocated time for mpls wg to try to reach a point where a decision is possible.
 - Carlos, expect 10 min presentation on the technical issues (including questions)

Background

- 2nd WG LC on draft-ietf-mpls-bfd-directed
 - <https://mailarchive.ietf.org/arch/msg/mpls/iXsENQuPWmmgsueNiUMlyRCcu4U>
- “authors decided to remove sections Segment Routing”
 - [https://mailarchive.ietf.org](https://mailarchive.ietf.org/arch/msg/mpls/iXsENQuPWmmgsueNiUMlyRCcu4U)

Background

- (3rd) Re: [mpls] WGLC for draft-ietf-mpls-bfd-directed
 - <https://mailarchive.ietf.org/arch/msg/mpls/iv-rCHGjiRVAndWOHnOiJ9HO2YY>
 - https://mailarchive.ietf.org/arch/msg/mpls/VZE_a1wL299Rjxav3SPT_QlgKCM
- Authors then allowed multiple sub-TLVs
 - <https://mailarchive.ietf.org/arch/msg/mpls/rluaKtXhkHCMUMEMeYyYGI7zgZjcLQ>

Background

- [RTG-DIR] Rtgdir early review of draft-ietf-mpls-bfd-directed-07
 - <https://mailarchive.ietf.org/arch/msg/rtg-dir/CFfhXtm62K8j1BQPxNstqOPyYsA>
- Technical issues: lack specification for interoperable implementation

Technical Issues: Motivation


- “For the case where a LSP is explicitly routed it is likely that the shortest return path to the ingress BFD peer would not follow the same path as the LSP in the forward direction. *The fact that BFD control packets are not guaranteed to follow the same links and nodes in both forward and reverse directions is a significant factor in producing false positive defect notifications, i.e. false alarms, if used by the ingress BFD peer to deduce the state of the forward direction.*”

Technical Issues: Motivation

- “a failure detection by ingress node on the reverse path cannot be interpreted as bi-directional failure unambiguously and thus trigger, for example, protection switchover of the forward direction without possibility of being a false positive.
- To address this scenario the egress BFD peer would be instructed to use a specific path for BFD control packets.”

Technical Issues

- Motivation



New structure: FEC sub-TLVs

- “Reverse Path field contains a sub-TLV. Any Target FEC sub-TLV (already defined, or to be defined in the future) for TLV Types 1, 16, and 21 of MPLS LSP Ping Parameters registry MAY be used in this field. None, one or more sub-TLVs MAY be included in the BFD Reverse Path TLV. If none sub-TLVs found in the BFD Reverse Path TLV, the egress BFD peer MUST revert to using the default, i.e., over IP network, reverse path.”
- May it use the Nil FEC? PW 129 FEC? Multicast FECs?
- What if *some sub-TLVs” are found?
- Any issues if the return LSP terminates mid stream?



RFC 5884 does not specify a Default

- <https://tools.ietf.org/html/rfc5884#section-7>
- “The BFD Control packet sent by the egress LSR to the ingress LSR **MAY** be routed based on the destination IP address as per the procedures in [[BFD-MHOP](#)]. If this is the case, the destination IP address **MUST** be set to the source IP address of the LSP Ping Echo request message, received by the egress LSR from the ingress LSR.
- **Or the BFD Control packet sent by the egress LSR to the ingress LSR MAY be encapsulated in an MPLS label stack.** In this case, the presence of the fault detection message is indicated as described above. This may be the case if the FEC for which the fault detection is being performed corresponds to a bidirectional LSP or an MPLS PW. This may also be the case when there is a return LSP from the egress LSR to the ingress LSR. In this case, the destination IP address **MUST** be randomly chosen from the 127/8 range for IPv4 and from the 0:0:0:0:0:FFFF:7F00/104 range for IPv6.”

Persistency

- **Return Path Specified Label Switched Path (LSP) Ping**
- <https://tools.ietf.org/html/rfc7110>
- However,
 1. what if the return path is not initially available?
 2. Should the egress save this FEC stack?
 3. Or if it goes down while the BFD session is UP?

Session Establishment

- “If the egress LSR cannot find the path specified in the Reverse Path TLV it MUST send Echo Reply with the received Reverse Path TLV and set the Return Code to *“Failed to establish the BFD session. The specified reverse path was not found”* Section 3.3. *The egress BFD peer MAY establish the BFD session over IP network as defined in [RFC5884].”*

Additional State?

- How is the return path tracked? New state variable?

<https://tools.ietf.org/html/rfc5880#section-6.8.1>

Switching over to a different return path



- How exactly? Text?
- Does it depend on the ingress identifying a change at egress?
- Does this update RFC 5884 procedures?

Port Usage

- The BFD Control packet sent by the egress LSR MUST have a well-known destination port 4784, if it is routed [[BFD-MHOP](#)], or it MUST have a well-known destination port 3784 [[BFD-IP](#)] if it is encapsulated in a MPLS label stack. The source port MUST be assigned by the egress LSR as per the procedures in [[BFD-IP](#)].
- <https://tools.ietf.org/html/rfc5884#section-7>

Thank you!