

# BFD in Segment Routing Networks Using MPLS Dataplane

draft-mirsky-spring-bfd

Greg Mirsky  
Jeff Tantsura  
Mach Chen  
Ilya Varlashkin

IETF-100 November 2017, Singapore

# BFD over MPLS dataplane

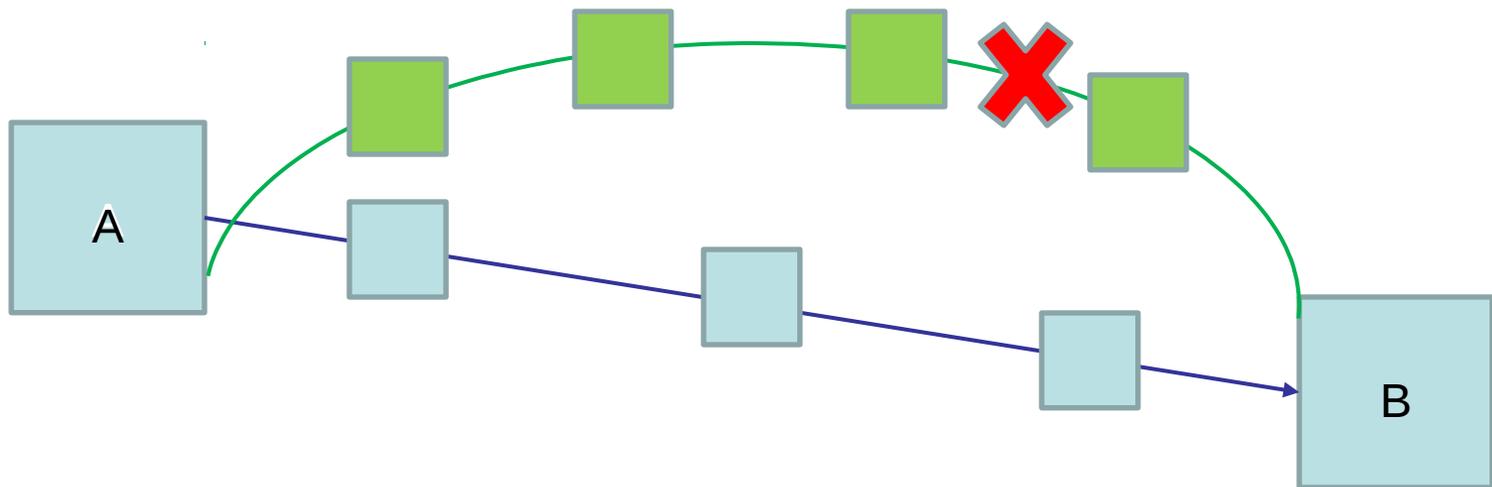
- RFC 5884 has defined use of BFD Asynchronous mode over MPLS LSP
- BFD over segment route with MPLS dataplane SHOULD use LSP Ping to bootstrap BFD session
- In addition to requirements stated in draft-ietf-mpls-spring-lsp-ping:
  - Initiator MUST include FEC(s) corresponding to the destination segment.
  - Initiator, i.e. ingress LSR, MAY include FECs corresponding to some or all of segments imposed in the label stack by the ingress LSR to communicate the segments traversed.

add:

- When LSP Ping is used to bootstrap a BFD session the FEC corresponding to the destination segment to be associated with the BFD session MUST be as the very last sub-TLV in the Target FEC TLV.
- BFD control packet encapsulation:
  - with IP/UDP header MUST:
    - destination IP address 128/8 for IPv4 address or 0:0:0:0:FFFF:7F00/104 for IPv6 address;
    - use UDP destination port 3784
  - ACH encapsulation use GAL and G-ACh type 0x0007

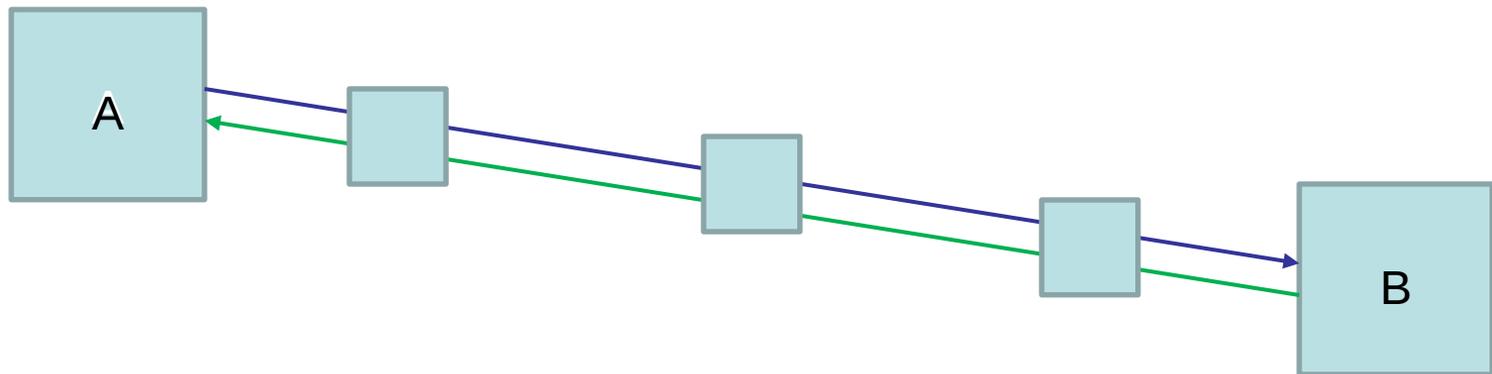
# BFD Reverse Path

- Ingress LER A periodically transmits BFD control messages over MPLS LSP
- Egress LER B periodically transmits BFD control messages, per RFC 5884, over path selected based on local policy:
  - IP network using UDP destination port 4784
  - reverse path segment route with IP/UDP encapsulation (UDP destination port 3784) or ACH encapsulation
- Failure in the reverse path of the BFD session may be interpreted as LSP failure



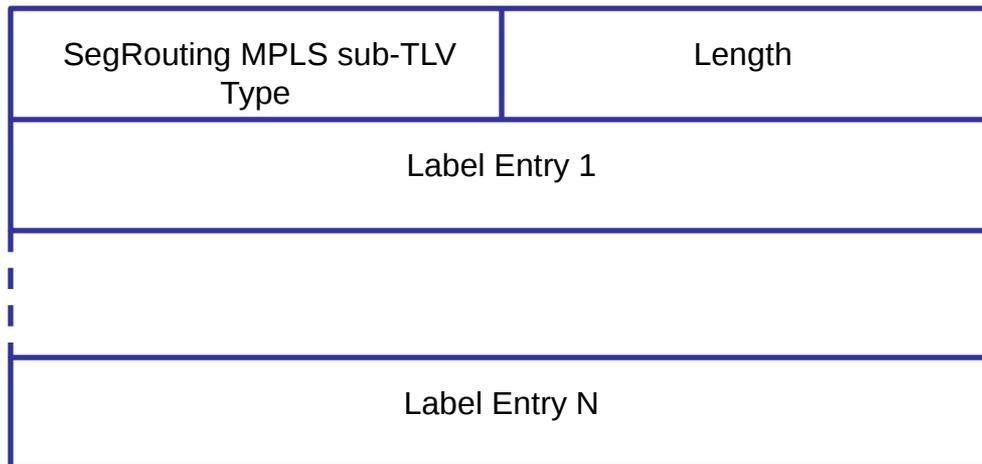
# Control BFD Reverse Path

- New optional BFD Reverse Path TLV
- Used with BFD Discriminator TLV
- Instructs egress BFD to transmit BFD control packets over the specified MPLS LSP
- Re-use sub-TLVs defined in draft-ietf-mpls-spring-lsp-ping
- BFD Reverse Path TLV may contain none, one or more sub-TLVs
- If none sub-TLV has been found in the BFD Reverse Path TLV, then the egress BFD MUST transmit BFD control packets over IP network



# New Segment Routing Static MPLS Tunnel sub-TLV

- Ordered list of Label Stack Elements with the top of the stack label as Label Entry 1 and the bottom of the stack label – Label Entry N
- BFD Reverse TLV MAY include zero or one SR Static MPLS Tunnel sub-TLV
- If no sub-TLVs present in the BFD Reverse Path TLV – the egress MUST switch the reverse BFD session to be transmitted over IP network
- If more then one SR Static MPLS Tunnel sub-TLVs present in the BFD Reverse Path TLV, the remote peer MUST send MPLS LSP Echo Reply with Return Code value set to “Too Many TLVs Detected” (new code)



# Next steps

- Your comments, suggestions, questions always welcome and greatly appreciated
- Which WG to anchor – MPLS or SPRING?