

LSP Self-Ping

draft-bonica-mpls-self-ping-03

R. Torvi, R. Bonica

Juniper Networks

M.Conn, D. Pacella, L. Tomotaki, M. Wygant

Verizon

The High Order Bits (Stay Awake!)

- Ingress LSR's use LSP Self-ping to verify that an LSP is ready to carry traffic
- The ingress LSR probes an LSP under test by sending an MPLS Echo message, addressed to itself, through the LSP
- The egress LSR forwards the MPLS Echo message back to the ingress LSR, exactly as it would forward any other packet
- The ingress LSR determines that the LSP under test is ready to carry traffic when it receives the MPLS Echo message
- Mind-numbing details follow

Why Do We Need This?

- Many RSVP implementations allow LSRs to send a RESV messages upstream before installing forwarding state
- This optimization is desirable, because it allows LSRs to install forwarding state in parallel
- However, this optimization creates a race condition
 - When the ingress LSR receives a RESV message, some downstream LSRs may have not yet completed the process of forwarding state installation
 - If the ingress sends traffic over the LSP, the traffic will be black-holed until forwarding state has been installed on all downstream LSRs

Rejected Approaches

- MPLS Echo [RFC 4379]
 - Rejected because it consumes control plane resources at egress LSR
 - Control plane resources may be scarce during periods of network perturbation
- When the ingress LSR receives a RESV message, wait a configurable period of time before using the LSP
 - Rejected because it is impossible to determine the optimal timer setting
 - If the timer value is set too low, it does not prevent black-holing
 - If the timer value is set too high, it slows down the process of LSP signaling and setup
 - Optimal timer value is a moving target

Next Steps

- Please adopt as a WG item

LSP Ping Procedure

Backup Slides

LSP Ping Session

- Short-lived (order seconds)
- Session state is local to ingress LSR
 - Session-id: A 32-bit number that identifies the session
 - verification-status: A boolean variable indicating whether LSP readiness has been verified. The initial value of this variable is FALSE
 - retries: The number of times that the ingress LSR probes the LSP before giving up. The initial value of this variable is determined by configuration
 - retry-timer: The number of milliseconds that the LSR waits after probing the LSP. The initial value of this variable is determined by configuration.

Probing The LSP Under Test

- The ingress LSR executes the following procedure until verification status equals TRUE or retries is less than 1
 - Format a MPLS Echo [RFC4379] message
 - Send the MPLS Echo message through the LSP under test
 - Set a timer to expire in retry-timer milliseconds
 - Wait until either a MPLS Echo message associated with the session returns or the timer expires. If an MPLS Echo message associated with the session returns, set verification-status to TRUE. Otherwise, decrement retries. Optionally, increase the value of retry-timer according to an appropriate back of algorithm.

The MPLS Echo Message

- The MPLS Echo message is encapsulate in a UDP header
- If the protocol messages used to establish the LSP were delivered over IPv4, the UDP datagram is encapsulated in an IPv4 header
- If the protocol messages used to establish the LSP were delivered over IPv6, the UDP datagram is encapsulated in an IPv6 header
- In either case, message contents are as defined on the next slide

The MPLS Echo Message (continued)

- IP Source Address is configurable. By default, it is the address of the egress LSR
- IP Destination Address is the address of the ingress LSR
- IP Time to Live (TTL) / Hop Count is 255
- IP DSCP is configurable. By default, it is equal to CS6 (0x48)
- UDP Source and destination ports are any port selected from the dynamic range (49152-65535)
- MPLS Echo Global Flags are clear (i.e., set to 0)
- MPLS Echo Type is equal to "MPLS Echo Reply" (2)
- MPLS Echo Reply Mode is "Reply via an IPv4/IPv6 UDP packet" (2)
- MPLS Echo Senders Handle is equal to the Session-ID
- MPLS Echo Sequence Number is equal to retries
- MPLS Echo Time Stamp Sent is equal to the current time