

LSP Ping Relay Reply

<draft-ietf-mpls-lsp-ping-relay-reply-07>

L. Jin

J. Luo

T. Nadeau

G. Swallow

Changes in draft-07

- LSP Ping Relayed Echo Reply Stack Object: Added
Source Address of Replying Router
Destination Address Pointer
- LSP Ping Relayed Echo Reply Stack Entry
Eliminated Address Length since it is implied by the type
- Relay Node Address Stack description
Moved procedural information into Section 4, Procedures
- Generalized MTU exceeded error
Generally useful, wider applicability

Changes in draft-07 (2)

- Procedures
 - Updated text (both technical changes and clarifications)
 - Changed many SHOULDs to MUSTs
- Section 4.7, Impact to Traceroute
 - Updated to use Source Address of Replying Router
- Section 5, LSP Ping Relayed Echo Reply Example
 - Updated example to use changed procedures

Relayed Echo Reply Stack Object

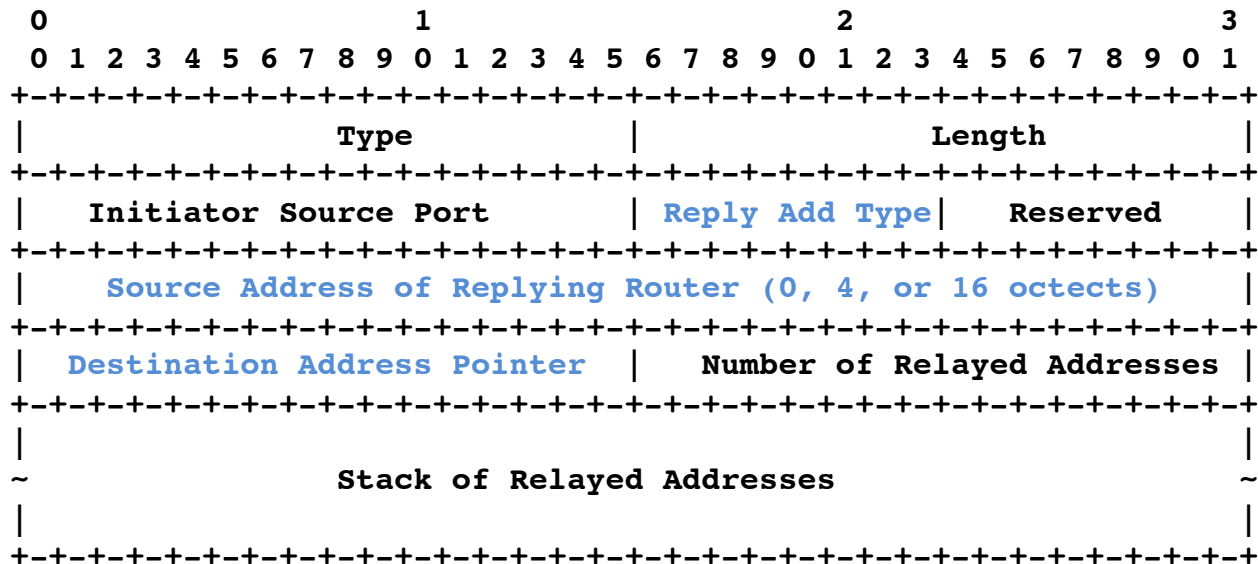
Source Address of Replying Router

Had been carried in Address stack

Confused semantics, often did not even belong there


Used for loop-detection and display

Destination Address Pointer




Source Address of Replying Router

- Semantic Issue
 - Procedures mandated adding the IF address of replying router at the bottom of the stack
 - Adding other address optional
- In example below ASBR 2 MUST add 10.0.0.2 at BOS, but that address is not reachable from ASBR 4!
- Address needed for relaying must be reachable through the interface from which the LSP ping would have been forwarded if TTL had not expired



Loop Detection

- Procedures said a the next relay node address MUST be located before the source IP address of the received Relayed Echo Reply which MUST be also in the stack
- In example below address 10.0.0.1 is used in both ASes
- Added Destination Address Stack Pointer
Points to address to which the message is being sent



Mandatory address in stack

- Mandated adding an address reachable through the interface from which the LSP ping would have been forwarded at the BOS
 - If a router is
 - hiding its address
 - located at an addressing domain boundary
 - e.g. AS boundary or between areas using BGP-LU (RFC3107) vs IGP
- then a Null address **MUST** be added with the K bit set

Stack processing

- Procedures in two cases ignored the significance of the K bit
- When eliminating addresses, this can only be done from the last entry with the K bit set to the end
- When choosing the next relay router **MUST** not use one above the last entry with the K bit set

Next Steps

- Workgroup review (requested March 7)
- MPLS-RT review (in progress)
- Address comments
- Last call