

Proxy LSP Ping

<draft-lim-proxy-lsp-ping-00.txt>

IETF-85

swallow@cisco.com

History

- First draft October 2006
draft-swallow-mpls-remote-lsp-ping-00.txt
- Accepted as a WG document by April 2007
draft-ietf-mpls-remote-lsp-ping-00.txt
- Intension was to go to last call after IETF-73
November 2008 (rev -03)
- Three things occurred
 - Decided to complete p2mp ping first (RFC in Nov 2011)
 - Largest usecase is for mLDP which is currently being rolled out
 - Wanted implementation experience

Proxy Ping Motivation

- Scalability

 - Control

 - number of replies

 - network wide processing of MPLS Echo Requests

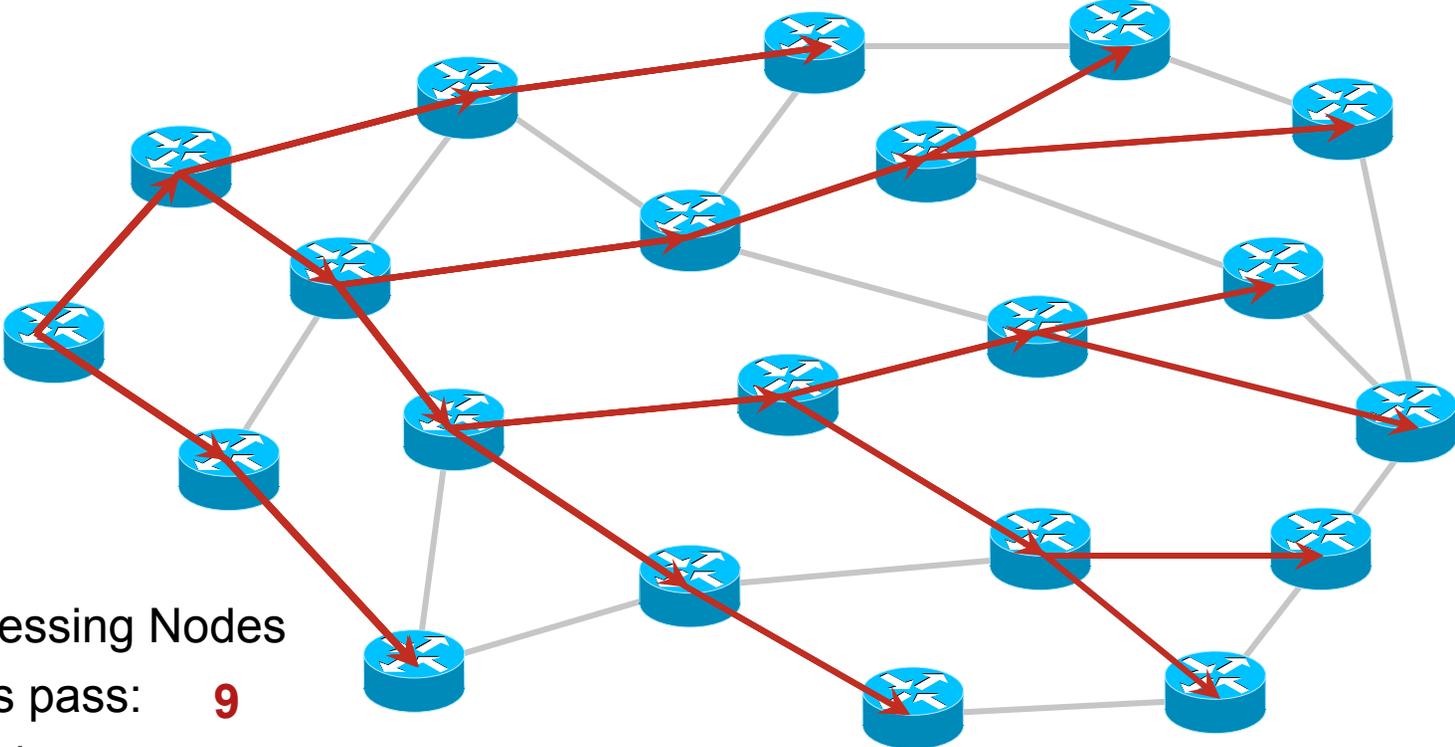
 - Reply scoping can limit the number of replies but still requires processing at the receiving node

- mLDP support

 - Trace from leaf to root

 - Previous Hop Information

Mcast LSP Trace



Processing Nodes
This pass: **9**
Total: **30**

Proxy LSP Ping

- Proxy Echo Request

- Out of band message to a to a “proxy” node

- presumably a node on the path of the LSP of interest

- Requests node to insert an inband MPLS Echo Request

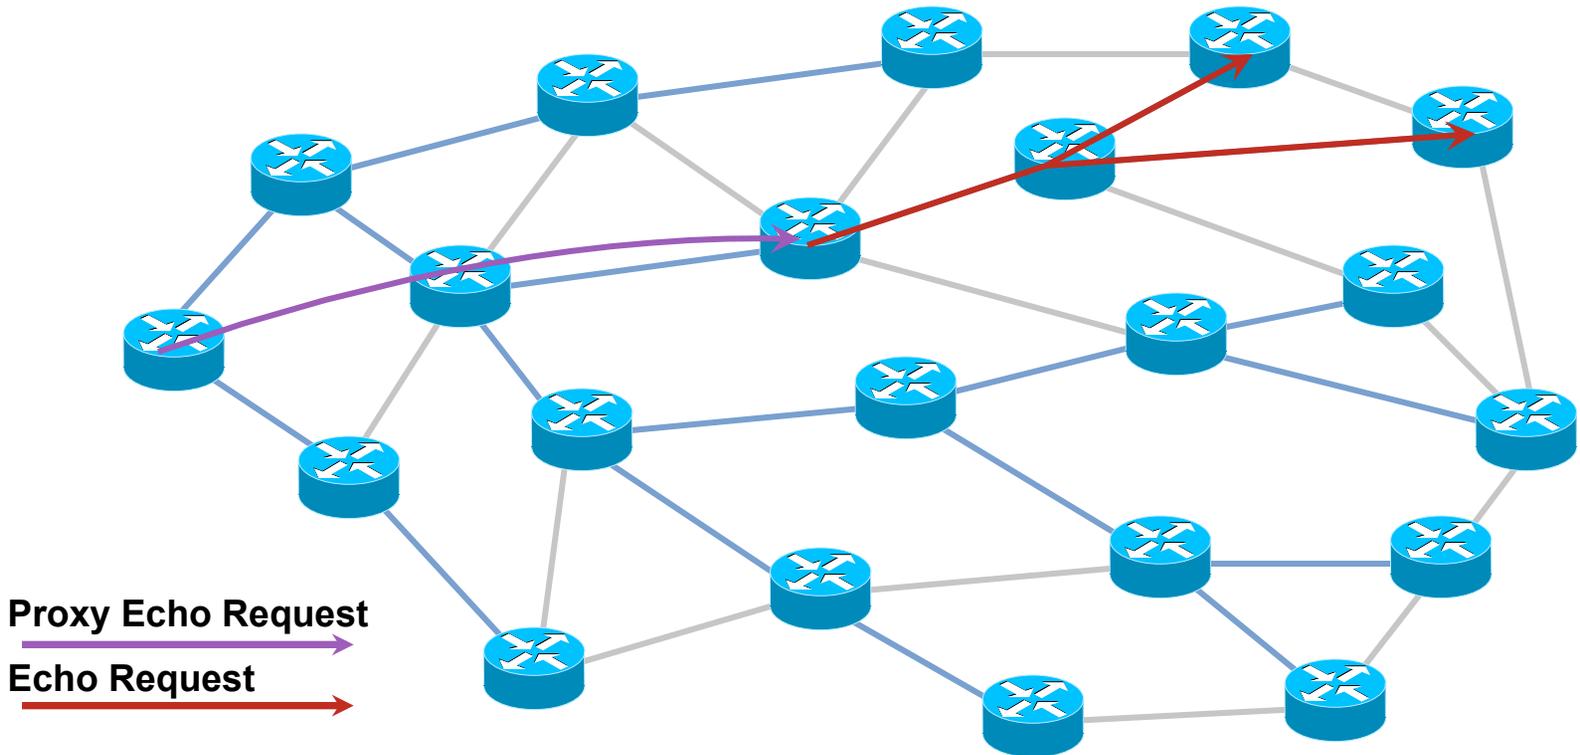
- Or asks for a Proxy Echo Reply for details on outgoing ports or upstream neighbor

- Proxy Echo Reply

- Verifies that FEC was valid

- Supplies PHOP information, downstream mapping, downstream neighbors

Proxy LSP Ping



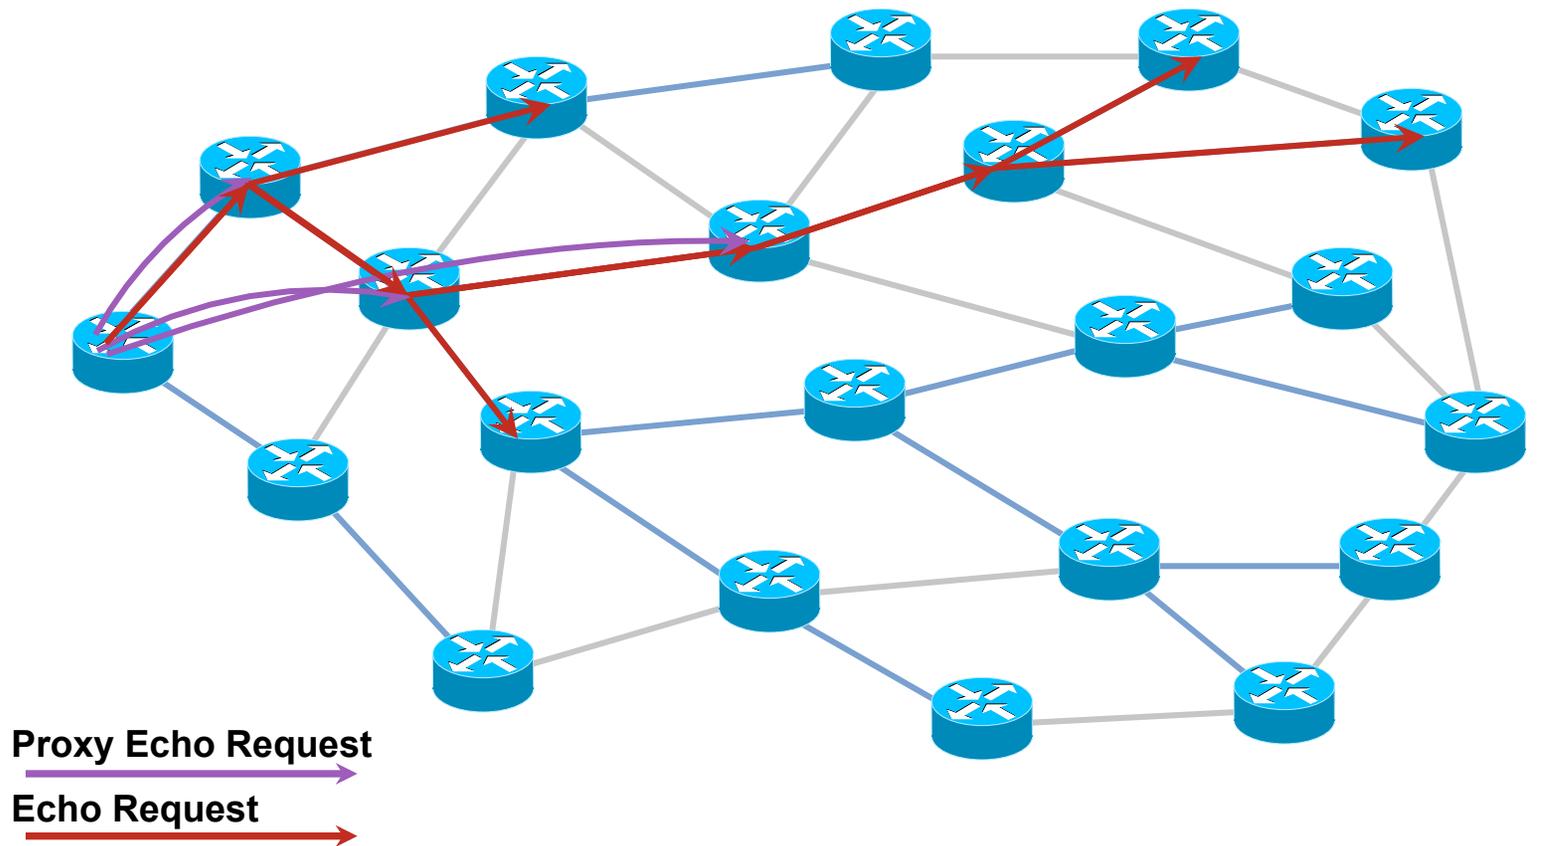
Next-Hop Scoping

- Further scalability by limiting which downstream branches are explored
- Requester
 - Adds the next-hop IP address to the Proxy Echo Parameters Object
 - Sends Proxy Echo Request to proxy node
- Proxy node
 - Only sends an MPLS Echo Request on the listed interfaces

P2MP-TE Fault Localization

- In Traffic Engineering, the topology of the tree is known
- The branches which proceed towards the node reporting trouble can be determined
- Traffic on extra branches can be avoided
 - By initiating trace further into the tree
 - By scoping to the branch of interest

P2MP-TE with Next-Hop Scoping



Note: Echo Replies not shown

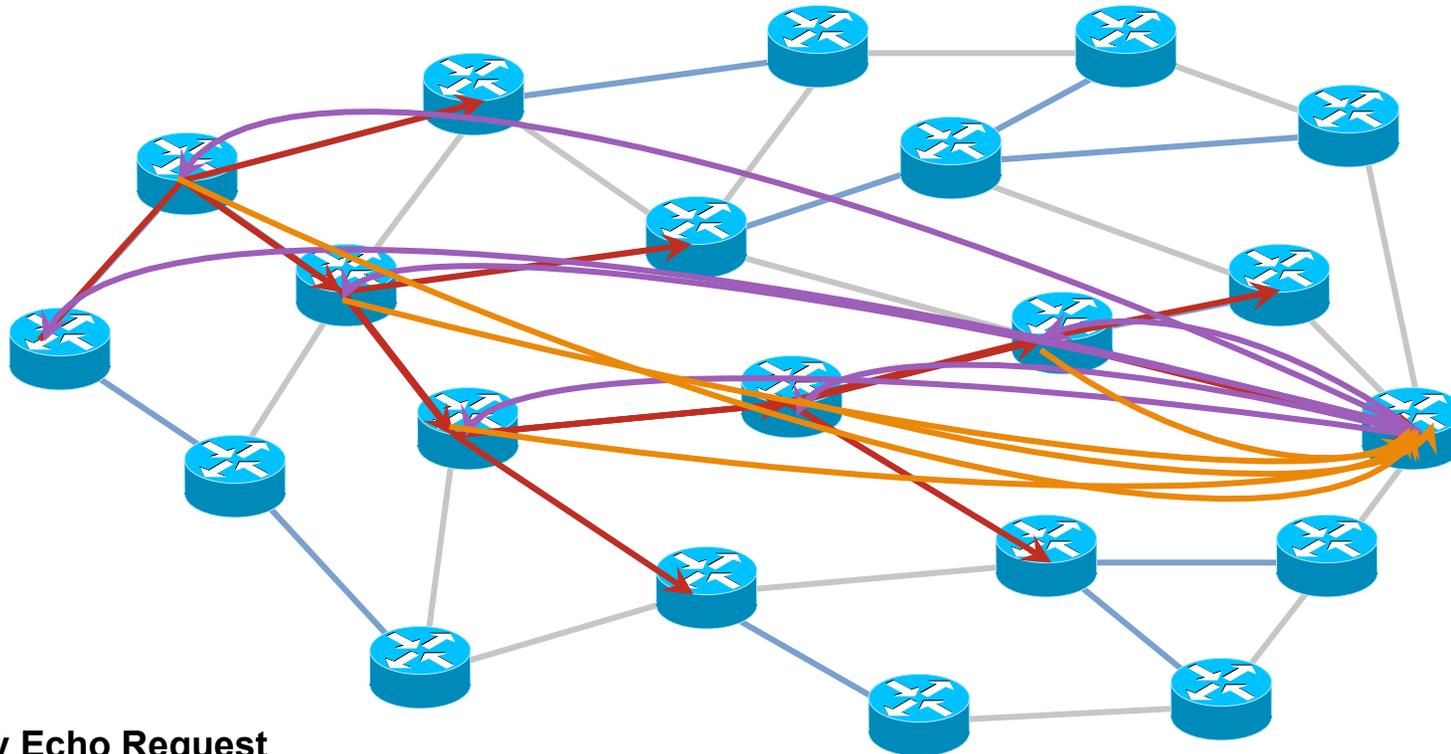
Multicast LDP Tracing

- mLDP LSPs are initiated by the leaf nodes
- Root node may not know the tree topology
- Failure is most likely to be detected at leaf
- Leaf is a logical place to begin tracing

Previous Hop Address Object

- Needed to communicate the address by which the upstream node knows this node
- Generally this would be the address used in the protocol by which the label for the bottom most FEC was exchanged
- This address would be used to specify the next hop when doing a reverse trace

mLDP Traceroute



Proxy Echo Request

Proxy Echo Reply

Echo Request

Echo Reply

- Tree is traced from leaf to root
- PHOPs are learned as trace progresses

Next Steps

- Not anticipating further functional changes
Fairly stable (this is the 7th revision)
- Would like WG review
- Request (re-)adoption as a WG document