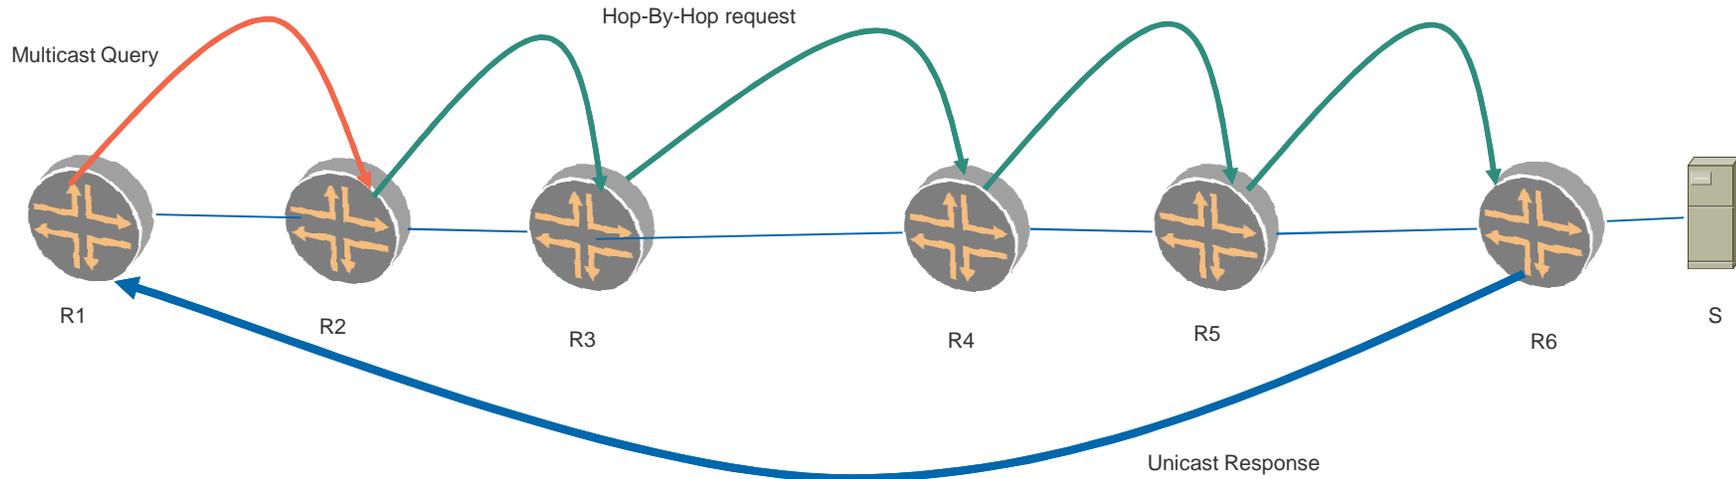


MULTICAST TRACEROUTE FOR MVPNS

ROBERT KEBLER
PAVAN KURAPATI

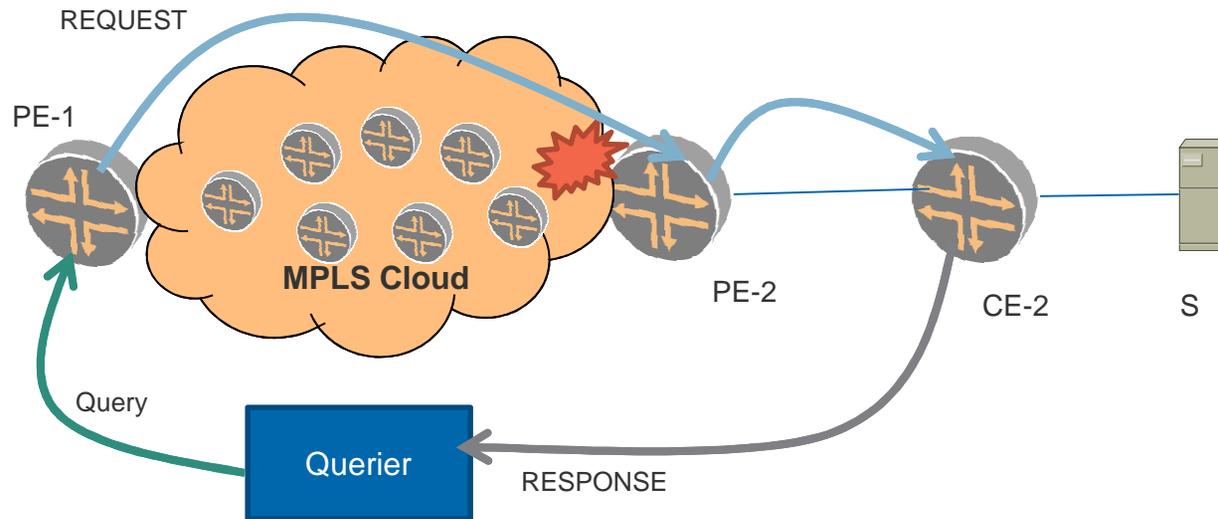


MTRACE – THE CURRENT VERSION



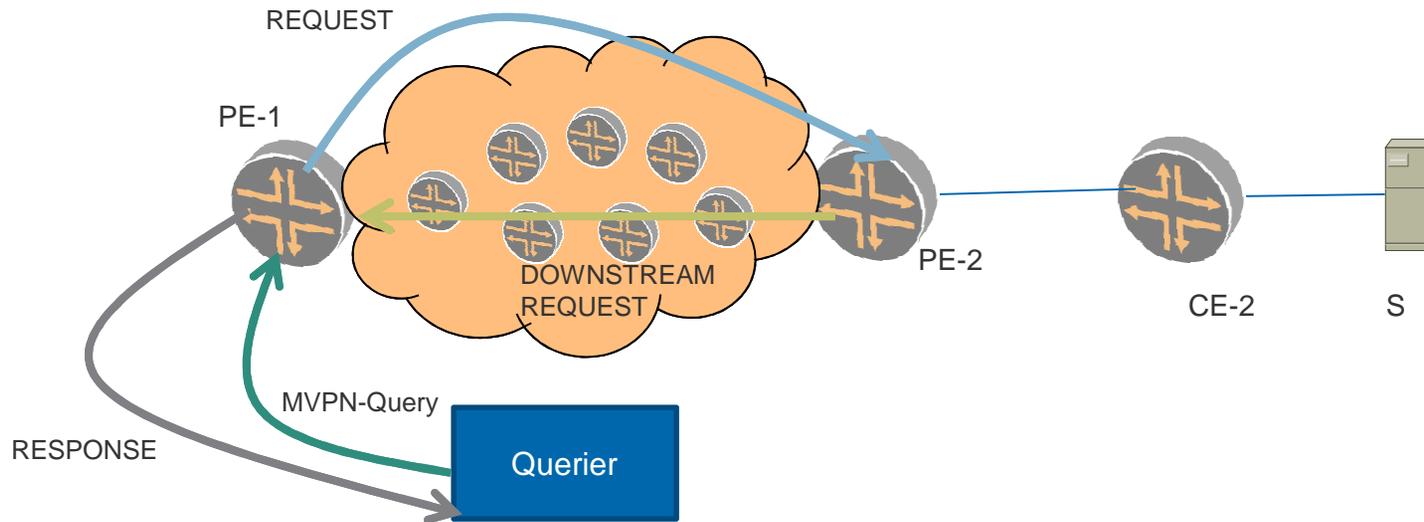
- Starts at the “Last-Hop” and travels hop-by-hop upstream towards the source.
- Mtrace travels upstream along the same interfaces that the data would use
- Multicast problems interpreted based on the response block and error codes added

MTRACE – APPLIED TO BGP-MVPN DEPLOYMENTS



- Routers in the path are not directly connected through interfaces.
- Routers separated several hops away through tunnels.
- mtrace message does not travel on same path as data.
- For example, PE-2 may be sending traffic on wrong tunnel because of which traffic may never reach PE-1. However, Mtrace will result in successful response
- mtrace extends beyond the Provider Network.
- No BGP-MVPN specific error codes available

MVPN MTRACE – WHAT IS NEW?

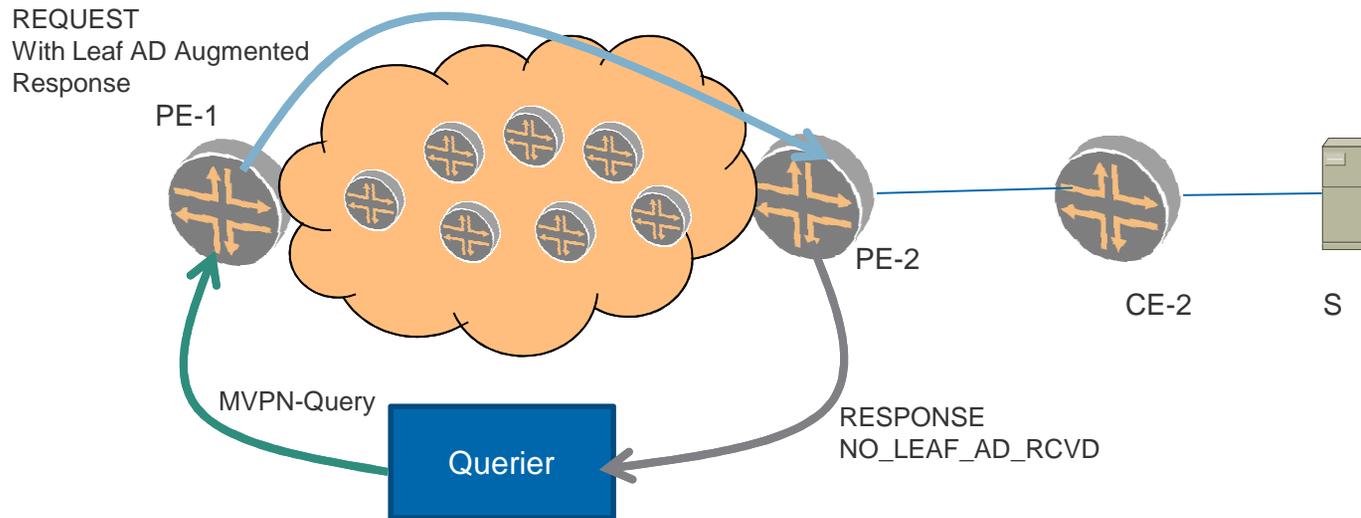


- Initiated within provider space
- New 'MVPN Extended Query block' to identify MVPN specific Mtrace Query
- New Mtracev2 TLV called DOWNSTREAM REQUEST defined
- Downstream Request is sent on same provider tunnel that the traffic is sent on
 - Ensures the verification of data path
- New Augmented Response Blocks to carry Leaf-AD and PMSI tunnel Attributes which is used by ingress PE for control plane validations
- New MVPN specific error codes added to identify MVPN related issues
- Applied to various MVPN deployment scenarios: inter-AS, inter-area, Virtual Hub&Spoke

MVPN SPECIFIC ERROR CODES

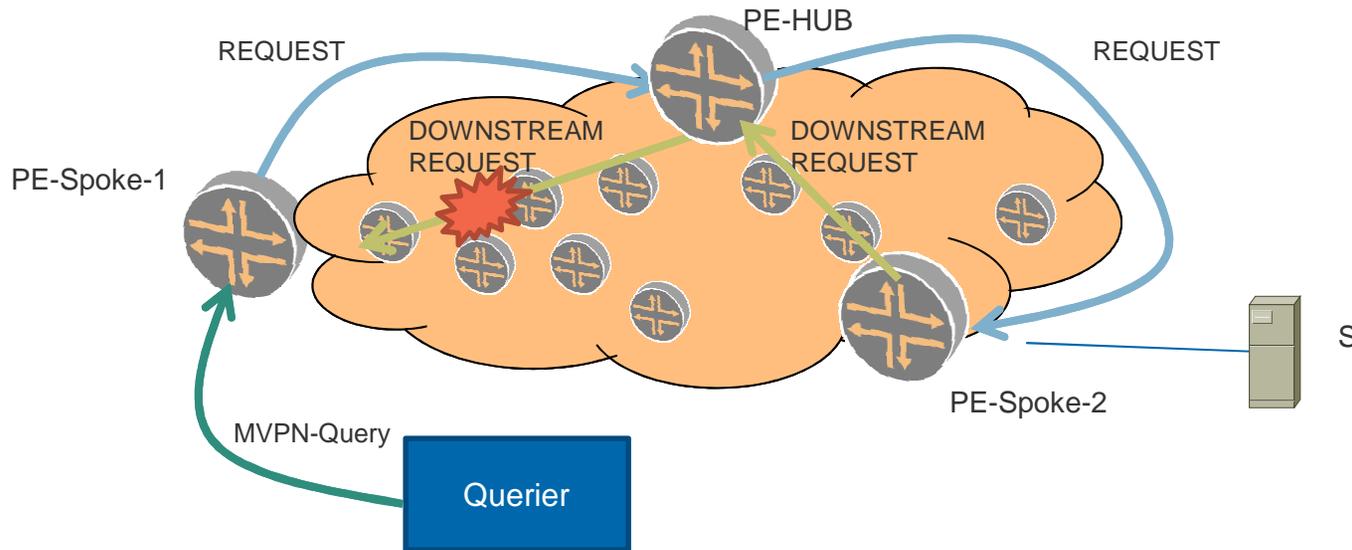
Name	Description
PTUNNEL_DOWN	The provide tunnel for this S,G is down
NO_LEAF_AD_RCVD	The S-PMSI has not been joined by downstream neighbor
BAD_LEAF_AD	The Leaf A-D route does not match the expected values
BAD_RD	The RD is known to not exist on this PE
UNEXPECTED_MVPN	The MVPN traceroute message is unexpected
BAD_PMSI_ATTR_FLAG	Error matching the PMSI attribute flag
BAD_PMSI_ATTR_TYPE	Error matching the PMSI attribute type
BAD_PMSI_ATTR_LABEL	Error matching the PMSI attribute label
BAD_PMSI_ATTR_ID	Error matching the PMSI attribute tunnel identifier
UNKNOWN_INTER_AS	Could not locate the Inter-AS provider tunnel segment.
NO_UPSTREAM_PE	No valid upstream PE or route
NO_CMCAST_STATE	No C-Mcast route for the requested query
NO_WILD_CARD_SPMSI_AD_RCVD	No Wild Card SPMSI SPMSI AD is received from the upstream PE
NO_WILD_CARD_SPMSI_LEAD_AD_SENT	PE did not send LEAF-AD route for the wild card SPMSI

MVPN MTRACE – ERROR CONDITIONS



- MVPN Mtrace handles MVPN specific error conditions.
- For example, consider situation where PE-1 sent a Leaf AD route in response to SPMSI AD advertised by PE-2
- Due to BGP related issues at RRs or PEs, PE-2 did not receive the LEAF-AD route from PE-1
- A Mtrace Query for this (C-S,C-G) from PE-1 would result in a response with NO_LEAF_AD_RCVD. This is sent from PE-2 to the Querier

MVPN MTRACE – ERROR CONDITIONS



- Querier sends the Mtrace Query towards LHR (PE-Spoke-1).
- Egress PE sends Request to V-HUB.
- V-HUB sends the request to Ingress Spoke PE.
- Ingress Spoke PE sends Downstream Request to V-HUB.
- V-HUB sets the 'D' bit in its PMSI Tunnel Attributes Augmented Response Block
- V-HUB sends Downstream request to ingress spoke, but not received by the ingress PE.
- The result of first 4 steps is that Querier did not receive the response. This makes the Querier fall back to TTL method.
- Querier reduces the TTL and the result will show that the hop from V-HUB to ingress spoke is missing thereby pointing the issue at the right place.

NEXT STEPS

- Request Working Group to adopt this as a work item
- Review and Feedback. L3VPN and MBONED
- <http://www.ietf.org/id/draft-kebler-kurapati-l3vpn-mvpn-mtrace-00.txt>