

UDP Path for In-band Performance Measurement for Segment Routing Networks

draft-gandhi-spring-udp-pm-01

Rakesh Gandhi - Cisco Systems (rgandhi@cisco.com) - Presenter

Clarence Filsfils - Cisco Systems (cfilsfil@cisco.com)

Sagar Soni - Cisco Systems (sagsoni@cisco.com)

Daniel Voyer - Bell Canada (daniel.voyer@bell.ca)

Stefano Salsano - Universita di Roma "Tor Vergata" (stefano.salsano@uniroma2.it)

Pier Luigi Ventre - CNIT (pierluigi.ventre@cnit.it)

Patrick Khordoc - Cisco Systems (pkhordoc@cisco.com)

Zafar Ali - Cisco Systems (zali@cisco.com)

Daniel Bernier - Bell Canada (daniel.bernier@bell.ca)

Dirk Steinberg - Steinberg Consulting (dws@dirksteinberg.de)

Agenda

- Requirements and Scope
- Probe Query Message
- Probe Response Message
- Return Path Segment List
- ECMP Support
- Next Steps

Requirements and Scope

Requirements:

- Delay and Loss Performance Measurement (PM) for SR links and end-to-end SR Policies
- Agnostic to data plane (SR-MPLS/SRv6/IP)
- No need to bootstrap PM session (e.g., to negotiate UDP port) - spirit of SR
- Stateless on egress node - spirit of SR
- One-way and two-way measurements
- Handle ECMP for SR Policies

Scope:

- Use RFC 6374 defined **probe message formats**
- Use RFC 7876 (IP/UDP return path) defined probe response messages
- Define IP/UDP path for PM probe query messages

Probe Query Messages

- IP/UDP path is defined for PM probe query messages for delay and loss measurements for SR links and end-to-end SR Policies.
- For **end-to-end** performance measurement, the probe query messages are sent in-band with MPLS label stack for SR-MPLS Policies and SRv6 SRH with SID list for SRv6 Policies.
- Payload contains [RFC6374] defined message for DM or LM.
- UDP port IANA-TBA1 is used for identifying DM probe packets.
- UDP port IANA-TBD2 is used for identifying LM probe packets.

```
+-----+
| IP Header |
. Source IP Address = Querier IPv4 or IPv6 Address .
. Destination IP Address = Responder IPv4 or IPv6 Address .
. Protocol = UDP .
. IP TTL = 1 .
. Router Alert Option Not Set .
.
+-----+
| UDP Header |
. Source Port = As chosen by Querier .
. Destination Port = TBA1 by IANA for DM, TBA2 for LM .
.
+-----+
| Payload = Message as specified in RFC 6374 for DM and LM |
.
+-----+
```

Probe Response Messages

- Probe response messages can be sent in-band (two-way measurement) or out-of-band (one-way measurement) for SR links and SR Policies.
- Use the information from the UDP Return Object (URO) TLV [RFC7876] from the received Probe query message payload, otherwise use the IP/UDP information (Source IP Address and Source UDP port) from the received Probe query message header.

```
+-----+
| IP Header
. Source IP Address = Responder IPv4 or IPv6 Address
. Destination IP Address = URO.Address
. Protocol = UDP
. Router Alert Option Not Set
.
+-----+
| UDP Header
. Source Port = As chosen by Responder
. Destination Port = URO.UDP-Destination-Port
.
+-----+
| Message as specified in RFC 6374 Section 3.2 for DM, or
. Message as specified in RFC 6374 Section 3.1 for LM
.
+-----+
```

```
+-----+
| IP Header
. Source IP Address = Responder IPv4 or IPv6 Address
. Destination IP Address = Source IP Address from Query
. Protocol = UDP
. Router Alert Option Not Set
.
+-----+
| UDP Header
. Source Port = As chosen by Responder
. Destination Port = Source Port from Query
.
+-----+
| Message as specified in RFC 6374 Section 3.2 for DM, or
. Message as specified in RFC 6374 Section 3.1 for LM
.
+-----+
```

Return Path Segment List

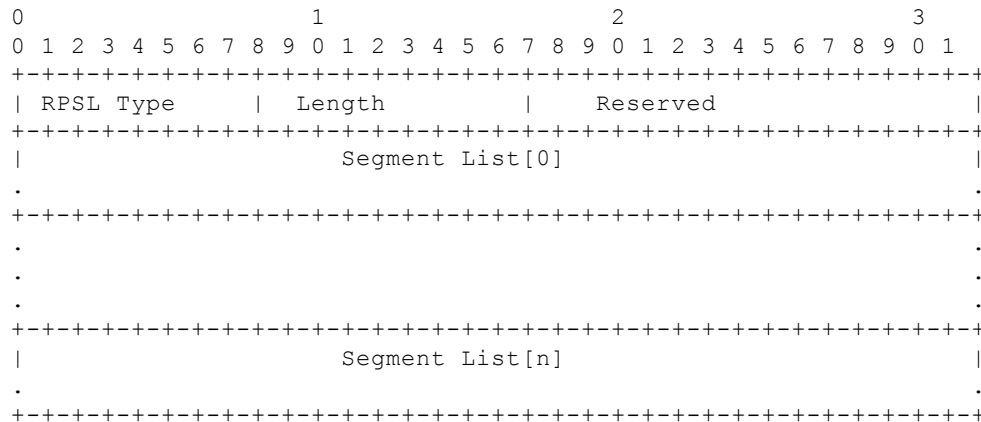
- For two-way end-to-end performance measurement of SR Policies, the responder node needs to send the probe response messages in-band on a specific reverse SR path.
- Querier node can request the responder node to send the probe response messages back on a given reverse path (e.g. co-routed path) by adding a Return Path Segment List (RPSL) TLV in the probe query messages.
- This does not require any SR Policy state on the egress node.

- TBA3: SR-MPLS RPSL

- TBA4: SRv6 RPSL

- TBA5: SR-MPLS BSID

- TBA6: SRv6 BSID



ECMP Support

- SR Policy can have ECMP between the ingress and transit nodes, between transit nodes and between transit and egress nodes.
- Sending PM probe queries that can take advantage of the hashing function in forwarding plane.
- Existing forwarding mechanisms are applicable to PM probe messages:
 - For IPv4 and SR-MPLS
 - Different Destination Addresses in the range of 127/8 or different Source Addresses or different Source UDP ports in IP/UDP header.
 - For SR-MPLS
 - Entropy label.
 - For IPv6
 - 3-tuple of Flow Label, Source Address and Destination Address fields in the IPv6 header.
 - For SRv6
 - Flow Label in SRH.

Next Steps

- Welcome your comments and suggestions
- Implementations of building blocks already exist (e.g. RFC6374, IP/UDP paths for probes)
- Request for WG adoption

Thank you.