Simple TWAMP (STAMP) Extensions for Segment Routing Networks

draft-gandhi-ippm-stamp-srpm-02

Rakesh Gandhi - Cisco Systems (rgandhi@cisco.com) - Presenter
Clarence Filsfils - Cisco Systems (cfilsfil@cisco.com)
Daniel Voyer - Bell Canada (daniel.voyer@bell.ca)
Mach(Guoyi) Chen - Huawei (mach.chen@huawei.com)
Bart Janssens - Colt (Bart.Janssens@colt.net)

Agenda

- Requirements and Scope
- Summary of Extensions
- Next Steps

Requirements, Goals and Scope

Requirements:

- In-band Performance Measurement for links and SR paths

Goals:

- Avoid maintaining each test session on Session-Reflector
- Avoid control protocol for signaling dynamic parameters

Scope:

- STAMP [RFC 8762]
- STAMP Extensions [RFC 8972]

Updates Since Version-00

Updates:

- ✓ Updated terminology to align with STAMP
- ✓ Moved direct measurement messages to *draft-gandhi-ippm-simple-direct-loss*
- ✓ Moved Control Code to Return Path TLV
- ✓ Various editorial changes to address review comments

Open Items:

- None

STAMP Destination Node Address TLV

Destination Node Address TLV (value TBA1):

- Indicates the address of the intended destination of the Session-Sender test packet
- STAMP Session-Reflector that supports this TLV, **MUST NOT** send reply if it is not the intended destination of the Session-Sender test packet
- Useful when test packet is sent with 127/8 destination address (e.g. sweeping ECMP paths)

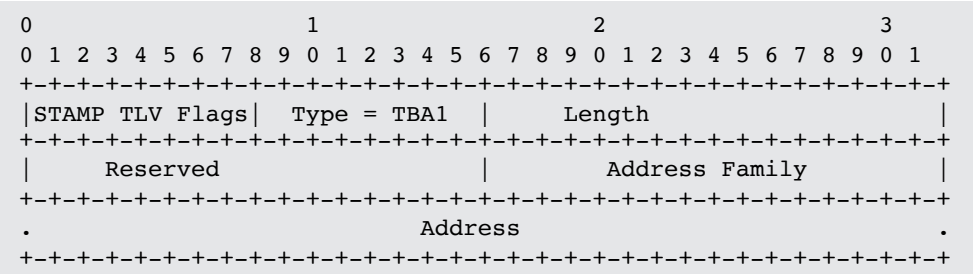

Figure: Destination Node Address TLV Format

STAMP Return Path TLV

Return Path TLV (value TBA2) to carry Sub-TLV for return path:

Return Path Sub-TLVs Types:

- Type (value 1): Return Path Control Code. Reply test packet based on the control code flags
  - 0x0: No Reply Requested
  - 0x1: In-band Reply Requested
- Type (value 2): Return Address. Destination address for the reply; different than the Source Address in the Session-Sender test packet
- Type (value 3): SR-MPLS Label Stack of the Return SR Path
- Type (value 4): SR-MPLS Binding SID [draft-ietf-pce-binding-label-sid] of the Return SR Policy
- Type (value 5): SRv6 Segment List of the Return SR Path
- Type (value 6): SRv6 Binding SID [draft-ietf-pce-binding-label-sid] of the Return SR Policy

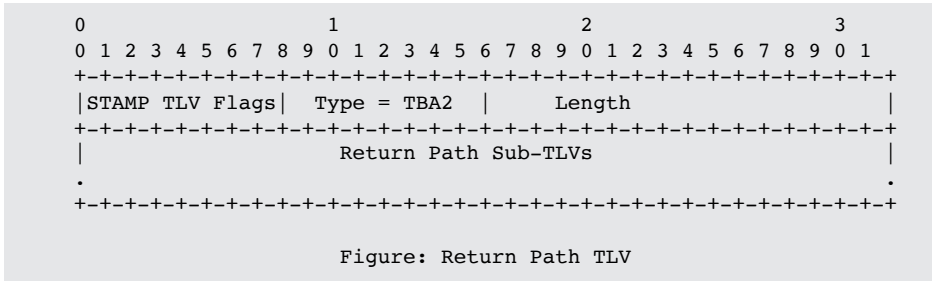

Figure: Return Path TLV

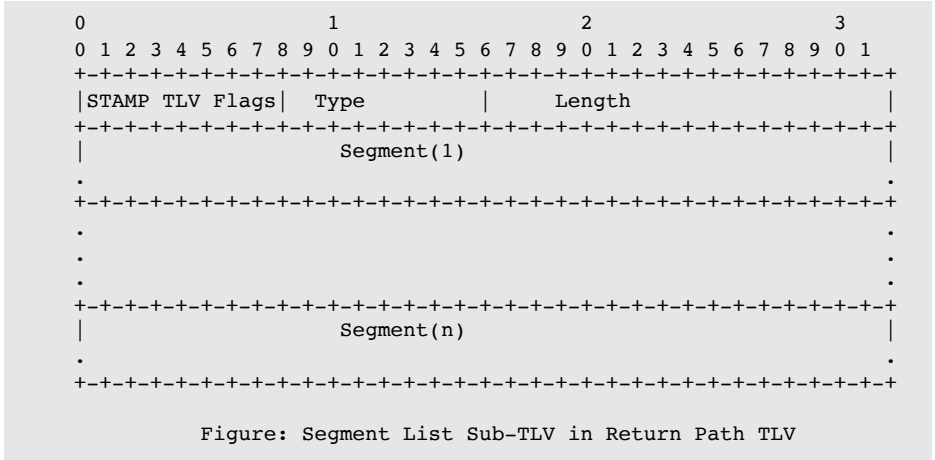

Figure: Segment List Sub-TLV in Return Path TLV

STAMP Return Path Control Code Sub-TLV - Usage

- Avoid maintaining each test session (session id, source-address) on Session-Reflector
- In-band Reply Requested:
  - For link delay measurement
    - STAMP Session-Reflector transmits test packet in-band on the same incoming link in the reverse direction
    - Link can be Virtual, LAG or LAG member
- No Reply Requested:
  - STAMP Session-Reflector does not transmit reply test packet to the STAMP Session-Sender and terminates the Session-Sender test packet

STAMP Return Address Sub-TLV - Usage

- Avoid maintaining each test session (session id, source-address) on Session-Reflector
- STAMP Session-Reflector reply test packet may be transmitted to a different node than the Session-Sender
- STAMP Session-Sender can specify in the test packet the receiving destination address for the STAMP Session-Reflector reply test packet

STAMP Return Path Segment List Sub-TLVs - Usage

- For an SR path, STAMP Session-Reflector reply test packet may need to be sent in-band on a specific return SR path
- Dynamically computed SR paths can change based on topology change, link/node failure, etc. in the network
- Avoid signaling and maintaining dynamic state on STAMP Session-Reflector for the return path for each STAMP test session (each session-id, source-address)
  - Can be order of 10K SR Policy (that can also have ECMPs)

Next Steps

- Welcome your comments and suggestions
- Requesting WG adoption

Thank you

Performance Measurement Using Simple TWAMP for Segment Routing Networks

draft-gandhi-spring-stamp-srpm-05

Rakesh Gandhi - Cisco Systems (rgandhi@cisco.com) - Presenter
Clarence Filsfils - Cisco Systems (cfilsfil@cisco.com)
Daniel Voyer - Bell Canada (daniel.voyer@bell.ca)
Mach(Guoyi) Chen - Huawei (mach.chen@huawei.com)
Bart Janssens - Colt (Bart.Janssens@colt.net)

Agenda

- Requirements and Scope
- Summary of Procedure
- Next Steps

Requirements, Goals and Scope

Requirements:

- In-band Performance Delay and Loss Measurement
  - ✓ Links and end-to-end P2P/P2MP SR paths
    - ✓ Links include physical, virtual, LAG, LAG member links
  - ✓ Applicable to SR-MPLS/SRv6 data planes
- One-way, two-way, round-trip delay and packet loss metrics

Goals:

- Avoid maintaining each test session on Session-Reflector
- Avoid control protocol for signaling dynamic parameters

Scope:

- STAMP [RFC 8762]
- STAMP Extensions [RFC 8972]
- STAMP Extensions for SR [draft-gandhi-ippm-stamp-srpm]

Updates Since Version-02

Updates:

- ✓ Draft status - Informational
- ✓ Updated terminology to align with STAMP
- ✓ Added (synthetic) packet loss section
- ✓ Removed stand-alone direct measurement messages
- ✓ Removed text for IPv6/UDP test packet with zero checksum
- ✓ Various editorial changes to address review comments

Open Items:

- None

Example STAMP Reference Model

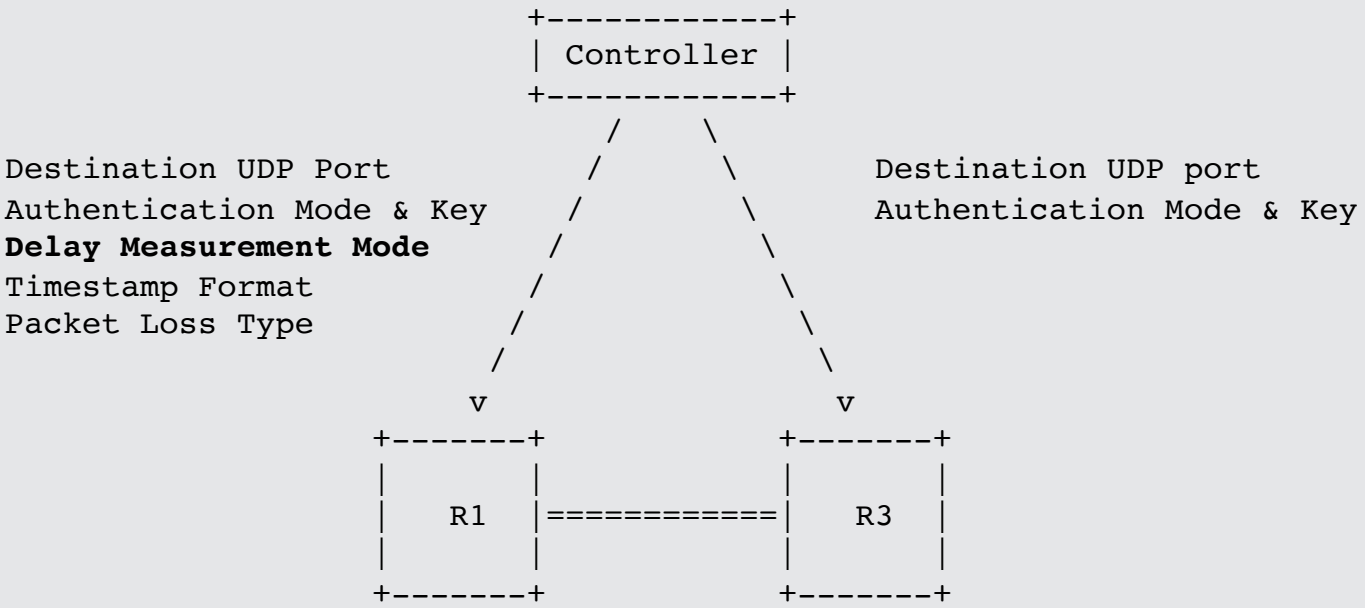

STAMP Session-Sender Test Packet for Links

- For links, STAMP Session-Sender test packets are transmitted over the links using local and remote link addresses
- User-configured destination UDP port is used for STAMP test packets (or port 862)
- IPv4 TTL /IPv6 Hop-limit is set to 1
- Applicable to physical, virtual, LAG, LAG member links

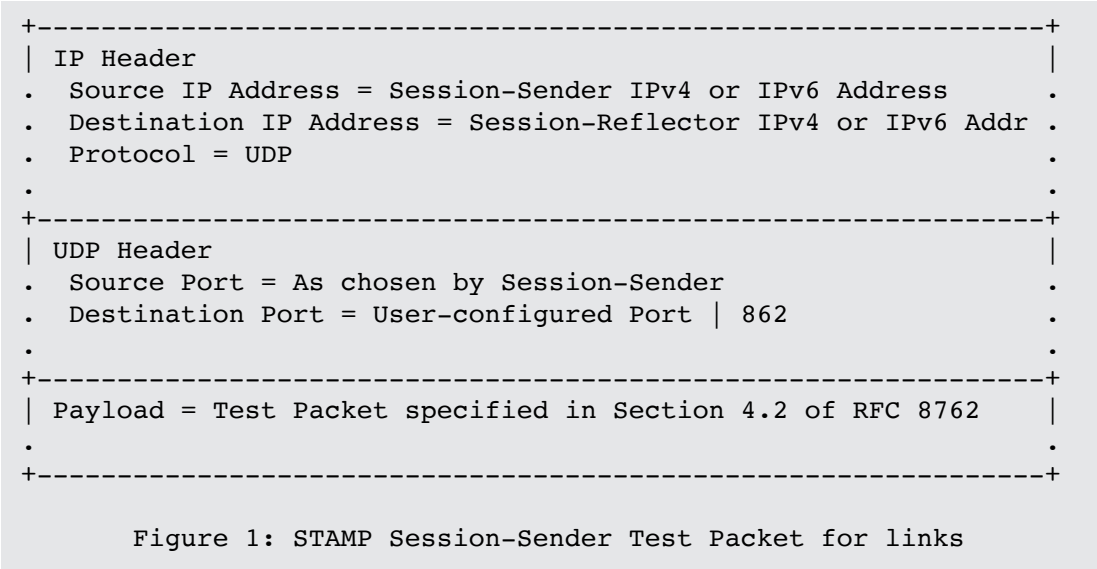

Figure 1: STAMP Session-Sender Test Packet for links

STAMP Session-Sender Test Packet for SR-MPLS and SRv6 Policy

- For end-to-end SR Policy, STAMP Session-Sender test packets are transmitted with:
  1. MPLS label stack of SR-MPLS Policy
  2. SRv6 SRH [RFC 8754] with Segment List of SRv6 Policy
     - Using upper-layer processing (for UDP header) defined in SRv6 network programming
- User-configured destination UDP port is used for STAMP test packets (or port 862)
- IPv4 TTL/IPv6 Hop-limit is set to 255
- Color-Only Destination Steering:
  - IPv4
    - Destination Address in 127/8 range
    - TTL is set to 1
  - IPv6
    - Destination Address set to ::1/128
    - Hop Limit is set to 1

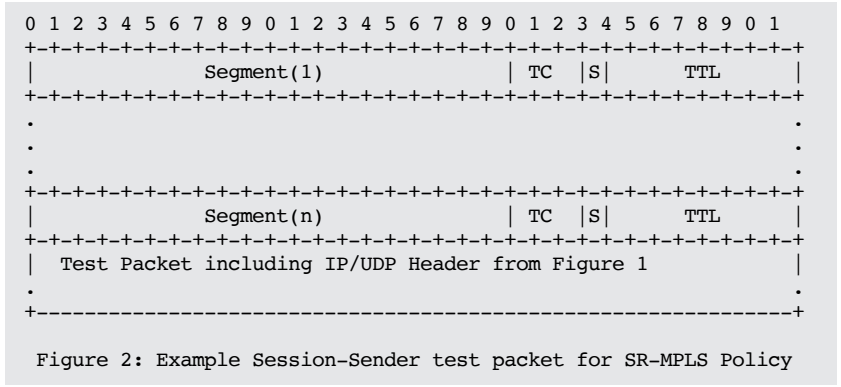

Figure 2: Example Session-Sender test packet for SR-MPLS Policy

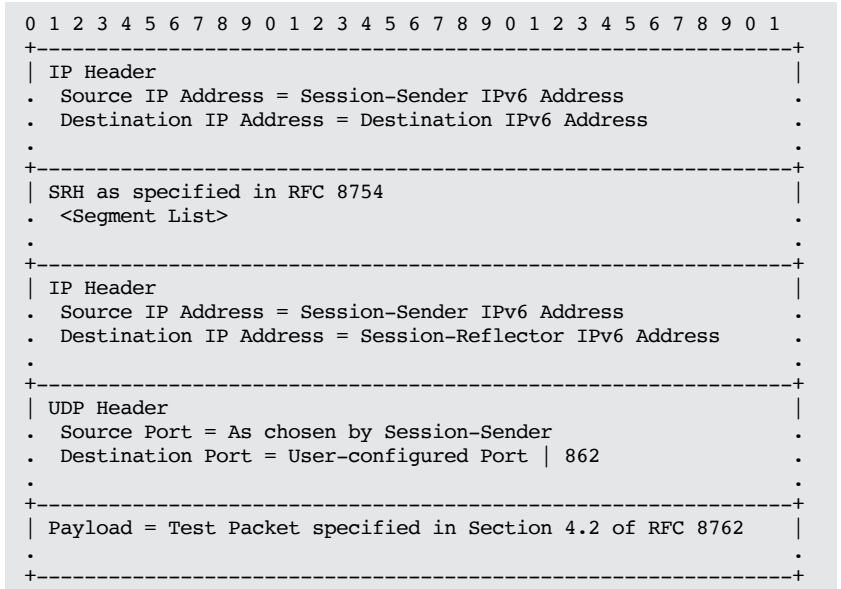

Figure 3: Example Session-Sender test packet for SRv6 Policy

STAMP Session-Sender Test Packet for P2MP SR-MPLS Policy

- For end-to-end P2MP SR-MPLS Policy, STAMP Session-Sender test packets are transmitted with:
  - Tree-SID of the P2MP SR-MPLS Policy
  - IPv4 destination address selected from 127/8 range
  - IPv4 TTL is set to 1

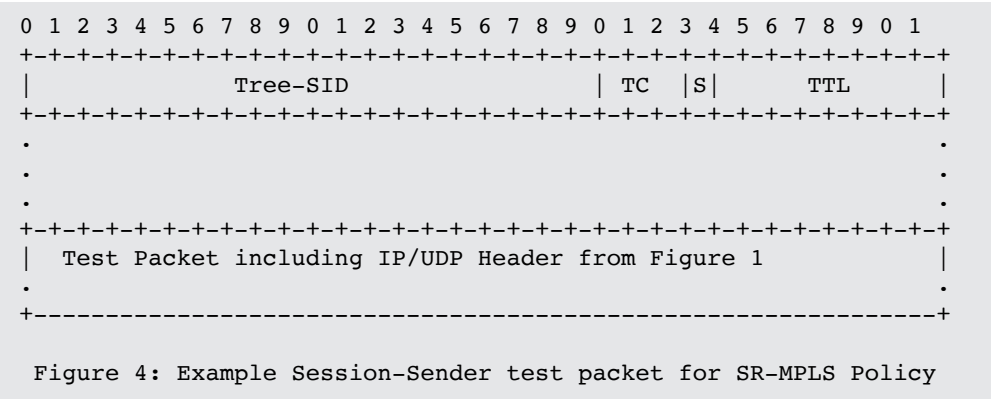

Figure 4: Example Session-Sender test packet for SR-MPLS Policy

STAMP Session-Reflector Test Packet

- STAMP Session-Reflector reply test packet is transmitted using the IP/UDP information from the received test packet.

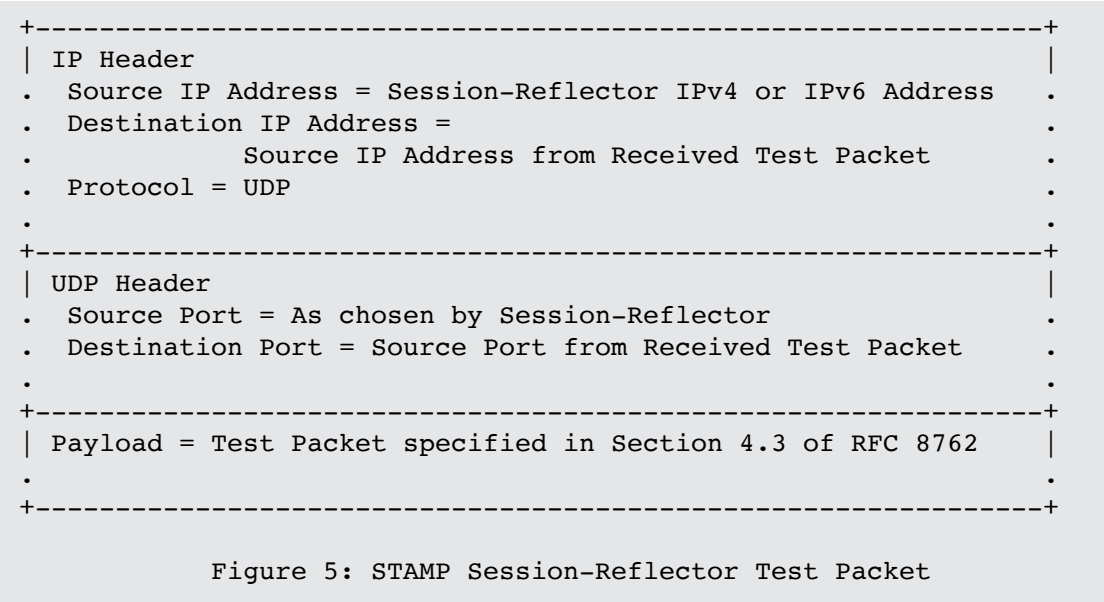

Figure 5: STAMP Session-Reflector Test Packet

Performance Measurement Modes

- Need to measure in-band one-way, two-way and round-trip delay metrics in SR networks
- One-way Delay Measurement Mode
  - Existing (default) behavior
- Two-way Delay Measurement Mode
  - STAMP Session-Reflector test packet sent “in-band” on reverse path
  - Avoid per test session state on Session-Reflector
  - Link: Use Control Code Sub-TLV in the Return Path TLV from the received test packet.
  - E2E SR path: Use Segment List Sub-TLV in the Return Path TLV from the received test packet.
- Round-trip Delay Measurement Mode
  - STAMP Session-Sender test packet sent in loopback mode, carries the return path in the packet header

ECMP Support for SR Path

- SR Path can have ECMP between the ingress and transit nodes, between transit nodes and between transit and egress nodes
- Sending STAMP test packets that can take advantage of the hashing function in forwarding plane
- Existing forwarding mechanisms are applicable to test packets. Examples are:
  - For IPv4
    - Sweeping destination address in IPv4 header (e.g. 127/8)
    - Identify intended actual destination node in “Destination Node Address TLV”
  - For IPv6
    - Sweeping flow label in IPv6 header

Example PM Metrics

- Compute following example (one-way, two-way, round-trip) delay metrics:
  - Minimum delay
  - Maximum delay
  - Average delay
  - Delay variance
- Compute following example loss metrics:
  - Packet loss (i.e., synthetic packet loss)
  - Direct measurement packet counters
  - Session status succeeded/failed (i.e., measurement is active)

Next Steps

- Welcome your comments and suggestions
- Requesting WG adoption

Thank you

Simple Two-Way Direct Loss Measurement Procedure

draft-gandhi-ippm-simple-direct-loss-00

Rakesh Gandhi - Cisco Systems (rgandhi@cisco.com) - Presenter
Clarence Filsfils - Cisco Systems (cfilsfil@cisco.com)
Daniel Voyer - Bell Canada (daniel.voyer@bell.ca)
Mach(Guoyi) Chen - Huawei (mach.chen@huawei.com)
Bart Janssens - Colt (Bart.Janssens@colt.net)
Stefano Salsano - Universita di Roma "Tor Vergata" (stefano.salsano@uniroma2.it)

Agenda

- Requirements and Scope
- Summary
- Next Steps

Requirements and Scope

Requirements:

- Direct Loss Measurement (DLM) for accurate data packet loss
- Support Alternate-Marking Method (AMM) [RFC 8321]
- High scale for number of sessions and faster packet loss detection interval
  - Support hardware implementation

Goals:

- Avoid maintaining each test session on Session-Reflector
- Avoid control protocol for signaling dynamic parameters

Scope:

- Follow STAMP [RFC 8762] approach

Alternate Marking Method for Packet Loss

- RFC 8321 - Alternate-Marking Method for Passive and Hybrid Performance Monitoring
- RFC 8957 - Synonymous Flow Label Framework

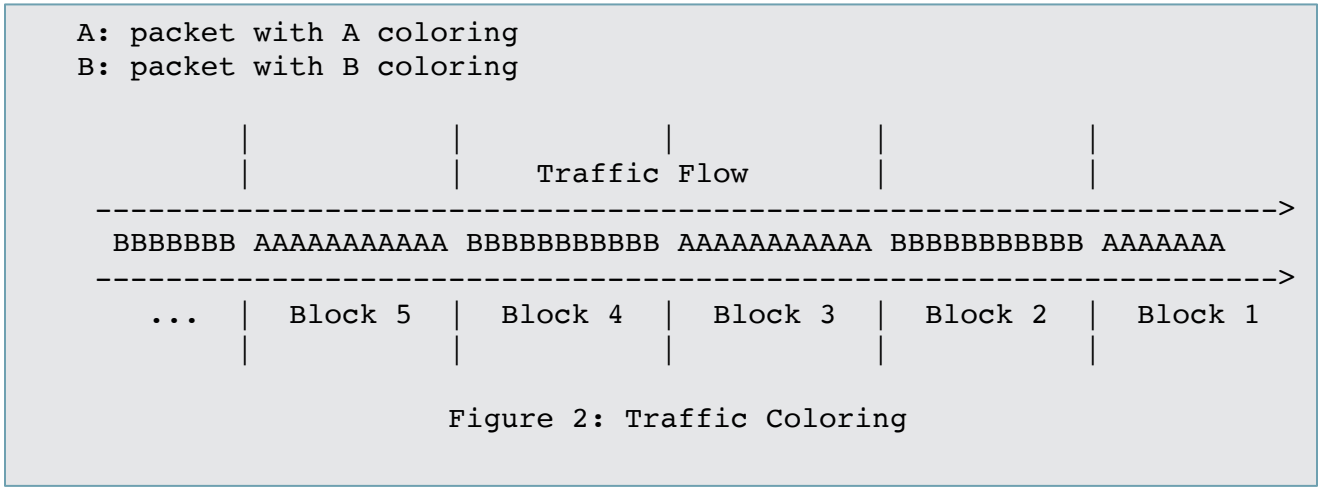

Figure 2: Traffic Coloring

Case 1. STAMP Test Packets with Direct Measurement TLV

Figure: STAMP Session-Sender Test Packet Format

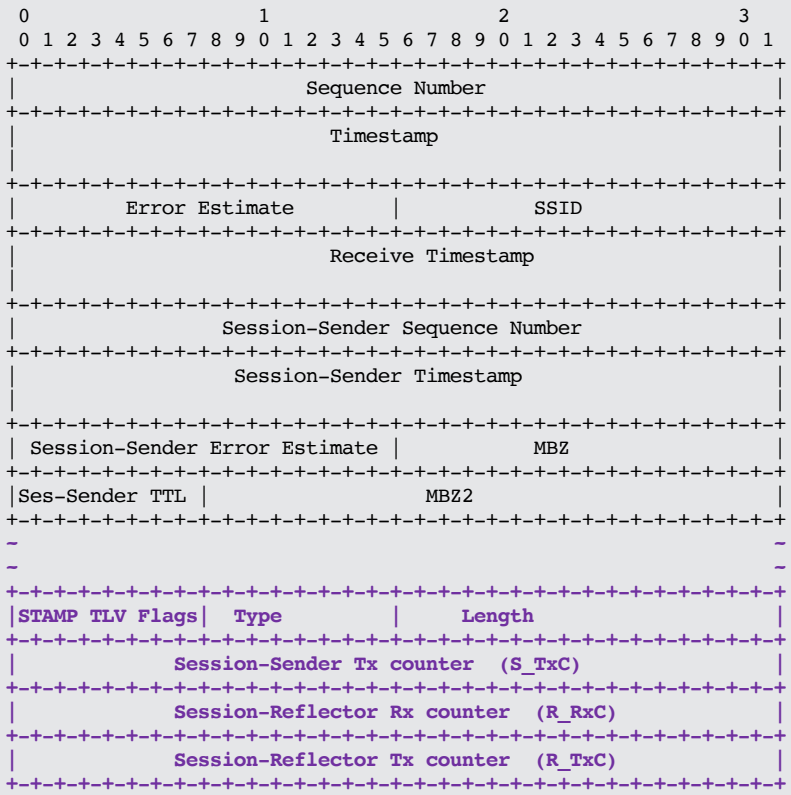

Figure: STAMP Session-Reflector Test Packet Format

Case 2. Direct Loss Measurement Probe Packet for Data Packet Loss Detection

- Base Direct Loss Measurement probe packet defined
  - Hardware efficient counter updating
    - Well-known locations for traffic counters
  - Block number of the counters for alternate-marking method [RFC 8321]
  - Traffic class of the counters for per class packet loss
  - DLM probe packet is also defined for authenticated mode
- User-configured destination UDP Port is used for identifying DLM probe packets (different than port 862 and the one used by STAMP)
- Sequence Number allows to monitor DLM session status, out of order probe packets and probe packet drops
- Flags
  - X set to 1 for 64-Bit Counter, set to 0 for 32-Bit Counter
  - B set to 1 for Byte Counter, set to 0 for Packet Counter
  - T set to 1 for Sender-DSCP scoped Counter

Figure: Session-Reflector Direct Loss Measurement Probe Packet

Direct Measurement TLV vs. Direct Loss Measurement Probe Packet

Attributes	Case 1. STAMP Direct Measurement TLV	*Case 1a. Define New STAMP Direct Measurement TLV2?*	Case 2. Direct Loss Measurement Probe Packet
Need to write timestamp (clock sync needed for one-way delay)	Yes	*Yes*	No
Counter at fixed location in the probe packet for hardware-based counter update	No (TLV-based)	*No (TLV-based)*	Yes
Reply probe packets with counter at fixed location for hardware-based counter update	No	*No*	Yes
Need to scan for DM TLV in each received probe packet on Session-Reflector in hardware (there can be multiple TLVs)	Yes	*Yes*	No
32-bit and 64-bit Byte counters	No	*Yes*	Yes
64-bit packet counters	No	*Yes*	Yes
Alternate-marking method packet loss - using block number for counters (out-of-order data packet support)	No	*Yes*	Yes
Per Traffic Class Counters	No	*Yes*	Yes

Data Packet Loss Calculation

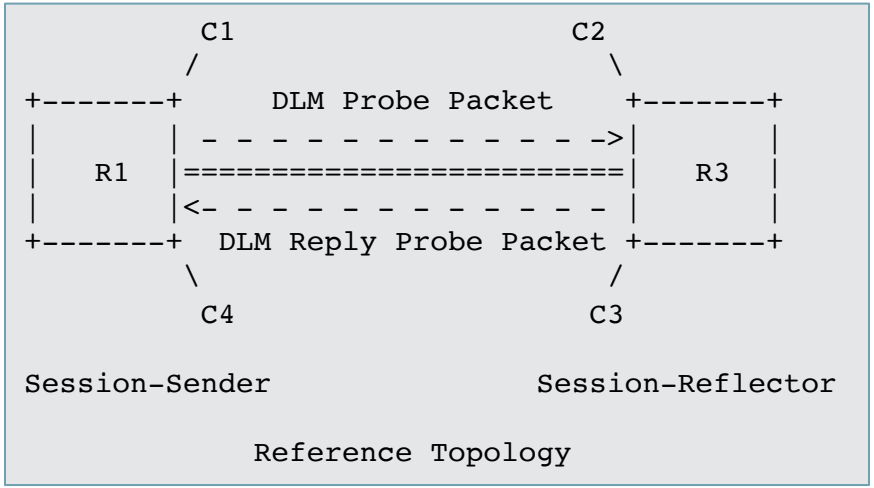

- Using the Counters C1, C2, C3 and C4 as per reference topology, from the n^{th} and $(n-1)^{th}$ Direct Loss Measurement probe packets.
  - Transmit Loss TxL[n-1, n] = (C1[n] - C1[n-1]) - (C2[n] - C2[n-1])
  - Receive Loss RxL[n-1, n] = (C3[n] - C3[n-1]) - (C4[n] - C4[n-1])
- When using Alternate-Marking Method, all Counters used for the loss calculation belongs to the same Block Number, as described in Section 3.1 of [RFC8321].

Next Steps

- Welcome your comments and suggestions
- Requesting WG adoption
- *Define New STAMP Direct Measurement TLV2?*

Thank you