

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: 12 February 2024

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
N. Vaghamshi
Reliance
M. Nagarajah
Telstra
R. Foote
Nokia
M. Chen
Huawei
A. Dhamija
Rakuten
11 August 2023

Enhanced Performance Measurement Using Simple TWAMP in Segment Routing
Networks
draft-gandhi-spring-enhanced-srpm-05

Abstract

Segment Routing (SR) leverages the source routing paradigm. SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. This document defines procedure for Enhanced Performance Measurement of end-to-end SR paths including SR Policies for both SR-MPLS and SRv6 data planes using Simple Two-Way Active Measurement Protocol (STAMP) defined in RFC 8762. The procedure allows to improve the scale for number of sessions and the interval for measurement.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 12 February 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	3
2.1. Requirements Language	3
2.2. Abbreviations	4
2.3. Reference Topology	4
3. Overview	5
3.1. Enhanced Loopback Mode Enabled with Network Programming Function	5
4. Enhanced Performance Measurement Procedure	6
4.1. Enhanced Performance Measurement Procedure for SR-MPLS Policies	6
4.1.1. Timestamp and Forward Network Action Assignment . . .	8
4.1.2. Node Capability for MNA Sub-Stack with Opcode MNA.TSF	8
4.2. Enhanced Performance Measurement Procedure for SRv6 Policies	8
4.2.1. Timestamp and Forward Endpoint Function Assignment .	10
4.2.2. Node Capability for Timestamp and Forward Endpoint Function	10
5. Security Considerations	10
6. IANA Considerations	10
7. References	10
7.1. Normative References	10
7.2. Informative References	11
Acknowledgments	12
Authors' Addresses	12

1. Introduction

Segment Routing (SR) leverages the source routing paradigm and greatly simplifies network operations for Software Defined Networks (SDNs). SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes [RFC8402]. SR Policies as defined in [RFC9256] are used to steer traffic through a specific, user-defined paths using a stack of Segments. A comprehensive SR Performance Measurement (PM) for delay and packet loss as well as Connectivity Verification (CV) is one of the essential requirements to measure network performance to provide Service Level Agreements (SLAs).

The Simple Two-Way Active Measurement Protocol (STAMP) provides capabilities for the measurement of various performance metrics in IP networks [RFC8762] without the use of a control channel to pre-signal session parameters. As described in [I-D.ietf-spring-stamp-srpm], STAMP can be used for performance measurement of one-way, two-way or round-trip delay and packet loss of end-to-end SR paths.

STAMP requires RFC8762 protocol support on the Session-Reflector to process the received test packets, and hence the received test packets need to be punted from the forwarding fast path and return test packets need to be generated. This limits the scale for number test sessions and the ability to provide faster measurement interval. The loopback measurement mode defined in [I-D.ietf-spring-stamp-srpm] does not require STAMP test packet processing on Session-Reflector, however, it does not provide accurate one-way delay.

This document defines procedure for Enhanced Performance Measurement of end-to-end SR paths including SR Policies for both SR-MPLS and SRv6 data planes, using Simple Two-Way Active Measurement Protocol (STAMP) defined in [RFC8762]. The procedure allows to improve the scale for number of sessions and the interval for measurement.

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

BSID: Binding Segment ID.

ECMP: Equal Cost Multi-Path.

I2E: Ingress-To-Egress.

IHS: Ingress-To-Egress, Hop-By-Hop or Select Scope.

MBZ: Must be Zero.

MNA: MPLS Network Action.

MPLS: Multiprotocol Label Switching.

PTP: Precision Time Protocol.

SID: Segment ID.

SR: Segment Routing.

SRH: Segment Routing Header.

SR-MPLS: Segment Routing with MPLS data plane.

SRv6: Segment Routing with IPv6 data plane.

STAMP: Simple Two-way Active Measurement Protocol.

TC: Traffic Class.

TS: Timestamp.

TSF: Timestamp and Forward.

TTL: Time To Live.

2.3. Reference Topology

In the reference topology shown in Figure 1, the STAMP [RFC8762] Session-Sender S1 initiates a Session-Sender test packet and the Session-Reflector R1 returns the test packet. The return test packet may be transmitted back to the Session-Sender S1 on the same path (same set of links and nodes) or a different path in the reverse direction from the path taken towards the Session-Reflector R1.

The Session-Sender S1 and Session-Reflector R1 are connected via an SR path [RFC8402]. The SR path can be an SR Policy [RFC9256] on node S1 (called head-end) with destination to node R1 (called tail-end).

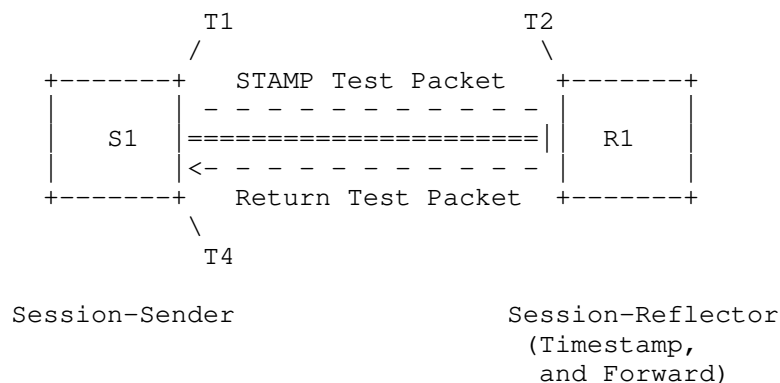


Figure 1: Loopback Mode Enabled with Network Programming Function

3. Overview

As described in [I-D.ietf-spring-stamp-srpm], in loopback mode, the STAMP Session-Sender S1 initiates Session-Sender test packets and the Session-Reflector R1 forwards them back to the Session-Sender S1. The received STAMP test packets are not punted out of the fast path in forwarding at the Session-Reflector. At the Session-Reflector, the loopback function simply makes the necessary changes to the encapsulation including IP and UDP headers to return the STAMP test packet to the Session-Sender S1. No STAMP test session is created on the Session-Reflector R1. As described in [I-D.ietf-spring-stamp-srpm], only loopback delay can be measured in the loopback mode. In SR networks, there is also a need to measure the one-way delay metric to provide low latency services.

3.1. Enhanced Loopback Mode Enabled with Network Programming Function

This document defines a new STAMP measurement mode, called enhanced loopback mode, that is loopback mode enabled with network programming function. In this mode, both transmit (T1) and receive (T2) timestamps in data plane are collected by the Session-Sender test packets as shown in Figure 1. The network programming function optimizes the "operations of punt test packet and generate return test packet" on the Session-Reflector as timestamping is implemented in forwarding fast path in hardware. This helps to achieve higher number of STAMP test session scale and faster measurement interval.

The Session-Sender adds transmit timestamp (T1) in the payload of the Session-Sender test packet. The Session-Reflector adds the receive timestamp (T2) in the payload of the received test packet in forwarding fast path in hardware without punting the test packet (e.g., to slow path or control-plane). The network programming function carried by the test packet enables the Session-Reflector to add the receive timestamp (T2) at the specific offset in the payload of the test packet.

4. Enhanced Performance Measurement Procedure

For enhanced performance monitoring of an end-to-end SR path including SR Policy, STAMP Session-Sender test packets are transmitted in loopback mode enabled with network programming function to timestamp and forward the packet.

For SR Policy, the Session-Sender test packets are transmitted using the Segment List of the Candidate-Path [RFC9256]. When a Candidate-Path has more than one Segment List, multiple Session-Sender test packets MUST be transmitted, one using each Segment List as described in [I-D.ietf-spring-stamp-srpm].

4.1. Enhanced Performance Measurement Procedure for SR-MPLS Policies

An SR-MPLS Policy Candidate-Path may contain a number of Segment Lists. A Session-Sender test packet MUST be transmitted using each Segment List of the SR-MPLS Policy. The content of an example Session-Sender test packet for an end-to-end SR-MPLS Policy is shown in Figure 3.

The loopback measurement mode for SR-MPLS Policies defined in Section 4.2.3.3 of [I-D.ietf-spring-stamp-srpm] is used and enhanced as described below.

MPLS Network Action (MNA) Sub-Stack defined in [I-D.ietf-mpls-mna-hdr] is used for SR-MPLS data plane to enable network programming function of "timestamp and forward" for the received test packet. The MNA Sub-Stack carries the MNA Label (value TBA1) as defined in [I-D.ietf-mpls-mna-hdr]. A new MNA Opcode (value MNA.TSF) is defined for the Timestamp and Forward network action.

In the Session-Sender test packets for SR-MPLS Policies, the MNA Sub-Stack with Opcode MNA.TSF is added in the MPLS header as shown in Figure 3, to collect "Receive Timestamp" field in the payload of the test packet. The Ingress-to-Egress (I2E), Hop-By-Hop (HBH), Select scope (IHS) is set to "I2E" when return path is IP/UDP and set to "Select" when the return path is SR-MPLS. The Network Action Sub-Stack Length (NASL) is set to 0 when there is no Label Stack Entry

(LSE) after the MNA.TSF Opcode in the MNA Sub-Stack. The U flag is set to skip the network action and forward the packet (and not drop the packet). The Label Stack for the reverse direction SR-MPLS path can be added after the MNA Sub-Stack (not shown in the Figure 3) to receive the return test packet on a specific path.

When a Session-Reflector receives a packet with MNA Sub-Stack with Opcode MNA.TSF, after timestamping the packet in STAMP payload at the specific offset, the Session-Reflector pops the MNA Sub-Stack (after completing any other network actions) and forwards the packet using the next label or IP header in the packet (just like the data packets for the normal traffic).

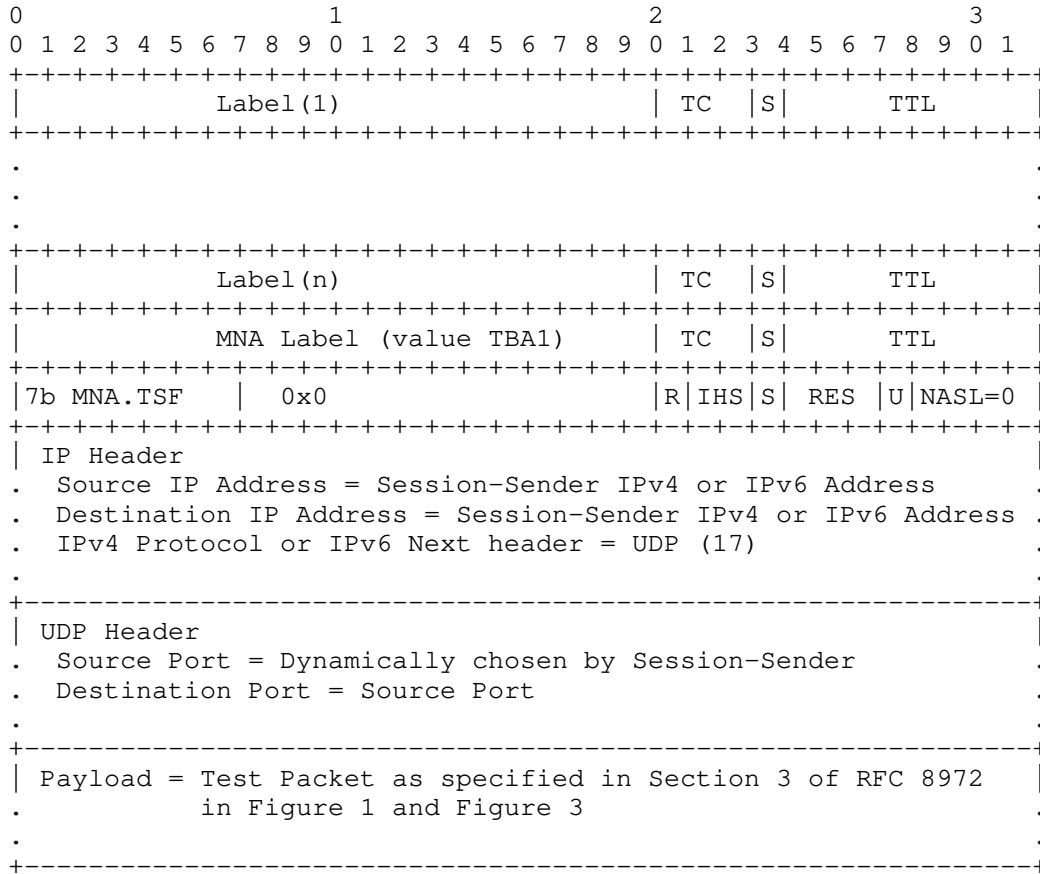


Figure 2: Example STAMP Test Packet with MNA for TSF for SR-MPLS

4.1.1. Timestamp and Forward Network Action Assignment

New MPLS Network Action Opcode is defined called "Timestamp and Forward Network Action, MNA.TSF". The MNA.TSF Opcode is statically configured on the STAMP Session-Reflector node with a value from "Private Use from Range 111-126". The timestamp format for 64-bit PTPv2 and NTP to be added in the STAMP payload is statically configured for MNA.TSF. The offset in the STAMP payload (e.g., for unauthenticated mode (value 16)) is also statically configured for MNA.TSF.

4.1.2. Node Capability for MNA Sub-Stack with Opcode MNA.TSF

The STAMP Session-Sender needs to know if the Session-Reflector can process the MNA Sub-Stack with Opcode MNA.TSF to avoid dropping the test packets. The signaling extension for this capability exchange or local configuration are outside the scope of this document.

4.2. Enhanced Performance Measurement Procedure for SRv6 Policies

An SRv6 Policy Candidate-Path may contain a number of Segment Lists. Each Segment List may contain a number of SRv6 SIDs as defined in [RFC8986]. A Session-Sender test packet MUST be transmitted using each Segment List of the SRv6 Policy. An SRv6 Policy may contain an SRv6 Segment Routing Header (SRH) carrying a Segment List as described in [RFC8754]. The content of an example Session-Sender test packet for an end-to-end SRv6 Policy using an SRH is shown in Figure 4.

The loopback measurement mode for SRv6 Policies defined in Section 4.2.3.4 of [I-D.ietf-spring-stamp-srpm] is used and enhanced as described below.

The [RFC8986] defines SRv6 Endpoint Behaviours for SRv6 nodes. A new Timestamp and Forward Endpoint Behaviour is defined for Segment Routing Header (SRH) [RFC8754] to enable "Timestamp and Forward (TSF)" function for the received test packets.

In the Session-Sender test packets for SRv6 Policies, Timestamp and Forward Endpoint Function (End.TSF) is carried with the target Segment Identifier (SID) in SRH [RFC8754] as shown in Figure 4, to collect "Receive Timestamp" field in the payload of the test packet. The Segment List for the reverse direction path can be added after the target SID to receive the return test packet on a specific path. When a Session-Reflector receives a packet with Timestamp and Forward Endpoint (End.TSF) for the target SID, which is local, after timestamping the packet at the specific offset, the Session-Reflector forwards the packet using the next SID in the SRH or inner IPv6 header in the packet (just like the data packets for the normal traffic).

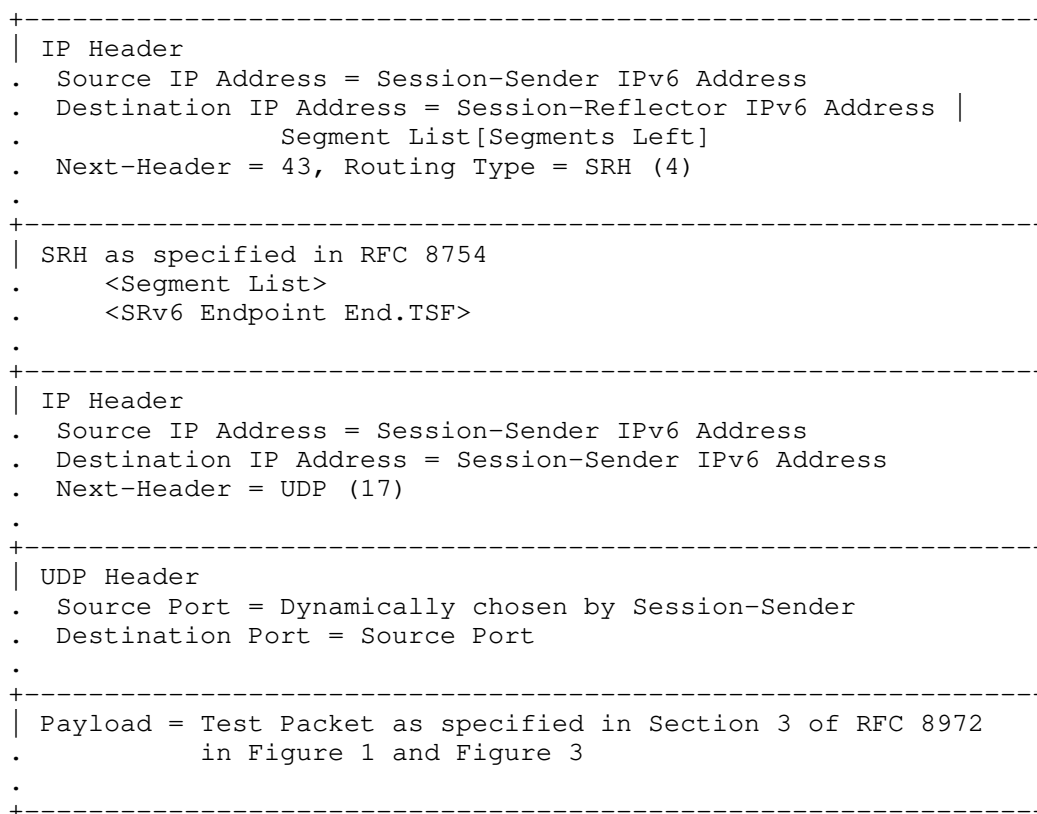


Figure 3: Example STAMP Test Packet with Endpoint Function for TSF for SRv6

4.2.1. Timestamp and Forward Endpoint Function Assignment

New SRv6 Endpoint Behavior is defined called "Endpoint Behavior bound to SID with Timestamp and Forward (End.TSF)". The End.TSF is a node SID instantiated at STAMP Session-Reflector node. The End.TSF is statically configured on the STAMP Session-Reflector node and not advertised into the routing protocols. The timestamp format for 64-bit PTPv2 and NTP to be added in the STAMP payload is statically configured for End.TSF. The offset in the STAMP payload (e.g., for unauthenticated mode (value 16)) is also statically configured for End.TSF.

4.2.2. Node Capability for Timestamp and Forward Endpoint Function

The STAMP Session-Sender needs to know if the Session-Reflector can process the Timestamp and Forward Endpoint Function to avoid dropping test packets. The signaling extension for this capability exchange or local configuration are outside the scope of this document.

5. Security Considerations

The procedures defined in this document is intended for deployment in a single operator network domain. As such, the Session-Sender address, Session-Reflector address, and IP and SR forward and return paths are provisioned by the operator for the STAMP session. It is assumed that the operator has verified the integrity of the IP and SR forward and return paths used to transmit STAMP test packets.

The Security Considerations specified in [RFC8762] and [RFC8972] also apply to the procedures defined in this document.

The Security Considerations specified in [I-D.ietf-spring-stamp-srpm] are also applicable to the procedures defined in this document.

The Security Considerations specified in [I-D.ietf-mppls-mna-hdr] are also applicable to the procedures defined in this document.

The Security Considerations specified in [RFC8986] are also applicable to the procedures defined in this document.

6. IANA Considerations

This document does not require any IANA action.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8762] Mirsky, G., Jun, G., Nydell, H., and R. Foote, "Simple Two-Way Active Measurement Protocol", RFC 8762, DOI 10.17487/RFC8762, March 2020, <<https://www.rfc-editor.org/info/rfc8762>>.
- [RFC8972] Mirsky, G., Min, X., Nydell, H., Foote, R., Masputra, A., and E. Ruffini, "Simple Two-Way Active Measurement Protocol Optional Extensions", RFC 8972, DOI 10.17487/RFC8972, January 2021, <<https://www.rfc-editor.org/info/rfc8972>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [I-D.ietf-spring-stamp-srpm]
Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and R. Foote, "Performance Measurement Using Simple TWAMP (STAMP) for Segment Routing Networks", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-09, 7 August 2023, <<https://www.ietf.org/archive/id/draft-ietf-spring-stamp-srpm-09.txt>>.
- [I-D.ietf-mpls-mna-hdr]
Rajamanickam, J., Ed., Gandhi, R., Ed., Zigler, R., Ed., Song, H., Ed., and K. Kompella, Ed., "MPLS Network Action Sub-Stack Solution", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-hdr-02, March 2023, <<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-hdr-02.txt>>.

7.2. Informative References

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

Acknowledgments

The authors would like to thank Greg Mirsky, Kireeti Kompella, and Adrian Farrel for providing useful comments.

Authors' Addresses

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Email: cfilsfil@cisco.com

Navin Vaghamshi
Reliance
Email: Navin.Vaghamshi@ril.com

Moses Nagarajah
Telstra
Email: Moses.Nagarajah@team.telstra.com

Richard Foote
Nokia
Email: footer.foote@nokia.com

Mach(Guoyi) Chen
Huawei
Email: mach.chen@huawei.com

Amit Dhamija
Rakuten
Email: amit.dhamija@rakuten.com