IP Performance Measurement
Internet-Draft
Updates: 8972 (if approved)
Intended status: Standards Track
Expires: 3 July 2026

G. White
CableLabs
G. Mirsky
Ericsson
X. Min
ZTE Corp.

30 December 2025

Update of the Simple Two-way Active Measurement Protocol Class-of-Service Extension - ECN
draft-ietf-ippm-stamp-cos-ecn-00

Abstract

The Simple Two-Way Active Measurement Protocol (STAMP) enables one-way and round-trip measurement of network metrics between IP hosts, and has a facility for defining optional extensions. This document updates the definition of the Class of Service TLV (originally defined in RFC 8972) to enable the measurement of manipulation of the value of the Explicit Congestion Notification (ECN) field of the IP header by middleboxes between two STAMP hosts, and to enable discovery and measurement of paths that provide differential treatment of packets depending on the value of their ECN field.

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 3 July 2026.

Copyright Notice

Table of Contents

1. Introduction

[RFC8972] defined several extensions to the Simple Two-way Active Measurement Protocol (STAMP). Among those is a Class of Service (CoS) TLV that enables measurements that utilize the Differentiated Services Code Point (DSCP) marking in both directions. Also, the CoS TLV supports outbound measurements that utilize the Explicit Congestion Notification (ECN) field, but it lacked support for such measurements on the return path. Experience deploying STAMP and its extensions demonstrated that it is helpful to an operator to monitor ECN's consistency in both directions. This specification updates the definition of the CoS TLV in a backward compatible manner to support monitoring of ECN in the return path of the STAMP test session.

Another STAMP extension header [I-D.ietf-ippm-stamp-ext-hdr] defines a mechanism that can be used by a Session-Sender to request that the Session-Reflector copy the entire IP header from the received test packet into the reflected test packet. That capability could be used in place of the CoS TLV to observe the DSCP and ECN values in the forward direction in situations where there is no need to control the return path DSCP or ECN value.

2. Conventions Used in This Document

2.1. Acronyms

CoS: Class of Service

DSCP: Differentiated Services Code Point

ECN: Explicit Congestion Notification

STAMP: Simple Two-way Active Measurement Protocol

2.2. 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.

3. Class of Service TLV

3.1. TLV Definition

The STAMP Session-Sender MAY include a CoS TLV in the STAMP test packet. The format of the CoS TLV is presented in Figure 1.

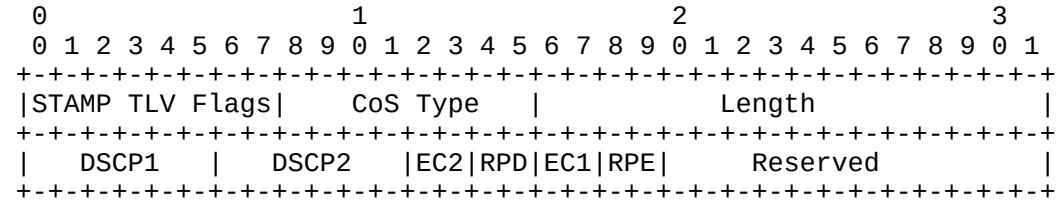

Figure 1: Class of Service TLV

Where the fields are defined as follows.

* STAMP TLV Flags: eight-bit field; format presented in Section 4 of [RFC8972].

* CoS (Class of Service) Type: one-octet field; the value MUST be set to 4.

* Length: two-octet field; set equal to the value 4 (octets).

* DSCP1: DSCP value intended by the Session-Sender to be used as the DSCP value of the reflected test packet.

* DSCP2: received value in the DSCP field at the ingress of the Session-Reflector.

* EC2: received value in the ECN field at the ingress of the Session-Reflector.

* RPD (reverse path DSCP): two-bit field indicating whether the Session-Reflector used DSCP1 or DSCP2 as the DSCP value of the reflected test packet; a Session-Sender MUST set the value of the RPD field to 0b00 on transmission.

* EC1: ECN value intended by the Session-Sender to be used as the ECN value of the reflected test packet.

* RPE (reverse path ECN): two-bit field indicating whether the Session-Reflector used EC1 as the ECN value of the reflected test packet; a Session-Sender MUST set the value of the RPE field to 0b00 on transmission.

* Reserved: twelve-bit field; MUST be zeroed on transmission and ignored on receipt.

3.2. Session-Reflector Behavior

A STAMP Session-Reflector that receives a test packet with the CoS TLV MUST include the CoS TLV in the reflected test packet. The Session-Reflector MUST copy the value of the DSCP and ECN fields of the IP header of the received STAMP test packet into the DSCP2 and EC2 fields, respectively, in the CoS TLV in the reflected test packet.

A Session-Reflector might have a local policy configuration that limits the DSCP values that it can use for reflected test packets. This local policy could be (for example) a system default policy, a global policy for the Session-Reflector, or a policy that is configured for specific destination addresses or networks. The Session-Reflector MUST verify whether the use of the value of the

DSCP1 field is permitted in the reflected test packet. If it is, the Session-Reflector MUST set the DSCP field's value in the IP header of the reflected test packet equal to the value of the DSCP1 field of the received test packet. Otherwise, the Session-Reflector MUST use the DSCP value of the received STAMP packet and set the value of the RPD field to 0b01. Upon receiving the reflected packet, if the value of the RPD field is 0b00, the Session-Sender will save the DSCP value for analysis of the CoS in the reverse direction. If the value of the RPD field in the received reflected packet is 0b01, only CoS in the forward direction can be analyzed.

A Session-Reflector might have a local policy configuration that limits the ECN values that it can use for reflected test packets. This local policy could be (for example) a system default policy, a global policy for the Session-Reflector, or a policy that is configured for specific destination addresses or networks. Additionally, a Session-Reflector could have platform limitations that prevent it from setting the ECN value in reflected test packets. The Session-Reflector MUST set the ECN value in the IP header of the reflected STAMP test packet to the value of the EC1 field, if it is permitted and capable to do so. If the Session-Reflector is able to set the ECN value in the IP header of the reflected STAMP test packet to the EC1 value, it MUST then set the RPE field in the CoS TLV in the reflected test packet to the value 0b11. If the Session-Reflector is unable to set the ECN value in the IP header of the reflected STAMP test packet to the EC1 value, it MUST instead set the RPE field in the CoS TLV in the reflected test packet to the value 0b10. As a result, the Session-Sender can detect whether the EC1 value was used by inspecting the value of the RPE field in the received reflected test packet.

3.3. Interoperability with RFC 8972

The extended CoS TLV defined in this draft is backward compatible with the specification in Section 4.4 of [RFC8972]. The handling of the DSCP1, DSCP2, EC2 and RPD fields defined here is identical to the handling defined for the equivalent fields in Section 4.4 of [RFC8972].

Consider a case when an implementation that supports this specification performs as Session-Sender, and the intended Session-Reflector's support of the CoS TLV is according to Section 4.4 of [RFC8972]. If the operator requires monitoring ECN in the reverse direction, the value of the EC1 field will be set to a non-zero value. Because the Session-Reflector would treat the EC1 field as part of the Reserved field, it would ignore its value as per Section 4.4 of [RFC8972]. Further, the Session-Reflector would treat the RPE field as part of the Reserved field and thus it would send

the value 0b00 in the reflected STAMP packet, rather than sending the value 0b10 or 0b11. Consequently, the Session-Sender will determine that the Session-Reflector does not implement the current version of the CoS TLV, and that the ECN value in the IP header of the reflected test packet was not set as requested.

Also, this specification supports the case when the Session-Reflector supports the extended CoS TLV as defined in this specification and the Session-Sender supports the CoS TLV according to Section 4.4 of [RFC8972]. In that scenario, the Session-Sender will set its [RFC8972] Reserved field to zeros. The Session-Reflector will interpret the first two bits of that field as the EC1 field, as shown in Figure 1 and thus will set the value of the ECN field in the IP header of the reflected packet to 0b00. Further the Session-Reflector will set the RPE field to 0b11. The Session-Sender will treat the RPE field as part of the Reserved field and will ignore its value.

3.4. Congestion Response

[RFC3168] and [RFC9330] mandate that applications that send packets marked with the ECT0 or ECT1 codepoints implement a response to congestion notifications from the network. While the STAMP protocol doesn’t include the concept of a connection between a Session-Sender and a Session-Reflector, there are situations in which the Session-Sender may send multiple STAMP packets to the same Session-Reflector within the round-trip time between them, and there are situations in which the Session-Sender may elicit multiple STAMP packets from a Session-Reflector in response to a single sent STAMP packet. As such, the Session-Sender is generally responsible for dictating both its sending rate, and the sending rate of the packets sent in response by an individual Session-Reflector, and so is not exempt from this mandate.

When using this CoS TLV with either an ECT0 or an ECT1 value in the ECN field of the IP header of the STAMP packet, the Session-Sender gains the ability (via observing a CE reflected value in the EC2 field) to detect that congestion was present on the path from the Session-Sender to the Session-Reflector. Similarly, when using this CoS TLV with either an ECT0 or an ECT1 value in the EC1 field, the Session-Sender gains the ability (via observing a CE in the ECN field of the IP header of the reflected packet) to detect that congestion was present on the path from the Session-Reflector to the Session-Sender.

If a Session-Sender sends multiple STAMP packets with the CoS TLV to a Session-Reflector within an RTT with either the ECT0 or ECT1 value in the ECN field of the IP header, it MUST observe the reflected EC2

field and reduce its sending rate upon observation of a CE value. If a Session-Sender sends multiple STAMP packets with the CoS TLV to a Session-Reflector within an RTT with either the ECT0 or ECT1 value in the EC1 field, it MUST observe the ECN value in the IP header of the reflected packets and reduce its sending rate upon observation of a CE value. If a Session-Sender sends a STAMP packet containing both the CoS TLV and the Reflected Test Packet Control TLV ([I-D.ietf-ippm-asymmetrical-pkts]) to a Session-Reflector, and specifies either the ECT0 or ECT1 value in the EC1 field, it MUST observe the ECN value in the IP header of the reflected packets and adjust the Reflected Test Packet Control parameters in any future STAMP packet sent to the same Session-Reflector based on the observation of CE values.

4. IANA Considerations

This document includes no request to IANA.

5. Security Considerations

This document extends the functionality of the CoS TLV ([RFC8972]) and inherits all the security considerations applicable to the base STAMP specification [RFC8762] and [RFC8972].

As this specification defines a mechanism to set ECN values, this document inherits all the security considerations discussed in [RFC3168] and [RFC9330]. Monitoring and optional control of ECN for a reflected STAMP test packet using the extended CoS TLV may be used across the Internet such that the Session-Sender and the Session-Reflector are located in different domains.

6. References

6.1. 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>.

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <https://www.rfc-editor.org/info/rfc2474>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.

[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>.

[RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <https://www.rfc-editor.org/info/rfc9330>.

[I-D.ietf-ippm-asymmetrical-pkts] Mirsky, G., Ruffini, E., Nydell, H., Foote, R. F., and W. Hawkins, "Performance Measurement with Asymmetrical Traffic Using STAMP", Work in Progress, Internet-Draft, draft-ietf-ippm-asymmetrical-pkts-08, 28 June 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-asymmetrical-pkts-08>.

6.2. Informative References

[RFC8311] Black, D., "Relaxing Restrictions on Explicit Congestion Notification (ECN) Experimentation", RFC 8311, DOI 10.17487/RFC8311, January 2018, <https://www.rfc-editor.org/info/rfc8311>.

[RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <https://www.rfc-editor.org/info/rfc9331>.

[I-D.ietf-ippm-stamp-ext-hdr]
Gandhi, R., Zhou, T., Li, Z., and W. Hawkins, "Simple Two-Way Active Measurement Protocol (STAMP) Extensions for Reflecting STAMP Packet IP Headers", Work in Progress, Internet-Draft, draft-ietf-ippm-stamp-ext-hdr-07, 6 November 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-stamp-ext-hdr-07>.

Acknowledgements

TBD

Contributors

Karthik Sundaresan, William Hawkins III, Gorry Fairhurst, R端diger Geib

Authors' Addresses

Greg White
CableLabs
858 Coal Creek Circle
Louisville, Colorado 80027
United States of America
Email: g.white@cablelabs.com
URI: http://www.cablelabs.com

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com

Xiao Min
ZTE Corp.
Nanjing
China
Email: xiao.min2@zte.com.cn