

Workgroup: NV03 Working Group  
Internet-Draft: draft-xiao-nv03-pm-geneve-09  
Published: 15 April 2024  
Intended Status: Standards Track  
Expires: 17 October 2024  
Authors: X. Min     G. Mirsky     S. Pallagatti  
          ZTE Corp.     Ericsson     VMware

## **Performance Measurement for Geneve**

### **Abstract**

This document describes the method to achieve Performance Measurement (PM) in point-to-point Generic Network Virtualization Encapsulation (Geneve) unicast tunnels used to make up an overlay network.

### **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 17 October 2024.

### **Copyright Notice**

Copyright (c) 2024 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](#)
- [2. Conventions Used in This Document](#)
  - [2.1. Abbreviations](#)
  - [2.2. Requirements Language](#)
- [3. PM Packet Transmission over Geneve Tunnel](#)
  - [3.1. PM Encapsulation With Inner Ethernet/IP/UDP Header](#)
  - [3.2. PM Encapsulation With Inner IP/UDP Header](#)
- [4. Demultiplexing of the PM packet](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Acknowledgements](#)
- [8. References](#)
  - [8.1. Normative References](#)
  - [8.2. Informative References](#)
- [Authors' Addresses](#)

### 1. Introduction

"Generic Network Virtualization Encapsulation" (Geneve) [[RFC8926](#)] provides an encapsulation scheme that allows building an overlay network of tunnels by decoupling the address space of the attached virtual hosts from that of the network.

This document describes the use of "Simple Two-way Active Measurement Protocol" (STAMP) [[RFC8762](#)] and "Simple Two-Way Active Measurement Protocol Optional Extensions" (STAMP Optional Extensions) [[RFC8972](#)], to enable measuring the performance of the path between two Geneve tunnel endpoints, like delay, delay variation, and packet loss.

Analogous to [[RFC9521](#)], in this document, Network Virtualization Edge (NVE) represents the Geneve tunnel endpoint, Tenant System (TS) represents the physical or virtual device attached to a Geneve tunnel endpoint from the outside, Virtual Access Point (VAP) represents the NVE side of the interface between the NVE and the TS, the usage of Management Virtual Network Identifier (VNI) is described in [[I-D.ietf-nvo3-geneve-oam](#)] and outside the scope of this document.

### 2. Conventions Used in This Document

#### 2.1. Abbreviations

FCS: Frame Check Sequence

Geneve: Generic Network Virtualization Encapsulation

NVE: Network Virtualization Edge

PM: Performance Measurement

SSID: STAMP Session Identifier

STAMP: Simple Two-way Active Measurement Protocol

TS: Tenant System

VAP: Virtual Access Point

VNI: Virtual Network Identifier

## 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. PM Packet Transmission over Geneve Tunnel

The PM session is originated and terminated at the VAP of an NVE. The selection of the PM packet encapsulation is based on how the VAP encapsulates the data packets. This document defines two formats of PM packet encapsulation in Geneve:

\*with Ethernet and IP/UDP encapsulation;

\*with IP/UDP encapsulation.

### 3.1. PM Encapsulation With Inner Ethernet/IP/UDP Header

If the VAP that originates the PM packets is used to encapsulate Ethernet data frames, then the PM packets are encapsulated in Geneve as described below. Note that here the PM packets are STAMP test packets.




Figure 1: Geneve Encapsulation of PM Packet With the Inner Ethernet/IP/UDP Header

The STAMP test packet MUST be carried inside the inner Ethernet frame of the Geneve packet. The inner Ethernet frame carrying the STAMP test packet has the following format:

The inner Ethernet header and IP header are encoded as defined in Section 4 of [[RFC9521](#)].

The destination UDP port MUST follow the STAMP UDP port usage defined in Section 4.1 of [[RFC8762](#)].

The STAMP test packet can be STAMP Session-Sender test packet or STAMP Session-Reflector test packet. The STAMP test packet is encoded as specified in [[RFC8762](#)] and [[RFC8972](#)].

When the PM packets are encapsulated in Geneve in this way, the values in the Geneve header MUST be set as specified in Section 4 of [[RFC9521](#)].

### **3.2. PM Encapsulation With Inner IP/UDP Header**

If the VAP that originates the PM packets is used to encapsulate IP data packets, then the PM packets are encapsulated in Geneve as described below. Note that here the PM packets are STAMP test packets.




Figure 2: Geneve Encapsulation of PM Packet With the Inner IP/UDP Header

The STAMP test packet MUST be carried inside the inner IP packet of the Geneve packet. The inner IP packet carrying the STAMP test packet has the following format:

The inner IP header is encoded as defined in Section 5 of [[RFC9521](#)].

The destination UDP port MUST follow the STAMP UDP port usage defined in Section 4.1 of [[RFC8762](#)].

The STAMP test packet can be STAMP Session-Sender test packet or STAMP Session-Reflector test packet. The STAMP test packet is encoded as specified in [[RFC8762](#)] and [[RFC8972](#)].

When the PM packets are encapsulated in Geneve in this way, the values in the Geneve header MUST be set as specified in Section 5 of [[RFC9521](#)].

#### 4. Demultiplexing of the PM packet

Once a packet is received, the NVE validates the packet as specified in Section 4.1 and 5.1 of [[RFC9521](#)], except that the received STAMP test packet would be processed by STAMP Session-Sender or STAMP Session-Reflector, instead of BFD.

Analogous to BFD over Geneve, multiple STAMP sessions for the same VNI may be running between two NVEs, so there needs to be a mechanism for demultiplexing received STAMP test packets to the proper session.

If the STAMP test packet is received with STAMP Session Identifier (SSID) equals to 0, then the procedure for demultiplexing the received STAMP test packets MUST follow the procedure for demultiplexing the received BFD packets with Your Discriminator equals to 0, which is specified in Section 4.1 and 5.1 of [[RFC9521](#)].

If the STAMP test packet is received with a non-zero SSID, then the STAMP session MUST be demultiplexed only with SSID as the key.

#### 5. Security Considerations

Security issues discussed in [[RFC8762](#)], [[RFC8972](#)], [[RFC8926](#)] and [[RFC9521](#)] apply to this document.

#### 6. IANA Considerations

This document has no IANA action requested.

#### 7. Acknowledgements

TBA.

#### 8. References

##### 8.1. 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>>.

[RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.

[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>>.

[RFC9521] Min, X., Mirsky, G., Pallagatti, S., Tantsura, J., and S. Aldrin, "Bidirectional Forwarding Detection (BFD) for Generic Network Virtualization Encapsulation (Geneve)", RFC 9521, DOI 10.17487/RFC9521, January 2024, <<https://www.rfc-editor.org/info/rfc9521>>.

## 8.2. Informative References

[I-D.ietf-nvo3-geneve-oam] Mirsky, G., Boutros, S., Black, D. L., and S. Pallagatti, "OAM for use in GENEVE", Work in Progress, Internet-Draft, draft-ietf-nvo3-geneve-oam-09, 6 December 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-nvo3-geneve-oam-09>>.

## Authors' Addresses

Xiao Min  
ZTE Corp.  
Nanjing  
China

Phone: [+86 18061680168](tel:+8618061680168)  
Email: [xiao.min2@zte.com.cn](mailto:xiao.min2@zte.com.cn)

Greg Mirsky  
Ericsson  
United States of America

Email: [gregimirsky@gmail.com](mailto:gregimirsky@gmail.com)

Santosh Pallagatti

VMware

India

Email: [santosh.pallagatti@gmail.com](mailto:santosh.pallagatti@gmail.com)