Stream: Internet Engineering Task Force (IETF)
RFC: 9772
Category: Standards Track
Published: April 2025
ISSN: 2070-1721
Authors: G. Mirsky, *Ericsson*; S. Boutros, *Ciena*; D. Black, *Dell EMC*; S. Pallagatti, *VMware*

RFC 9772
Active Operations, Administration, and Maintenance (OAM) for Use in Generic Network Virtualization Encapsulation (Geneve)

Abstract

Geneve (Generic Network Virtualization Encapsulation) is a flexible and extensible network virtualization overlay protocol designed to encapsulate network packets for transport across underlying physical networks. This document specifies the requirements and provides a framework for Operations, Administration, and Maintenance (OAM) in Geneve networks. It outlines the OAM functions necessary to monitor, diagnose, and troubleshoot Geneve overlay networks to ensure proper operation and performance. The document aims to guide the implementation of OAM mechanisms within the Geneve protocol to support network operators in maintaining reliable and efficient virtualized network environments.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9772.

Copyright Notice

Copyright (c) 2025 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

Geneve [RFC8926] is designed to support various scenarios of network virtualization. It encapsulates multiple protocols, such as Ethernet and IPv4/IPv6, and includes metadata within the Geneve message.

Operations, Administration, and Maintenance (OAM) protocols provide fault management and performance monitoring functions necessary for comprehensive network operation. Active OAM protocols, as defined in [RFC7799], utilize specially constructed packets injected into the

network. OAM protocols such as ICMP and ICMPv6 ([RFC0792] and [RFC4443] respectively), Bidirectional Forwarding Detection (BFD) [RFC5880], and the Simple Two-way Active Measurement Protocol (STAMP) [RFC8762] are examples of active OAM protocols. To ensure that performance metrics or detected failures are accurately related to a particular Geneve flow, it is critical that these OAM test packets share fate, i.e., are in-band, with the overlay data packets of that monitored flow when traversing the underlay network. In this document, "in-band OAM" is interpreted as follows:

- In-band OAM is an active or hybrid OAM method, as defined in [RFC7799], that traverses the same set of links and interfaces and receives the same Quality of Service treatment as the monitored object. In this context, the monitored object refers to either the entire Geneve tunnel or a specific tenant flow within a given Geneve tunnel.

Section 2.1 lists the general requirements for active OAM protocols in the Geneve overlay network. IP encapsulation meets these requirements and is suitable for encapsulating active OAM protocols within a Geneve overlay network. Active OAM messages in a Geneve overlay network are exchanged between two Geneve tunnel endpoints, which may be the tenant-facing interface of the Network Virtualization Edge (NVE) or another device acting as a Geneve tunnel endpoint. Testing end-to-end between tenants is out of scope. For simplicity, this document uses an NVE to represent the Geneve tunnel endpoint. Refer to [RFC7365] and [RFC8014] for detailed definitions and descriptions of an NVE.

The IP encapsulation of Geneve OAM defined in this document applies to an overlay service by introducing a Management Virtual Network Identifier (VNI), which can be used in combination with various values of the Protocol Type field in the Geneve header, such as Ethertypes for IPv4 or IPv6. The analysis and definition of other types of OAM encapsulation in Geneve are outside the scope of this document.

1.1. Conventions Used in This Document

1.1.1. 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.

1.1.2. Acronyms

Geneve: Generic Network Virtualization Encapsulation

NVO3: Network Virtualization over Layer 3

OAM: Operations, Administration, and Maintenance

VNI: Virtual Network Identifier

BFD: Bidirectional Forwarding Detection

STAMP: Simple Two-way Active Measurement Protocol

NVE: Network Virtualization Edge

2. The Applicability of Active OAM Protocols in Geneve Networks

2.1. Requirements for Active OAM Protocols in Geneve Networks

OAM protocols, whether part of fault management or performance monitoring, are intended to provide reliable information that can be used to detect a failure, identify the defect, and localize it, thus helping to identify and apply corrective actions to minimize the negative impact on service. Several OAM protocols are used to perform these functions; these protocols require demultiplexing at the receiving instance of Geneve. To improve the accuracy of the correlation between the condition experienced by the monitored Geneve tunnel and the state of the OAM protocol, the OAM encapsulation is required to comply with the following requirements:

Requirement 1: Geneve OAM test packets **MUST** share the same fate as the data traffic of the monitored Geneve tunnel. Specifically, the OAM test packets **MUST** be in-band with the monitored traffic and follow the same overlay and transport path as packets carrying data payloads in the forward direction, i.e., from the ingress toward the egress endpoint(s) of the OAM test.

An OAM protocol **MAY** be employed to monitor an entire Geneve tunnel. In this case, test packets could be in-band relative to a subset of tenant flows transported over the Geneve tunnel. If the goal is to monitor the conditions experienced by the flow of a particular tenant, the test packets **MUST** be in-band with that specific flow within the Geneve tunnel. Both scenarios are discussed in detail in Section 2.2.

Requirement 2: The encapsulation of OAM control messages and data packets in the underlay network **MUST** be indistinguishable from each other from the underlay network IP forwarding point of view.

Requirement 3: The presence of an OAM control message in a Geneve packet **MUST** be unambiguously identifiable to Geneve functionality, such as at endpoints of Geneve tunnels.

Requirement 4: OAM test packets **MUST NOT** be forwarded to a tenant system.

A test packet generated by an active OAM protocol, whether for defect detection or performance measurement, **MUST** be in-band with the tunnel or data flow being monitored, as specified in Requirement 1. In environments where multiple paths through the domain are available, underlay transport nodes can be programmed to use characteristic information to balance the load across known paths. It is essential that test packets follow the same route -- that is, traverse the same set of nodes and links as a data packet of the monitored flow. Therefore, the following requirement supports OAM packet fate-sharing with the data flow.

Requirement 5: There **MUST** be a way to encode entropy information into the underlay forwarding scheme so that OAM packets take the same data-flow paths as the transit traffic flows.

2.2. Defect Detection and Troubleshooting in Geneve Network with Active OAM

This section considers two scenarios where active OAM is used to detect and localize defects in a Geneve network. Figure 1 presents an example of a Geneve domain.

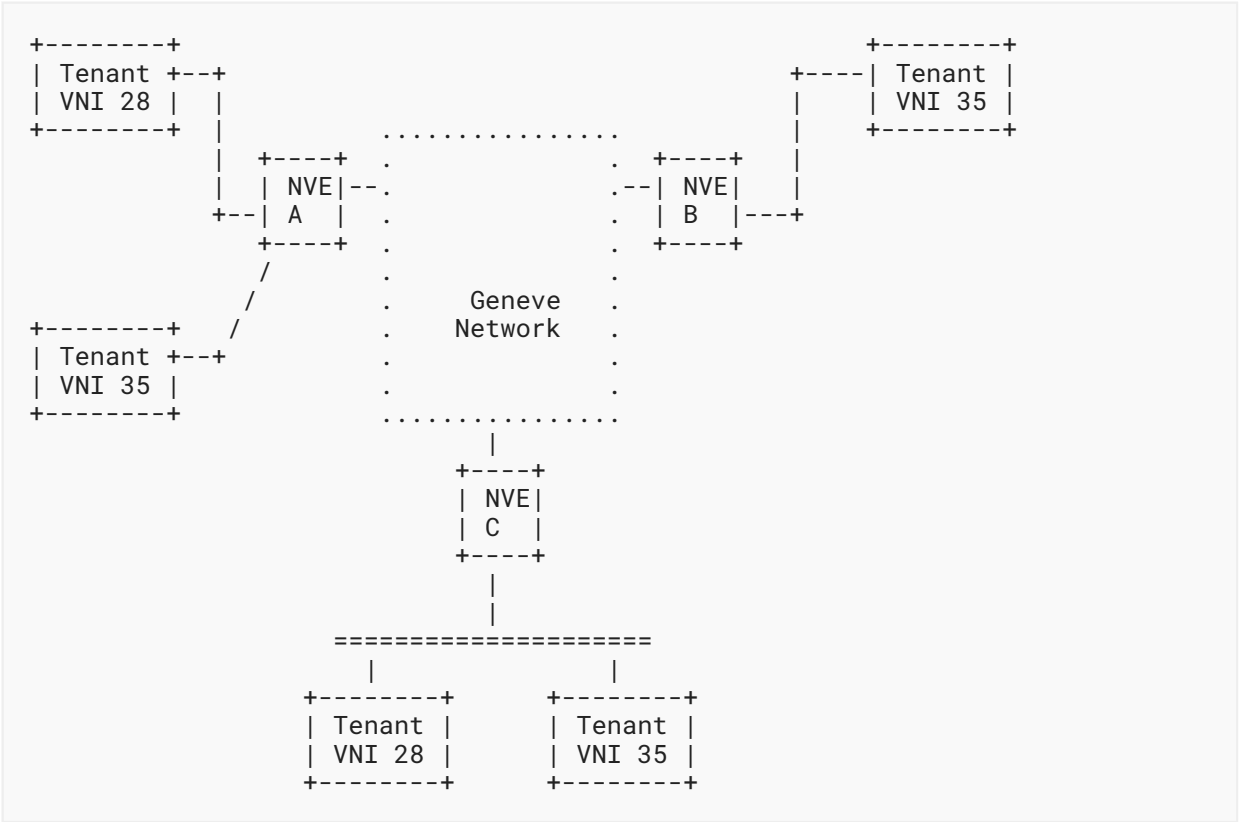

Figure 1: An Example of a Geneve Domain

In the first case, consider when a communication problem between Network Virtualization Edge (NVE) device A and NVE C exists. Upon investigation, the operator discovers that the forwarding in the IP underlay network is working accordingly. Still, the Geneve connection is unstable for all NVE A and NVE C tenants. Detection, troubleshooting, and localization of the problem can be done regardless of the VNI value.

In the second case, traffic on VNI 35 between NVE A and NVE B has no problems, as on VNI 28 between NVE A and NVE C. However, traffic on VNI 35 between NVE A and NVE C experiences problems, for example, excessive packet loss.

The first case can be detected and investigated using any VNI value, whether it connects tenant systems or not; however, to conform to Requirement 4, OAM test packets **SHOULD** be transmitted on a VNI that doesn't have any tenants. Such a Geneve tunnel is dedicated to carrying only control and management data between the tunnel endpoints, so it is referred to as a "Geneve control channel" and that VNI is referred to as the "Management VNI". A configured VNI **MAY** be used to identify the control channel, but it is **RECOMMENDED** that the default value 1 be used as the Management VNI. Encapsulation of test packets using the Management VNI is discussed in Section 2.3.

The control channel of a Geneve tunnel **MUST NOT** carry tenant data. As no tenants are connected using the control channel, a system that supports this specification **MUST NOT** forward a packet received over the control channel to any tenant. A packet received by the system that supports this specification over the control channel **MUST** be forwarded if and only if it is sent onto the control channel of the concatenated Geneve tunnel. Else, it **MUST** be terminated locally. The Management VNI **SHOULD** be terminated on the tenant-facing side of the Geneve encapsulation/decapsulation functionality, not the DC-network-facing side (per definitions in Section 4 of [RFC8014]), so that Geneve encap/decap functionality is included in its scope. This approach causes an active OAM packet, e.g., an ICMP echo request, to be decapsulated in the same fashion as any other received Geneve packet. In this example, the resulting ICMP packet is handed to NVE's local management functionality for the processing which generates an ICMP echo reply. The ICMP echo reply is encapsulated in Geneve (as specified in Section 2.3) for forwarding it back to the NVE that sent the echo request. One advantage of this approach is that a repeated ICMP echo request/reply test could detect an intermittent problem in Geneve encap/decap hardware, which would not be tested if the Management VNI were handled as a "special case" at the DC-network-facing interface.

The second case is when a test packet is transmitted using the VNI value associated with the monitored service flow. By doing that, the test packet experiences network treatment as the tenant's packets. An example of the realization of that scenario is discussed in [RFC9521].

2.2.1. Echo Request and Echo Reply in Geneve Tunnel

ICMP and ICMPv6 ([RFC0792] and [RFC4443] respectively), as noted above, are examples of an active OAM protocol. They provide required on-demand defect detection and failure localization. ICMP control messages immediately follow the inner IP header encapsulated in Geneve. ICMP extensions for Geneve networks use mechanisms defined in [RFC4884].

2.3. Active OAM Encapsulation in Geneve

Active OAM over a Management VNI in the Geneve network uses an IP encapsulation. Protocols such as BFD [RFC5880] and STAMP [RFC8762] use UDP transport. The destination UDP port number in the inner UDP header (Figure 2) identifies the OAM protocol. This approach is well-known and has been used, for example, in MPLS networks [RFC8029]. To use IP encapsulation for an active OAM protocol, the Protocol Type field of the Geneve header **MUST** be set to 0x0800 (IPv4) or 0x86DD (IPv6). [RFC9521] describes the use of IP encapsulation for BFD.

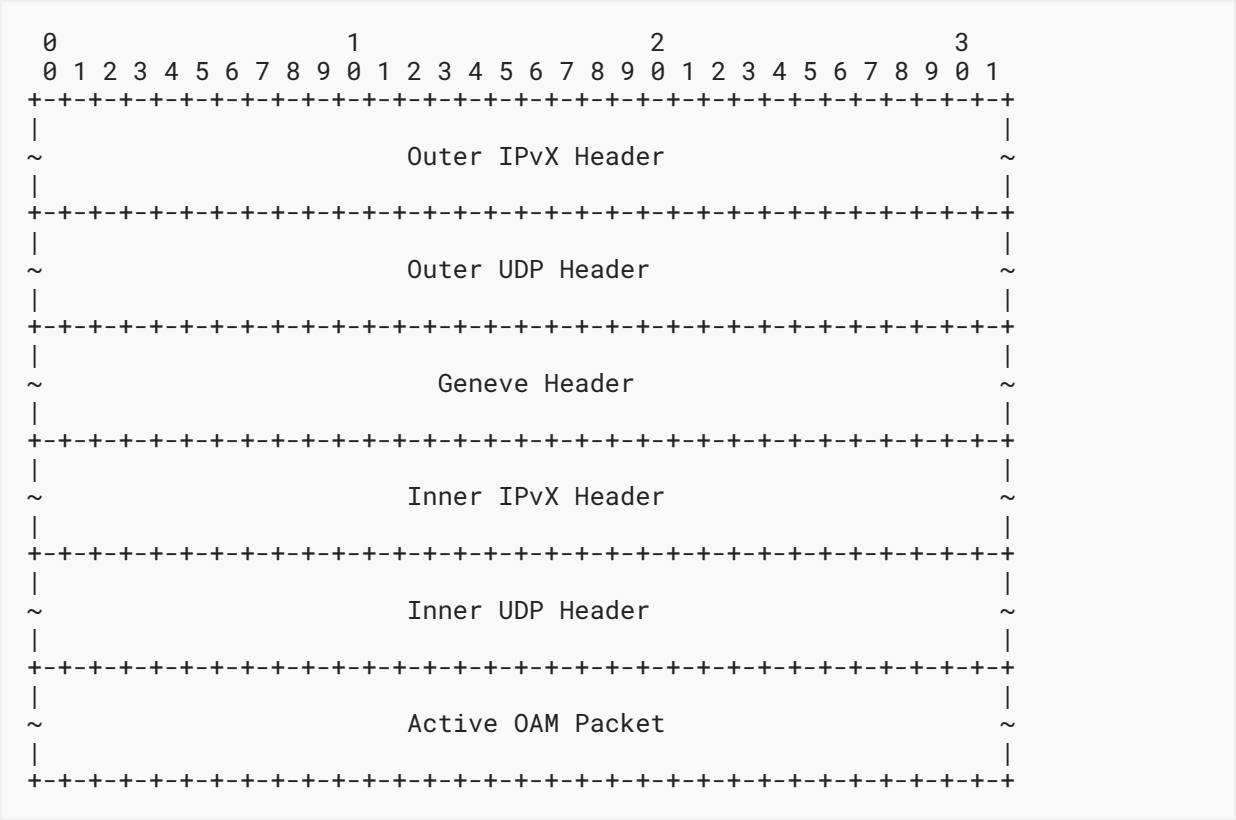

Figure 2: An Example of Geneve IP/UDP Encapsulation of an Active OAM Packet

Inner IP header:

Destination IP: The IP address **MUST** be set to the loopback address 127.0.0.1/32 for IPv4 version. For IPv6, the address **MUST** be selected from the Dummy IPv6 Prefix for IPv6 *Dummy-IPv6-Prefix*. A source-only IPv6 dummy address is used as the destination to generate an exception and a reply message to the request message received.

[Note to RFC Editor: Please replace *Dummy-IPv6-Prefix* with the actual value allocated (requested in draft-ietf-mpls-p2mp-bfd) in IANA IPv6 Special-Purpose Address Registry.]

Source IP: IP address of the NVE.

TTL or Hop Limit: **MUST** be set to 255 per [RFC5082]. The receiver of an active OAM Geneve packet with IP/UDP encapsulation **MUST** drop packets whose TTL/Hop Limit is not 255.

3. IANA Considerations

This document has no IANA actions.

4. Security Considerations

As part of a Geneve network, active OAM inherits the security considerations discussed in [RFC8926]. Additionally, a system **MUST** provide control to limit the rate of Geneve OAM packets punted to the Geneve control plane for processing in order to avoid overloading that control plane.

OAM in Geneve packets uses the General TTL Security Mechanism [RFC5082], and any packet received with an inner TTL / Hop Count other than 255 **MUST** be discarded.

5. References

5.1. 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>.

[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>.

5.2. Informative References

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, DOI 10.17487/RFC4884, April 2007, <https://www.rfc-editor.org/info/rfc4884>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <https://www.rfc-editor.org/info/rfc5082>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", RFC 7365, DOI 10.17487/RFC7365, October 2014, <https://www.rfc-editor.org/info/rfc7365>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <https://www.rfc-editor.org/info/rfc7799>.

[RFC8014] Black, D., Hudson, J., Kreeger, L., Lasserre, M., and T. Narten, "An Architecture for Data-Center Network Virtualization over Layer 3 (NVO3)", RFC 8014, DOI 10.17487/RFC8014, December 2016, <https://www.rfc-editor.org/info/rfc8014>.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <https://www.rfc-editor.org/info/rfc8029>.

[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>.

[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>.

Acknowledgments

The authors express their appreciation to Donald E. Eastlake 3rd for his suggestions that improved the readability of the document.

Authors' Addresses

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com

Sami Boutros
Ciena
Email: sboutros@ciena.com

David Black
Dell EMC
176 South Street
Hopkinton, MA, 01748
United States of America
Email: david.black@dell.com

Santosh Pallagatti
VMware
Email: santosh.pallagatti@gmail.com