NVO3
Internet-Draft
Intended status: Standards Track
Expires: December 29, 2017

D. Migault
Ericsson

June 27, 2017

Geneve Protocol Security Requirements
draft-mglt-nvo3-geneve-security-requirements-00

Abstract

This draft lists the security requirements associated to the Generic Network Virtualization Encapsulation (Geneve) [I-D.ietf-nvo3-geneve].

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 http://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 December 29, 2017.

Copyright Notice

Copyright (c) 2017 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 (http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Introduction

A cloud provider may administrate Tenant Systems belonging to one or multiple tenants using an Geneve overlay network. The Geneve overlay enables multiple Virtual Networks to coexist over a shared infrastructure, and a Virtual Network may be distributed within a single data center or between different data centers. The Geneve overlay network is constituted by Geneve forwarding elements as well as Network Virtualization Edges (NVE) [RFC8014]. Traffic with a Virtual Network is thus steered between NVEs using Generic Network Virtualization Encapsulation (Geneve) [I-D.ietf-nvo3-geneve].

This document analyses and lists the security requirements to securely deploy Geneve overlay networks. It is expected that these requirements will help design the appropriated security mechanisms for Geneve as well as provides some basis security notion for further Geneve deployments.

In addition, when a tenant subscribes to a cloud provider for hosting its Tenant Systems, the cloud provider manages the Geneve overlay network on behalf of the tenant [RFC7365]. It may also, but not necessarily manage the infrastructure supporting the overlay network. The Geneve security requirements listed in this document aims at

providing the cloud provider the necessary options to ensure the tenant:

1. Tenants Isolation, that is Tenant System inside a Virtual Network are isolated from other Tenants Virtual Networks. This Tenants isolation mostly prevents traffic from one tenant to be redirected to another tenant as well as traffic from one tenant being injected into another tenant.

2. Geneve robustness, that is a rogue elements of the Geneve overlay network will have limited impacts over the Geneve overlay network itself as well as over Tenants Systems.

3. Geneve isolation of the infrastructure, that is information in transit is not subject to passive monitoring. Information in transit concerns both information associated to the Geneve overlay network as well as information exchanged by the Tenant' Systems. Hiding information of the overlay network to the infrastructure is typically required when the overlay network provider and the infrastructure provider are different entities.

3. Tenants Isolation

Tenant Systems isolation prevents communications from one tenant to leak into another Tenant Systems' virtual network. This section is focused on an Geneve overlay network perspective which means:

1. Only communications between NVEs are considered. In other words, the transmission from the NVE to the Tenant System itself is out of the scope of this section. Similarly the security used by the infrastructure to steer Geneve Packets from a NVE to Geneve forwarding element is out of scope either.

2. Isolation is only broken by rogue NVE or rogue Geneve forwarding elements. In other words, breaking isolation using other elements or other protocol layers are out of scope of this section

3. A Geneve NVE SHOULD be able to set different security policies to different flows. These flows MUST be characterized from the Geneve Header and Geneve Options as well as some inner traffic selectors. Typically an NVE SHOULD be able to selectively authenticate only the sections that are not authenticate by the Tenant System. If the Tenant Systems authenticates its communications with TLS, only the IP, transport (TCP / UDP) and TLS/DTLS section should be encrypted while only the IP header and ESP header is expected to be encrypted when tenants' communications are encrypted with ESP.

Suppose Tenant A and Tenant B are two distinct tenants and are expected to remain isolated by the Geneve overlay network. The attacks breaking the isolation considered in this section are the injection of traffic into one virtual network as well as the redirection of one tenant's traffic to a third party.

3.1. Traffic Injection

Traffic injection can target a specific element on the overlay network such as, for example, an NVE, a Geneve forwarding element or eventually specific Tenant System. On a overlay network perspective, the difference of targeting a Tenant's System requires valid MAC and IP addresses of the Tenant's System.

In order to provide integrity protection, Tenant's System may protect their communications using IPsec or TLS. Such protection protects the Tenants from receiving spoofed packets, as any injected packet is expected to be discarded by the destination Tenant's System. Such protection is independent from the Geneve overlay network and as such provides protection against any node outside the virtual network including the nodes of the Geneve overlay network to inject packets to a Tenant System. On the other such protection does not protect the virtual network from receiving illegitimate packets that may disrupt the Tenant's System performances.

When Tenant Systems are protected against spoofed packets, the Geneve overlay network may still prevent such spoofed Geneve Packet to be steered into the virtual network. In addition, when the Tenant's System have not enabled such protections, the overlay network should be able to provide a secure infrastructure for hosting these virtual networks and prevent a third party to inject traffic into the overlay. In this section the third party is a node on the infrastructure hosting the Geneve overlay network. In addition, this node could be any Geneve element except the legitimate NVEs (source or destination).

A Geneve overlay network is composed of multiple Geneve forwarding elements steering a Geneve Packet between the two NVEs. The Geneve Packet is forwarded according to the information carried in the Geneve Packet as well as routing tables associated to this information. For that reason, the information carried in the Geneve Header, including Geneve Option MUST be accessible by the intermediary nodes.

In order to prevent traffic injection in one virtual network, the destination Geneve NVE MUST be able to authenticate the incoming traffic sent by the source NVE. Note that this threat model assumes

that the third party injecting traffic does not inject traffic through the NVEs.

Authentication of the whole Geneve Packet may raise the cost of security unnecessarily. In fact it is expected that the Tenant Systems will also protect their end-to-end traffic, as a result, corruption of the Geneve Payload can be detected by the System Tenant. In addition, for the ease of processing, an authenticated Geneve Packet should not impact the processing of the intermediary nodes, unless they are able to check the authentication themselves. A key advantage of validating the authentication by intermediary nodes is that detection can occur earlier, however such requirement may require the use of asymmetric cryptography, which may be balanced by its low performance over symmetric cryptography. As a result the following requirements are associated with the authentication:

REQ1: A Geneve NVE MUST be able to authenticate the Geneve Header including the immutable Geneve Options.

REQ2: A Geneve NVE MUST be able to agreement that authentication includes or not the Geneve Payload, and if so it SHOULD also be able to indicate that only a portion of it is authenticated.

REQ3: A Geneve intermediary forwarding element MAY be able to validate the authentication before the packet reaches the Geneve destination tunnel end point.

REQ4: A Geneve intermediary forwarding element MUST be able to insert an authenticated Geneve Option into a authenticated Geneve Packet - protected by the source Geneve tunnel termination point.

REQ5: A Geneve intermediary forwarding element not supporting authentication MUST NOT be impacted by the authentication of the Geneve Packet and should be able to handle the Geneve Packet as an non-authenticated Geneve Packet.

REQ6: A Geneve NVE SHOULD be able to set different security policies to different flows. These flows MUST be characterized from the Geneve Header and Geneve Options as well as some inner traffic selectors. Typically an NVE SHOULD be able to selectively encrypt only the sections that are not encrypted by the Tenant System. If the Tenant Systems encrypts its communications with TLS, only the IP, transport (TCP / UDP) and TLS/DTLS section should be encrypted while only the IP header and ESP header is expected to be encrypted when tenants' communications are encrypted with ESP.

3.2. Traffic Redirection

A rogue element of the overlay Geneve network under the control of an attacker may leak and redirect the traffic from a virtual network to the attacker for passive monitoring, or for actively re-injecting a modified Geneve Packet into the overlay.

Avoiding leaking information is hard to enforced at a Geneve level. However, the Geneve overlay network and the Tenants Systems can lower the consequences of such leakage in case these occurs. The Tenant System may protect partly the data carries over the overlay network using end-to-end encryption such as IPsec/TLS. Doing so provides integrity protection as well as confidentiality for the Tenant's information. Such protection applies even if the source or destination NVE are corrupted.

The purpose of the Geneve overlay network is to limit the information it is aware of to leak. When Tenant Systems are enforcing confidentiality of the information in transit with IPsec or TLS for example, they are still some information revealed the MAC and IP headers of the inner packet may remain unprotected. IN this case, the Geneve overlay network should be able to maintain this information confidential. When Tenant's have not enforced such security the Geneve overlay network should be able to provide a secure infrastructure and prevent leakage of information outside the virtual network. In addition, the information carried by the Geneve Header may also reveal some information on the overlay network itself, its deployment as well as states from the Tenant System. In this the Geneve overlay network should also be able to protect such Geneve Options.

Note that when the overlay network is hosted on an architecture that belongs to another administrative domain, the administrator of the infrastructure is typically able to perform passive monitoring attacks.

In order to protect the Geneve communications between the Geneve tunnel terminating points here are the following requirements:

REQ7: A Geneve NVE MUST be able to agree that the Geneve Payload or portion of it is encrypted as well as as immutable Geneve Options not intended for the intermediary Geneve nodes.

REQ8: A Geneve intermediary forwarding element MUST be able to insert an encrypted Geneve Option into a authenticated Geneve Packet - protected by the source Geneve tunnel termination point.

REQ9: A Geneve intermediary forwarding element MUST be able to insert an encrypted Geneve Option into an encrypted Geneve Packet - protected by the source Geneve NVE.

REQ10: A Geneve intermediary forwarding element not supporting encryption MUST NOT be impacted by the authentication of the Geneve Packet and should be able to handle the Geneve Packet as an non-protected Geneve Packet.

Re-injection through a Geneve intermediary node is prevented by the authentication. On the other hand, if the re-injection is performed through one of the Geneve NVE, the protection provided by encryption as well as authentication does not apply. The authentication is intended to check integrity toward the data provided by the source Geneve NVE. If that point is corrupted, it is likely to inject corrupted traffic with integrity protection. On the other hand, if the destination Geneve NVE is expected to validate the data, as a result if traffic is injected through that node it is likely to bypass the integrity validation.

4. Overlay Network Robustness

While Tenant isolation prevents one Tenant to inject packets into another Tenants, it does not prevent a rogue or misconfigured node to replay a packet, to load a specific Tenant System with a modified Geneve payload or to abuse the Geneve overlay network.

1. A rogue Geneve overlay forwarding element on path of one Tenants traffic may replay a valid packet to load the network. This can typically be seen as a volumetric attack in order to disrupt the tenants domain, a specific Tenant System or the multi Tenant infrastructure itself. In some cases, especially when the tenants costs are evaluate on the necessary computing resources, such attacks may target an increase of the tenants costs.

2. When traffic between tenants is not protected, a rogue Geneve overlay element may forward a modified packet over a valid Geneve Header. The crafted packet may for example, include a specifically crafted application payload intended for a specific Tenant Systems application. Other examples includes a larger randomly craft payload intended to load one specific application.

3. The Geneve forwarding policies are engineered according to the various types of flows with their associated volumetry and requirements. For example, some OAM flows are expected to be associated with a higher priority then standard data plane flows. Similarly, the use of various Geneve Header parameters or options may introduce different treatments. Updating the Geneve header

may result in counter all optimizations used to setup a performant infrastructure and thus affect the tenants.

Note nodes that may address such attacks MUST be provided means to authenticate the Geneve Packet. More specifically,

In order to avoid the above mentioned attacks, the following requirements should be considered:

REQ11: Geneve Header SHOULD be bound to the forwarded payload. By reading the Geneve Header and the Payload, the Geneve forwarding element SHOULD be able to validate the Geneve Header corresponds to the Geneve payload. In case of mismatch the Geneve forwarding element is expected to discard the packet.

REQ12: Geneve forwarding element SHOULD be provided anti replay mechanisms. By reading the Geneve Header, the Geneve forwarding element is expected to detect a packet has been replayed or at least limit the replay windows. When a packet is detected as being replayed, the Geneve forwarding element is expected to discard this packet.

5. Infrastructure Isolation

The cloud provider managing the Geneve overlay network may be willing to isolate the communications between Tenant Systems as well as the organization of the Geneve overlay from the infrastructure. Such isolation may be performed by encrypting the data in transit within the Geneve overly network.

5.1. Tenants Communication

The main purpose for encrypting tenants communication inside the Geneve overlay network is to prevent that external parties such a infrastructure providers may access to the information exchanged between Tenant System exchanged via the Geneve overlay network. A typical example comes would be the infrastructure provider used by the Geneve overlay network.

In addition, encryption of the data in transit in the Geneve overlay network may also be one way to prevent the leakage of information when tenant isolation is broken. Encryption is not expected to enforce tenant isolation, but if information can hardly be used by another tenant it may limit the interest in breaking such isolation to still information as well as it might reduce the risks of leaking some confidential information.

The requirements correspond to the those protecting against the redirection or passive monitoring attacks in Section 3.2.

IPsec or TLS provides end-to-end encryption for NVE communications. However, as the Geneve Header would be encrypted, these mechanisms cannot be used are general mechanisms for the overlay network.

Encrypting Geneve payload by the NVE prevents disclosing the Geneve payload to third party in case of leakage. However, such service is provided by the cloud provider and the tenant has little control over it. In most cases, if the tenant is willing to enforce data confidentiality, it is recommended that it encrypts communications between Tenants systems using IPsec or TLS. By doing so, the cloud provider would not even have access to such information. While encryption is being performed by the tenant, a cloud provider may be willing to avoid re-encrypting that same content. Instead, the cloud provider may prefer to only encrypt the tenants informations that have not been encrypted by TLS or IPsec. Doing so is expected to reduce the necessary resource for encrypting.

The requirements correspond to the those protecting against the redirection or passive monitoring attacks in Section 3.2.

5.2. Overlay Network Architecture

In addition, to the information exchanged between Tenant Systems, the cloud provider may also avoid revealing the distribution of the Tenant Systems through the data center. In fact a passive attacker may observe the NVI in the Geneve header in order to derive the communication pattern between the Tenant Systems. Other parameters or options may reveal other kind of informations. One possibility is to encrypt the information, but other transformations may also apply.

The requirements correspond to the those protecting against the redirection or passive monitoring attacks in Section 3.2.

6. IANA Considerations

There are no IANA consideration for this document.

7. Security Considerations

The whole document is about security.

Limiting the coverage of the authentication / encryption provides some means for an attack to craft special packets.

8. Acknowledgment

9. References

9.1. Normative References

[I-D.ietf-nvo3-geneve]
Gross, J., Ganga, I., and T. Sridhar, "Geneve: Generic Network Virtualization Encapsulation", draft-ietf-nvo3-geneve-04 (work in progress), March 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

9.2. Informational References

[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, <http://www.rfc-editor.org/info/rfc7365>.

[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, <http://www.rfc-editor.org/info/rfc8014>.

Author's Address

Daniel Migault
Ericsson
8400 boulevard Decarie
Montreal, QC H4P 2N2
Canada

Email: daniel.migault@ericsson.com