NVO3
Internet-Draft
Intended status: Informational
Expires: July 29, 2018

D. Migault
Ericsson
S. Boutros
D. Wing
VMware
S. Krishnan
Kaloom

January 25, 2018

Geneve Protocol Security Requirements
draft-mglt-nvo3-geneve-security-requirements-02

Abstract

The document defines the security requirements to protect tenants overlay traffic against security threats from the NVO3 network components that are interconnected with tunnels implemented using Generic Network Virtualization Encapsulation (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 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 July 29, 2018.

Copyright Notice

Copyright (c) 2018 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 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", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

The network virtualization overlay over Layer 3 (NVO3) as depicted in Figure 1, allows an overlay cloud provider to provide a logical L2/L3 interconnect for the Tenant Systems TSes that belong to a specific tenant network. A packet received from a TS is encapsulated by the ingress Network Virtualization Edge (NVE). The encapsulated packet is then sent to the remote NVE through a tunnel. When reaching the egress NVE of the tunnel, the packet is decapsulated and forwarded to the target TS. The L2/L3 address mappings to the remote NVE(s) are distributed to the NVEs by a logically centralized Network Virtualization Authority (NVA) or using a distributed control plane such as Ethernet-VPN. In a datacenter, the NVO3 tunnels can be implemented using Generic Network Virtualization Encapsulation (Geneve) [I-D.ietf-nvo3-geneve]. Such Geneve tunnels establish NVE-to-NVE communications, may transit within the data center via Geneve

Transit Nodes (GTN). The Geneve tunnels overlay network enable multiple Virtual Networks to coexist over a shared underlay infrastructure, and a Virtual Network may span a single data center or multiple data centers.

The underlay infrastructure on which the multi-tenancy overlay networks are hosted, can be owned and provided by an underlay provider who may be different from the overlay cloud provider.

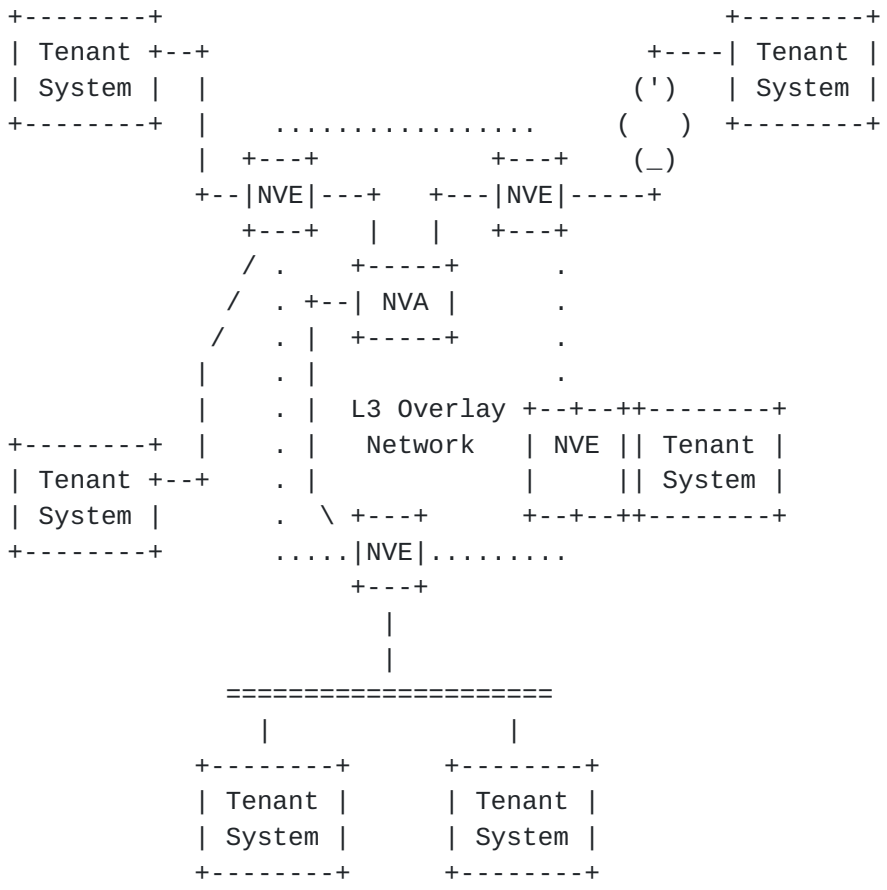

Figure 1: Generic Reference Model for Network Virtualization Overlays [RFC7365]

This document discusses the security risks that a Geneve based NVO3 network may encounter and tries to provide a list of essential security requirements that needs to be fulfilled. In addition, this document lists the requirements to protect the Geneve packet components defined in [I-D.ietf-nvo3-geneve] that include the Geneve tunnel IP and UDP header, the Geneve Header, Geneve options, and inner payload. Protecting the complete Geneve packet - that is the full IP packet or the full outer UDP payload for example - is out of scope of this document, given that this can be supported using existing mechanisms.

This document assumes that a tenant subscribes to an overlay cloud provider for hosting its Tenant Systems, the cloud provider manages the Geneve overlay network on behalf of the tenant. The overlay network will be hosted on an underlay network infrastructure, that may be managed by another underlay cloud provider.

The security requirements in this document aims at providing the overlay cloud provider the necessary options to ensure:

1. Delivering tenant payload traffic and ensuring privacy and integrity of the overlay traffic, and isolation between the overlay and underlay networks, as well preventing tenant traffic from being redirected or injected to other tenants.

2. Protecting tenant traffic from rogue devices in the providers of Geneve overlay or underlay networks.

In summary, the document defines the security requirements to protect tenants overlay traffic against security threats from the NVO3 network components that are interconnected with tunnels implemented using Geneve. As well, the document strongly recommend to re-use existing security protocols like IP Security (IPsec) [RFC4301] and Transport Layer Security (TLS) [RFC5246], and existing encryption algorithms, and authentication protocols.

Authentication requirements for NVO3 devices, automated key management, as well as packet level security providing confidentiality, integrity and authorization requirements defined in [I-D.ietf-nvo3-security-requirements] are also requirements for this document.

3. Terminology

This document uses the terminology of [RFC8014], [RFC7365] and [I-D.ietf-nvo3-geneve]

4. Security Threats

Attacks from compromised NVO3 and underlay network devices, and attacks from compromised tenant systems defined in [I-D.ietf-nvo3-security-requirements] are considered for the Geneve overlay network. Furthermore, the attackers knowing the details of the Geneve packets can perform their attacks by changing fields in the Geneve tunnel header, base header, Geneve options and Geneve packet inner payload.

Threats include traffic analysis, sniffing, injection, redirection, and replay. Based on these threats, this document enumerates the security requirements.

Threats are divided into two categories: passive attack and active attack.

4.1. Passive Attacks

Passive attacks include traffic analysis (noticing which workloads are communicating with which other workloads, how much traffic, and when those communications occur) and sniffing (examining traffic for useful information such as personally-identifyable information or protocol information (e.g., TLS certificate, overlay routing protocols).

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 [RFC7258].

Avoiding leaking information is hard to enforced and the security requirements expect to mitigate such attacks by lowering the consequences, typically making leaked data unusable to an attacker..

4.2. Active Attacks

Active attacks involve modifying packets, injecting packets, or interfering with packet delivery (such as by corrupting packet checksum).

There are multiple motivations to inject illegitimate traffic into a tenants network. When the rogue element is on the path of the TS traffic, it may be able to inject and receive the corresponding messages back. On the other hand, if the attacker is not on the path of the TS traffic it may be limited to only inject traffic to a TS without receiving any response back. When rogue element have access to the traffic in both directions, the possibilities are only limited by the capabilities of the other on path elements - GTN, NVE or TS - to detect and protect against the illegitimate traffic. On the other hand, when the rogue element is not on path, the surface for such attacks remains still quite large. For example, an attacker may target a specific TS or application by crafting a specific packet that can either generate load on the system or crash the system or application. TCP syn flood typically overload the TS while not requiring the ability to receive responses. Note that udp application are privileged target as they do not require the establishment of a session and are expected to treat any incoming packets.

Traffic injection may also be used to flood the virtual network to disrupt the communications between the TS or to introduce additional cost for the tenant, for example when pricing considers the traffic inside the virtual network. The two latest attacks may also take advantage of applications with a large factor of amplification for their responses as well as applications that upon receiving a packet interact with multiple TS. Similarly, applications running on top of UDP are privileged targets.

Note also that an attacker that is not able to receive the response traffic, may use other channels to evaluate or measure the impact of the attack. Typically, in the case of a service, the attacker may have access, for example, to a user interface that provides indication on the level of disruption and the success of an attack, Such feed backs may also be used by the attacker to discover or scan the network.

Preventing traffic to cross virtual networks, reduce the surface of attack, but rogue element main still perform attacks within a given virtual network by replaying a legitimate packet. Some variant of such attack also includes modification of unprotected parts when available in order for example to increase the payload size.

5. Requirements for Security Mitigations

The document assumes that Security protocols, algorithms, and implementations provide the security properties for which they are designed, an attack caused by a weakness in a cryptographic algorithm is out of scope.

Protecting network connecting TSes and NVEs which could be accessible to outside attackers is out of scope.

An attacker controlling an underlying network device may break the communication of the overlays by discarding or delaying the delivery of the packets passing through it. The security consideration to prevent this type of attack is out of scope of this document.

Securing communication between NVAs and NVEs is out of scope.

Selectively providing integrity/authentication, confidentiality/ encryption of only portions of the Geneve packet is in scope. This will be the case if the Tenant Systems uses security protocol to protect its communications.

5.1. Protection Against Traffic Sniffing

A passive network observer can determine two virtual machines are communicating by manipulating activity or network activity of other virtual machines on that same host. For example, the attacker could control (or be otherwise aware of) network activity of the other VMs running on the same host, and deduce other network activity is due to a victim VM. Comparing application TLS to guest IPsec ESP to NVE IPsec ESP, each provides stronger protection from traffic analysis in the same order. Application TLS exposes TCP port numbers to a passive observer, guest IPsec ESP encrypts the inner transport header but still identifies the communicating VM's IP address, while NVE IPsec ESP encrypts the entire inner payload.

To protect packet payloads from passive listeners, application-level encryption (e.g., JSON Web Encryption [RFC7516]), application TLS, guest IPsec ESP, or hypervisor IPsec ESP can be used. Each provides the same protection against a passive listener.

To protect against the above-described traffic sniffing attacks, we require:

REQ1: The NVE MUST ensure the traffic leaving the NVE has its payload encrypted.

REQ2: To provide best protection from traffic analysis, the NVE SHOULD encrypt the VM's inner IP address, transport header, and payload.

5.2. Protecting Against Traffic Injection

Traffic injection from a rogue non legitimate NVO3 Geneve overlay device or a rogue underlay transit device can target an NVE, a transit underlay device or a Tenant System. Targeting a Tenant's System requires a valid MAC and IP addresses of the Tenant's System.

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 does not protect the tenant system from receiving illegitimate packets that may disrupt the Tenant's System performance.

The Geneve overlay network MAY still need to prevent such spoofed Tenant's system packets from being steered to the Tenant's system.

When the Tenant's System are not protecting their communications, the Geneve overlay network SHOULD be able to to prevent a rogue device from injecting traffic into the overlay network.

In order to prevent traffic injection to one virtual network, the destination legitimate Geneve NVE MUST be able to authenticate the incoming Geneve packets from the source NVE. Authenticated Geneve Packet MAY be checked by underlay intermediary nodes.

Based on a policy partial authentication MAY be performed on Geneve packets if tenant's system is protecting it's communication.

This leads to the following security requirements:

REQ3: A Geneve NVE MUST be able to authenticate the Geneve tunnel Header, and/or the Geneve base header, and/or the immutable Geneve Options, and/or the Geneve payload.

REQ4: A Geneve NVE MAY be able to authenticate only a portion of the Geneve payload if the Tenant's system is protecting its communication.

REQ5: A GTN MAY be able to validate the authentication before the packet reaches the Geneve destination NVE.

REQ6: A GTN MUST be able to insert an authenticated Geneve Option into a authenticated Geneve Packet - protected by the source Geneve NVE.

REQ7: A GTN MUST be capable of forwarding the Geneve authenticated packet as an non-authenticated Geneve Packet.

REQ8: A Geneve NVE SHOULD be able to set different security policies for different flows. These flows MUST be identified from the Geneve Header and/or Geneve Options as well as some inner traffic selectors.

REQ9: In the case when Tenant systems secure their communications using protocols such as TLS or IPsec. A Geneve NVE MAY be able to selectively encrypt and/or authenticate only the sections that are not encrypted/authenticated by the Tenant System. For example, only the IP, transport (TCP / UDP) in case of TLS/DTLS MAY be encrypted/authenticated, while only the IP header and ESP header MAY be encrypted/authenticated.

5.3. Protecting Against Traffic Redirection

A rogue device of the NVO3 overlay Geneve network or the underlay network may redirect the traffic from a virtual network to the attacker for passive or active attacks.

To prevent an attacker located in the middle between the NVEs and modifying the tunnel address information in the data packet header to redirect the data traffic, the solution need to provide confidentiality protection for data traffics exchanged between NVEs.

Based on a policy partial encryption MAY be performed on Geneve packets if tenant's system is protecting it's communication.

This leads to the following security requirements:

REQ10: A Geneve NVE MUST be able encrypt Geneve base Header, and/or Geneve Payload and/or Geneve Options not intended for the GTN.

REQ11: A Geneve NVE MAY be able encrypt portion of Geneve Payload as well as as Geneve Options not intended for the GTN.

REQ12: A transit underlay intermediary node MUST be able to insert an encrypted Geneve Option into an encrypted/authenticated Geneve Packet - protected by the source Geneve NVE.

REQ13: A Geneve NVE SHOULD be able to assign different cryptographic keys to protect the unicast tunnels between NVEs respectively.

REQ14: If there are multicast packets, a Geneve NVE SHOULD be able to assign distinct cryptographic group keys to protect the multicast packets exchanged among the NVEs within different multicast groups. Upon receiving a data packet, an egress Geneve NVE MUST be able to verify whether the packet is sent from a proper ingress NVE which is authorized to forward that packet.

5.4. Protecting Against Traffic Replay

A rogue device of the NVO3 overlay Geneve network or the underlay network may replay a Geneve packet, to load the network and/or a specific Tenant System with a modified Geneve payload. In some cases, such attacks may target an increase of the tenants costs.

When traffic between tenants is not protected, the rogue device may forward the modified packet over a valid Geneve Header. The crafted packet may for example, include a specifically crafted application

payload for a specific Tenant Systems application, with the intention to load the tenant specific application.

Updating the Geneve header and option parameters such as setting an OAM bit, adding bogus option TLVs, or setting a critical bit, may result in different processing behavior, that could greatly impact performance of the overlay network and the underlay infrastructure and thus affect the tenants traffic delivery.

The NVO3 overlay network and underlay network nodes that may address such attacks MUST provide means to authenticate the Geneve packet components.

This leads to the following security requirements:

REQ15: A Geneve NVE or a GTN SHOULD be able to validate the Geneve Header corresponds to the Geneve payload, and discard such packets.

REQ16: A Geneve NVE or a GTN SHOULD provide anti replay mechanisms and discard replayed packet.

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. References

8.1. Normative References

[I-D.ietf-nvo3-geneve]
Gross, J., Ganga, I., and T. Sridhar, "Geneve: Generic Network Virtualization Encapsulation", draft-ietf-nvo3-geneve-05 (work in progress), September 2017.

[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>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/info/rfc4301>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <https://www.rfc-editor.org/info/rfc5246>.

[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>.

8.2. Informational References

[I-D.ietf-nvo3-security-requirements] Hartman, S., Zhang, D., Wasserman, M., Qiang, Z., and M. Zhang, "Security Requirements of NVO3", draft-ietf-nvo3-security-requirements-07 (work in progress), June 2016.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

[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>.

[RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <https://www.rfc-editor.org/info/rfc7516>.

[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>.

Authors' Addresses

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

Email: daniel.migault@ericsson.com

Sami Boutros
VMware, Inc.

Email: sboutros@vmware.com

Dan Wing
VMware, Inc.

Email: dwing@vmware.com

Suresh Krishna
Kaloom

Email: suresh@kaloom.com