INTERNET-DRAFT
Intended Status: Standard Track

Sami Boutros(Ed.)
VMware

Dan Wing
Calvin Qian
VMware

Expires: January 1, 2018

June 30, 2017

IPSec over Geneve encapsulation
draft-boutros-nvo3-ipsec-over-geneve-00

Abstract

This document specifies how Generic Network Virtualization Encapsulation (Geneve) can be used to carry IP Encapsulating Security Payload (ESP) and IP Authentication Header (AH) to provide secure transport over IP networks. Using IPSec ESP and AH will provide both Geneve header integrity protection and Geneve payload encryption.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Table of Contents

1. Introduction

The Network Virtualization over Layer 3 (NVO3) develop solutions for network virtualization within a data center (DC) environment that assumes an IP-based underlay. An NVO3 solution provides layer 2 and/or layer 3 overlay services for virtual networks enabling multi-tenancy and workload mobility. The Generic Network Virtualization Encapsulation [GENEVE] have been recently recommended to be the proposed standard for network virtualization overlay encapsulation.

Generic Network Virtualization Encapsulation (Geneve) does not have any inherent security mechanisms. An attacker with access to the underlay network transporting the IP packets has the ability to snoop or inject packets.

Within a particular security domain, such as a data center operated by a single provider, the most common and highest performing security mechanism is isolation of trusted components. Tunnel traffic can be carried over a separate VLAN and filtered at any untrusted boundaries. In addition, tunnel endpoints should only be operated in environments controlled by the service provider, such as the hypervisor itself rather than within a customer VM.

When crossing an untrusted link, such as the public Internet, IPsec [RFC4301] may be used to provide authentication and/or encryption of the IP packets formed as part of Geneve encapsulation. If the remote tunnel endpoint is not completely trusted, for example it resides on a customer premises, then it may also be necessary to sanitize any tunnel metadata to prevent tenant-hopping attacks.

This document describes how Geneve tunnel encapsulation [GENEVE] can be used to carry both the IP Encapsulation security payload (ESP) [RFC4303] and the IP authentication header (AH) [RFC4302] to provide secure transport over IP networks. Using IPSec ESP and AH will provide both Geneve header integrity protection and Geneve payload encryption.

2. Terminology

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 RFC 2119 [RFC2119].

3. Abbreviations

NVO3 Network Virtualization Overlays over Layer 3

OAM Operations, Administration, and Maintenance

TLV Type, Length, and Value

VNI Virtual Network Identifier

NVE Network Virtualization Edge

NVA Network Virtualization Authority

NIC Network interface card

IPSec IP Security

ESP IP Encapsulating Security Payload

AH IP Authentication Header

GRE Generic Routing Encapsulation (GRE)

EtherIP Tunneling Ethernet Frames in IP Datagrams

Transit device Underlay network devices between NVE(s).

4.0 Encapsulation Security Payload (ESP) over Geneve tunnel

The Geneve packet is encapsulated in UDP over either IPv4 or IPv6. The Geneve packet consists of a Geneve header which is then followed by a set of variable options. The Geneve header protocol type will indicate that the Geneve payload is the ESP protocol (IP Protocol 50), and the Geneve payload will consist of the ESP protocol data.

The ESP next header can carry only an IP protocol so can't carry the inner Ethernet frame, so given that (1) Generic Routing Encapsulation (GRE) [RFC2784] and EtherIP [RFC3378] are IP protocols and (2) GRE/EtherIP can carry Ethernet Frames, hence the need of EtherIP/GRE encapsulation for the inner Ethernet payload.

The ESP Next Header field will be set to inner payload protocol which can be either EtherIP (97) or to the Generic Routing Encapsulation (GRE) (47). The GRE protocol type will be set to the Ethernet protocol type.

Inner Ethernet packet, as sent/received by the virtual machine:

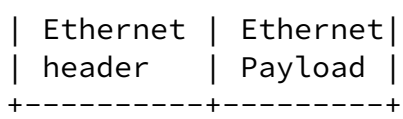

After applying Geneve and ESP, with ETH-IP header.

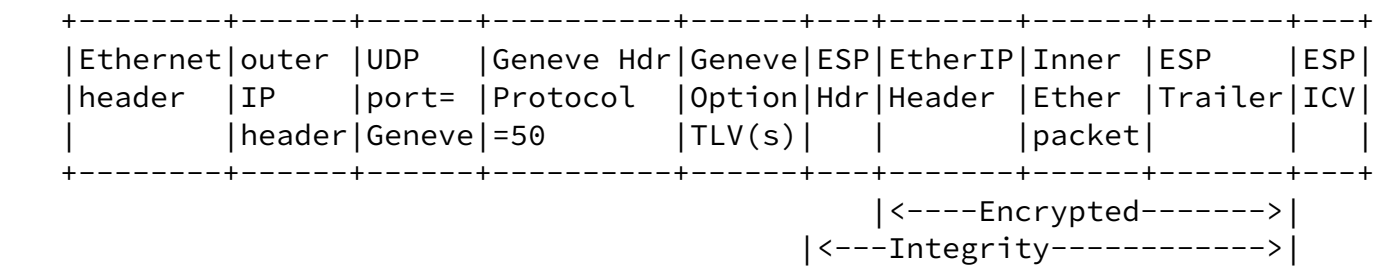

After applying Geneve and ESP, with GRE header.

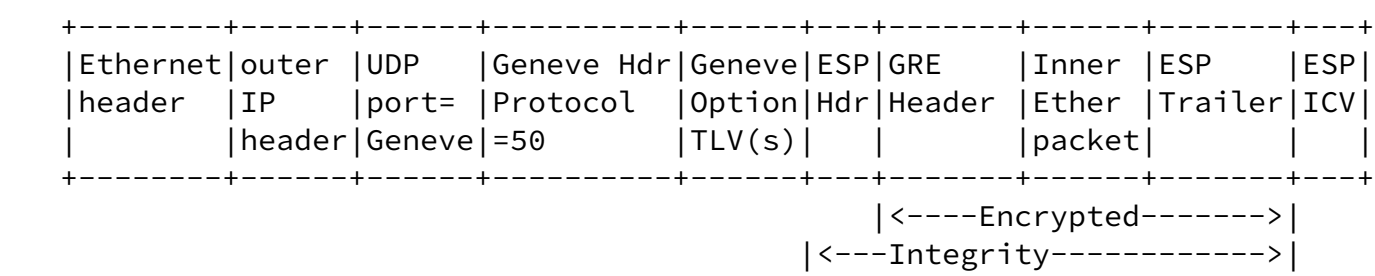

Figure 1: ESP over Geneve Packet Diagram

5.0 IP Authentication header (AH) over Geneve tunnel

The Geneve packet is encapsulated in UDP over either IPv4 or IPv6. The Geneve packet consists of a Geneve header which is then followed by a set of variable options. The Geneve header protocol type will indicate that the Geneve payload is the AH protocol (IP Protocol 51), and the Geneve payload will consist of the Authentication header (AH). The AH next header can carry only an IP protocol so can't carry the inner Ethernet frame, so given that (1) Generic Routing Encapsulation (GRE) [RFC2784] and EtherIP [RFC3378] are IP protocols and (2) GRE/EtherIP can carry Ethernet Frames, hence the need of EtherIP/GRE encapsulation for the inner Ethernet payload.

The AH Next Header field will be set to inner payload protocol which can be either EtherIP (97) or to the Generic Routing Encapsulation (GRE) (47). The GRE protocol type will be set to the Ethernet protocol type.

It is to be noted that some of the option TLV(s) in the Geneve header SHOULD be treated as mutable fields and not included in the AH authentication.

Inner Ethernet packet, as sent/received by the virtual machine:

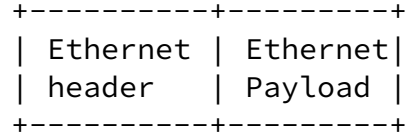

After applying Geneve and AH, with ETH-IP header.

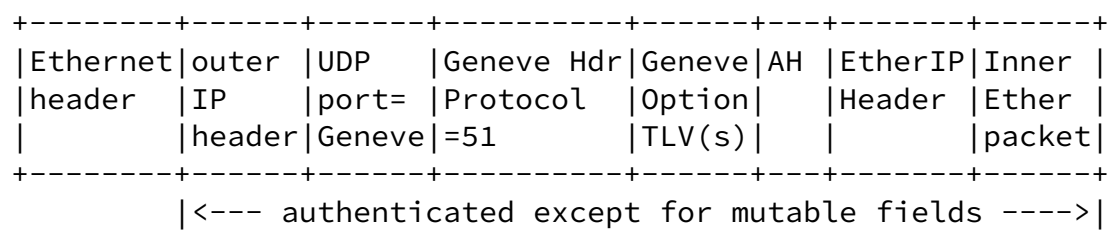

After applying Geneve and AH, with GRE header.

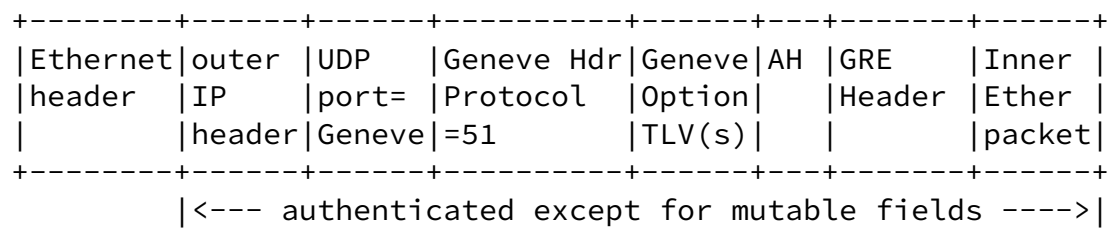

Figure 2: AH over Geneve Packet Diagram

6. Control Plane Considerations

A control plane extension could allow a Network Virtualization Endpoint (NVE) to express the next protocol that can be carried by Geneve to its peers.

In the datapath, a transmitting NVE MUST NOT encapsulate a packet destined to another NVE with any protocol the receiving NVE is not capable of processing.

In this document the next protocol signaled in control plane by NVE(s) can be ESP or AH.

Once 2 NVE(s) agree to carry ESP or AH as next protocol, Security Association and Key Management Protocol defined in [RFC2408] can be used to negotiate, establish, modify and delete Security Associations. As well, mechanisms to perform key exchange defined in [RFC2409] can be used.

7. Acknowledgements

The authors would like to thank T. Sridhar for his valuable comments.

8. Security Considerations

This document does not introduce any additional security constraints.

9. IANA Considerations

TBD

10. References

10.1 Normative References

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2 Informative References

[Geneve] "Generic Network Virtualization Encapsulation", [I-D.ietf-nvo3-geneve]

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

[RFC4302] Kent, S. "IP Authentication Header", RFC 4302, December 2005, <http://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S. "IP Encapsulating Security Payload", RFC 4303, December 2005, <http://www.rfc-editor.org/info/rfc4303>.

[RFC3378] R. Housley, et al. "EtherIP: Tunneling Ethernet Frames in IP Datagrams", RFC 3378, September 2002, <http://www.rfc-editor.org/info/rfc3378>.

[RFC2784] T. Li, et al. "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000, <http://www.rfc-editor.org/info/rfc2784>.

[RFC2408] D. Maughan, et al. "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998, <http://www.rfc-editor.org/info/rfc2408>.

[RFC2409] D. Harkins, et al. "The Internet Key Exchange (IKE)", RFC 2409, November 1998, <http://www.rfc-editor.org/info/rfc2409>.

Authors' Addresses

Sami Boutros
VMware
Email: sboutros@vmware.com

Dan Wing
VMware, Inc.
Email: dwing@vmware.com

Calvin Qian
VMware, Inc.
Email: calvinq@vmware.com