Network Working Group
Internet-Draft
Intended status: Standard Track

L. Yong
Huawei USA
T. Herbert
Google

Expires: April 2015

October 27, 2014

Generic UDP Encapsulation (GUE) for Network Virtualization Overlay
draft-hy-nvo3-gue-4-nvo-00

Abstract

This document describes network virtualization encapsulation scheme by use of generic UDP encapsulation (GUE) [GUE].

Status of This Document

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 April 27, 2015.

Table of Contents

1. Introduction

Network Virtualization Overlay (NVO3) [RFC7365] aims to a virtual network solution over an IP network in a DC with multi-tenant environment. Virtual network traffic between any pair of network virtualization edges (NVE) is encapsulated with a network virtualization header and is sent from ingress NVE to egress NVE as of an IP packet. This is known as a tunnel mechanism.

UDP based tunnel mechanism provides several merits for such tunneling applications.[GRE-in-UDP] This document specifies network virtualization encapsulation schema by use of generic UDP encapsulation (GUE) [GUE]. This allows NVEs to adopt GUE tunnel implementation.

This document specifies one flag (1 bit) for Network Virtualization Overlay (NVO) indication in GUE header and a Virtual Network ID field in GUE optional fields. It also specifies optional use of GUE secure transport capability for NVO.

2. Terminology

The terms defined in [RFC768] are used in this document.

2.1. Requirements Language

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].

3. Generic UDP Encapsulation (GUE) for NVO

Generic UDP Encapsulation adds a 32 bits basic GUE header after UDP header. GUE header contains some key fields that a UDP tunnel application needs. These key fields are version, control message indication (c), Header Length (HLen), and Protocol Type (or ctype). It also contains some undefined flags, which are reserved for tunnel applications. Figure 1 illustrates GUE structure and key fields. For the detail specification, see [GUE].

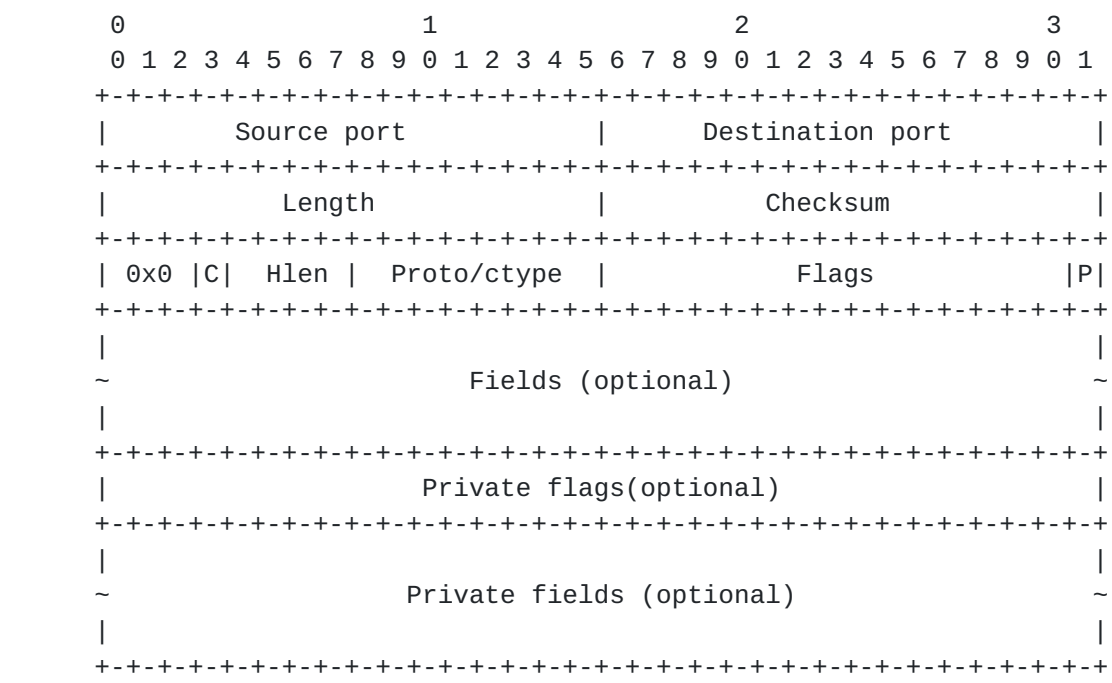

Figure 1 GUE Header Format

This document proposes to allocate one flag bit from GUE undefined flags for the Network Virtualization Overlay (NVO) and defines Virtual Network Identifier (VN ID) field for NVO in GUE optional fields. It also specifies use of GUE secure transport for NVO. The network virtualization header format is shown in figure 2 and the specification is followed.

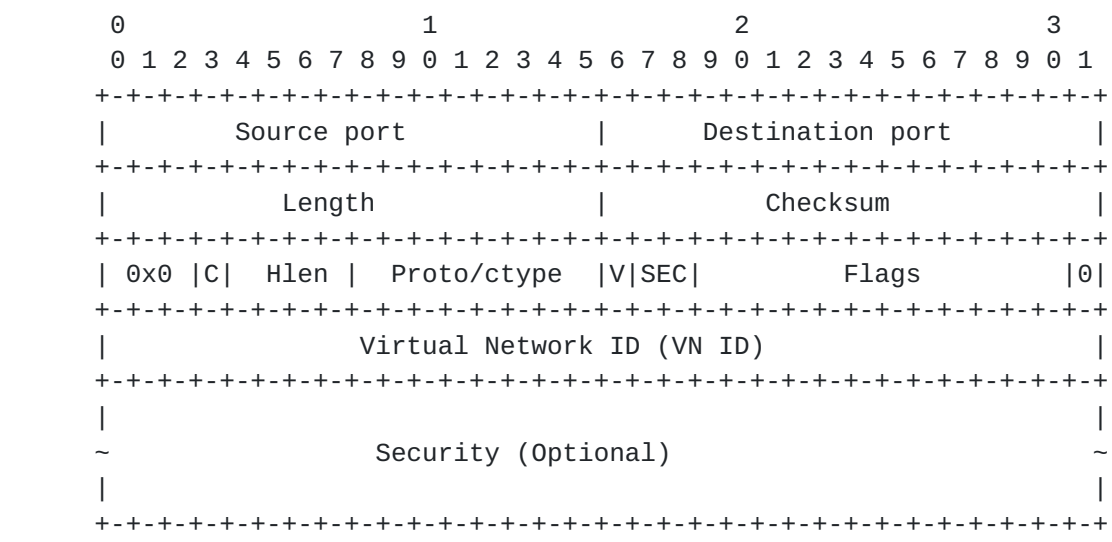

Figure 2 GUE for Network Virtualization Overlay

- 'V' Virtualization flag. Indicates presence of the Virtual Network Identifier (VN ID) field in GUE optional fields. This flag MUST be set when GUE is used for network virtualization overlay (NVO).

- Virtual Network ID (4 octets): Used in network virtualization overlay to identify a virtual network that packet was sent on. This field only presents if 'V' virtualization flag is set. Use and semantics of this field should be defined in separate documents.

- 'SEC' Security flags: Indicates presence of security field [GUE4SEC. It provides secure transport for a tunneled protocol. NVO MAY use it to provide secure transport. Thus this is optional fields for NVO. If the flag is set, i.e. not 00, the egress NVE MUST process the security field that is placed after VNI field. Use two bits for 'SEC' flag to convey the security field length as following.

  - 00 - No security field

  - 01 - 64 bit security field

  - 10 - 128 bit security field

  - 11 - 256 bit security field

The use of the security field is expected to be negotiated out of band between two NVEs. Potential uses of the security field for NVO is described in Section of Security Considerations.

The usage of the key fields in the GUE header [GUE] for network virtualization encapsulation is described as below:

- Type: Set to 0x0 for network virtualization overlay encapsulation.
- Control flag: When set, indicates the packet contains a control message. An OAM packet for the virtual network instance can be carried when it sets. Control or OAM processing MUST occur. The OAM protocol is out of scope for this document.
- Hlen: 1 if Security flags are clear. When Security flags are set, 1+ 2 ^ number(SEC flags)
- Protocol: Contain the protocol of the encapsulated payload packet, i.e. next header. The next header begins at the offset provided by Hlen. For network virtualization, the payload protocol can be Ethernet, IPv4 or IPv6.
- CType: Reserved for control message type. The VN ID can be used with CType to direct control message for the VN layer.
- 'P' Private flag. It is the last bit in the GUE header. This flag SHOULD be clear for the network virtualization encapsulation.

UDP header usage for network virtualization overlay is: UDP dst port SHOULD be filled with GUE port [GUE]; UDP src port MAY be filled with virtual network flow entropy. The checksum and length implementation MUST be compliant with GUE implementation [GUE].

4. Encapsulation/Decapsulation Operation

The network virtualization encapsulation schema specified in this document applies to both IPv4 and IPv6 underlay networks. The outer IP address must be NVE egress IP address (dst) and NVE ingress IP address (src). The network virtualization edge (NVE) implementation must compliant with the tunnel implementation specified in GUE [GUE] including GUE header process precedence.

When use of secure transport, NVE egress MUST perform security validation prior to the payload processing.

5. IANA Considerations

The document does not require any IANA action.

6. Security Considerations

Network Virtualization Edge (NVE) implements the UDP tunnel mechanism specified in [GUE] so it adopts the same security concern stated in Section of Security Considerations in [GUE].

Security option described in this document can be used improve the security in data plane for NVO applications. The security field may be used as a cookie. This would be similar to cookie mechanism described in L2TP [RFC3931], and the general properties should be the same. The cookie may be used to validate the encapsulation. The cookie is a shared value between ingress NVE and egress NVE which should be chosen randomly and may be changed periodically. Different cookies may used for logical flows between the ingress NVE and egress NVE, for instance packets sent with different VNIs in network virtualization might have different cookies.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC2119, March 1997.

[RFC3931] Lau, J., Townsley, M., et al, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC3931, 2005

[RFC7365] Lasserre, M., el al, "Framework for Data Center (DC) Network Virtualization".

[GUE] Herbert T. and Yong, L., "Generic UNP Encapsulation", draft-herbert-gue-02, work in progress.

[GUE4SEC] Yong, L., Herbert, T., "Generic UDP Encapsulation (GUE) for Secure Transport", draft-hy-gue-4-secure-transport-00, work in progress.

7.2. Informative References

[GRE-in-UDP] Grabbe, E., Yong, L., Xu, X., "Generic UDP Encapsulation for IP Tunneling", draft-ietf-tsvwg-gre-in-udp-encap-03, work in progress

8. Authors' Addresses

Lucy Yong
Huawei USA
5340 Legacy Dr.
Plano, TX 75024
US

Email: lucy.yong@huawei.com

Tom Herbert
Google
1600 Amphitheatre Parkway
Mountain View, CA
US

Email: therbert@google.com