

# Generic Network Virtualization Encapsulation

draft-gross-geneve-02

---

Jesse Gross, VMware

T. Sridhar, VMware

Pankaj Garg, Microsoft

Chris Wright, Red Hat

Ilango Ganga, Intel

Puneet Agarwal, Broadcom

Ken Duda, Arista

Dinesh Dutt, Cumulus

Jon Hudson, Brocade

IETF 91 Honolulu

November 10, 2014

# Geneve Overview

---

- Extensible encapsulation format to allow for future innovation
- Decouple control plane and data plane components to allow different rates of evolution
- Continue to use standard IP fabrics as an underlay
- Support for multiple encapsulated protocols and OAM

Geneve combines a UDP shim, small base header, and TLV options to achieve these goals.

# Draft Progress

---

- Many clarifications, corrections, and improvements based on feedback
- IANA has assigned UDP port 6081 to Geneve
- Still to be done: incorporate work from transport group on checksums and congestion control

# New Authors

---

Geneve now includes contributors from hardware OEMs, silicon, and software.

**Puneet Agarwal, Broadcom**

**Ken Duda, Arista**

**Dinesh Dutt, Cumulus**

**Jon Hudson, Brocade**

# Header Format

---

## UDP:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Source Port = xxxx           |           Dest Port = Fixed Port           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           UDP Length                    |           UDP Checksum                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

## Geneve Base Header:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Ver|  Opt Len |O|C|  Rsvd.  |           Next Protocol           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Virtual Network Identifier (VNI)           |           Reserved           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Variable Length Options                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

## Options:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Option Class           |           Type           |R|R|R| Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Variable Option Data                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

# Implementations

---

## **Software endpoints:**

Open vSwitch

Linux

## **Troubleshooting:**

Wireshark

tcpdump

## **NIC:**

Intel XL710 40Gbps adapter

Only 9 months in – many more to come in 2015.

# What's Next?

---

- Now that NVO3 has been rechartered, goal is adoption of Geneve as an NVO3 dataplane protocol for network virtualization
- Continue to gain experience through implementations and deployment