



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

NVo3 Interim Meeting

February 27, 2015



Geneve: Why?

- 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
- Existing encapsulation methods have created fractions in products for end users.

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

Geneve: Header



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           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```



Draft Progress



- No Significant changes since -02

Implementations

Software endpoints:

Open vSwitch

Linux



Troubleshooting:

Wireshark

tcpdump

NIC:

Intel XL710 40Gbps adapter



Goals

- Adoption of Geneve as an NVO3 dataplane protocol for network virtualization
- Continued evolution through increasing community support & experience.

