

Encapsulation Techniques: Generic Network Virtualization Encapsulation, VXLAN Generic Protocol Extension, and Network Service Header

Overview

Tagging and encapsulation techniques have long been a part of networking. As various approaches for virtual networks have gained support, network virtualization has become an increasingly popular topic in standards organizations such as the IETF. Standardization at the data-plane level is especially important to allow interoperability between physical and hypervisor-based switches and network services. Virtual Extensible LAN (VXLAN; RFC 7348) is one of the popular recent encapsulation technologies initially developed for data center environments.

Requirements for Encapsulation Protocols

As new encapsulation protocols are proposed to the IETF, they should be evaluated in the context of several principles:

- **Standardization:** A protocol must be a fully specified standard. This feature is essential so that hardware and software can be built to allow interoperability between vendors.
- **Hardware interoperability and efficiency:** The protocol must be efficiently implemented in hardware. Although software-based encapsulation approaches can always be used, modern data centers include a mix of hypervisor-based virtual machines, bare-metal workloads, and network services. All these components must be able to efficiently process packets and participate in an encapsulation protocol.

Generic Network Virtualization Encapsulation

The Generic Network Virtualization Encapsulation (Geneve) protocol¹ offers a new approach to encapsulation designed to offer control-plane independence between tunnel endpoints. The protocol specifies only a data-plane schema using a number of variable length options. Each packet contains an **option length** parameter, and the interpretation of different options is outside the specification. Figure 1 shows the Geneve header.

Figure 1. Geneve Header

Ver	Opt Len	O	C	Rsvd	Protocol Type
Virtual Network ID					Rsvd
Variable Length Options					

Geneve was designed to offer maximum flexibility across different control-plane implementations and use cases. It allows essentially any information to be encoded in a packet and passed between tunnel endpoints. Geneve is currently an IETF draft proposed by VMware, Microsoft, Red Hat, and Intel.

¹ <https://datatracker.ietf.org/doc/draft-gross-geneve>

Unfortunately, Geneve's approach to variable length options has several shortcomings:

- Due to the ambiguity of the length and meaning of the option fields, Geneve is not efficiently implemented in hardware and is not currently supported by a number of vendors of networking application-specific integrated circuits (ASICs). This limits Geneve's applicability to software-based solutions and makes support for bare-metal workloads and for network services difficult to achieve.
- The use cases for options are not specified. This may lead to limited interoperability between vendor solutions because there is no standardized set of options that can be supported across different platforms.
- Geneve assumes that only network endpoints need to parse the header and understand the connection state. In some cases, transit devices such as physical firewalls also need these capabilities, to enforce security requirements.

Cisco does not support Geneve in its Cisco Nexus[®] switches today and will be watching the Geneve protocol closely as it proceeds through the IETF. Cisco supports open standards and will continuously reevaluate support for Geneve as different use cases emerge.

VXLAN Generic Protocol Extension and Network Service Header

Cisco and its partners have been working within the IETF on two other encapsulation approaches: VXLAN Generic Protocol Extension (GPE)² and Network Service Header (NSH)³. These protocols address one of the common use cases to which Geneve may potentially be applied - specification of dynamic chains of network services - but they do so through a less generic approach, for which it is easier to design hardware platforms and help ensure multivendor interoperability.

VXLAN-GPE introduces a protocol type field into the VXLAN header (Figure 2). This protocol type describes the packet payload and currently defines types including IPv4, IPv6, Ethernet, and NSH.

Figure 2. VXLAN-GPE Header

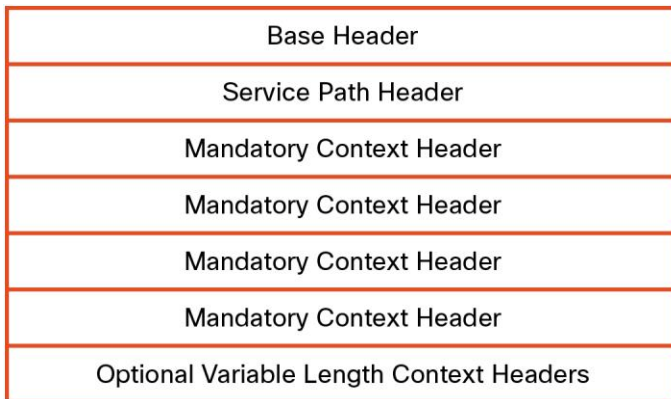
R	R	R	R	I	P	R	O	Ver	Rsvd	Next Protocol
Virtual Network ID									Rsvd	

Network Service Header was defined to offer a mechanism for network service path identification and metadata transmission (Figure 3). The goal of NSH is to create a topology-independent way of specifying a service path. NSH also includes a number of mandatory, fixed-size context headers designed to capture network platform information. NSH even contains an **optional** variable length metadata field for additional extensibility. However, unlike Geneve, NSH is designed to include all required information inside fixed-size fields. Although NSH works well with VXLAN-GPE, it is designed for multiple encapsulations (VXLAN, Cisco Locator/ID Separation Protocol [LISP], Multiprotocol Label Switching [MPLS], Generic Route Encapsulation [GRE], Network Virtualization Using GRE [NVGRE], etc.).

² <https://datatracker.ietf.org/doc/draft-quinn-vxlan-gpe>

³ <https://datatracker.ietf.org/doc/draft-quinn-sfc-nsh>

Figure 3. NSH



VXLAN-GPE is a draft currently proposed to the IETF by Broadcom, Cisco, Huawei, Intel, and Microsoft. NSH is a draft proposed by Broadcom, Cisco, Citrix, Intel, Microsoft, Rackspace, and Red Hat.

Conclusion

Geneve, VXLAN-GPE, and NSH are all recent protocol drafts proposed to the IETF. Geneve, while the most flexible of the three protocols, is not fully specified and has limitations related to multivendor interoperability and hardware support. VXLAN-GPE and NSH offer alternatives that can be more efficiently implemented in hardware with strong support for network service use cases.

Cisco is fully committed to promoting open networking standards and will continue to actively engage in standards bodies to help these protocols evolve.

For More Information

- <https://datatracker.ietf.org/doc/draft-gross-geneve/>
- <https://datatracker.ietf.org/doc/draft-quinn-vxlan-gpe>
- <https://datatracker.ietf.org/doc/draft-quinn-sfc-nsh>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)