

HVLAN

Hierarchical VLAN (HVLAN) is a proposed Ethernet standard that extends the use of enterprise Ethernet VLAN (802.1Q) to carrier networks. A number of developments have emerged in recent years to help bring Ethernet, a flexible and cost-efficient packet transport technology, to carrier networks. These developments include Q-in-Q (802.1ad), PBB (802.1ah), PBT (Provider Backbone Transport), and PBB-TE (Provider Backbone Bridge Traffic Engineering), which bring a set of features to traditional Ethernet to make it “carrier-grade”, adding to it high-availability, OA&M, and more.

While attempting to retain the core features that made Ethernet attractive in the first place, these technologies do not address other inefficiencies that could limit their use in the long term. This is especially true when considering the expected significant growth of multipoint network applications – IPTV, Private LANs, gaming, and others. The delivery of such services is better supported by PBB and associated protocols than alternatives such as MPLS, yet could hit scalability issues should services evolve as predicted.

HVLAN introduces the concept of hierarchical addressing schemes into the VLAN tag to provide both enterprise and carrier transport networks the characteristics they need in the long run.

Contents

Evolution of carrier-grade Ethernet standards

Ethernet

VLAN (802.1Q)

Q-in-Q (802.1ad)

Mac-in-Mac (802.1ah)

Conclusion

Links and References

Evolution of carrier-grade Ethernet standards

Ethernet

Ethernet is a connectionless technology. It does not have a routing mechanism and its address scheme is based on 48-bit MAC addresses. However, its flat address scheme results in a potential explosion of forwarding database entries and an uncontrolled flooding of broadcast messages throughout the network. In order to overcome Ethernet's scalability issues, a partitioning scheme, named VLAN, was introduced.

VLAN (802.1Q)

A **virtual LAN**, commonly known as VLAN, is a method of creating independent logical Ethernet networks within a physical network. Several VLANs can co-exist within such a network. This helps in reducing the broadcast domain and aids in network administration by separating logical segments of a LAN

(like company departments) that should not exchange data using a LAN (they still can exchange data by routing).

VLANs are configured through software rather than hardware, which makes them extremely flexible. Frames having a VLAN tag carry an explicit identification of the VLAN to which they belong. The value of the VLAN Identification (VID) in the tag header signifies the particular VLAN the frame belongs to. The main problem with VLAN is its limited VID space (4096). While this space may suffice for enterprise applications, it is much too small for carrier networks, which must support many customers and services.

Q-in-Q (802.1ad)

A number of solutions have been proposed to increase VLAN's scalability. A first proposal, called Q-in-Q, also known as Provider Bridge, VLAN stacking or tag stacking, allows service providers to **insert an additional VLAN tag** (referred to as provider VLAN) in the Ethernet frame in order to identify the service, resulting in a unique 24-bit length label. While this solution enables one, in theory, to identify up to 16 million services ($4094 * 4094$), in reality, one provider VLAN is dedicated to one customer, and therefore the number of supported customers is still limited to 4094.

Q-in-Q also introduces a scalability issue within the core of the carrier network, where every core switch needs to learn and maintain forwarding entries for every customer MAC address.

Mac-in-Mac (802.1ah)

PBB, PBT, and PBB-TE use an alternative proposed solution, known as MAC-in-MAC, described in the proposed IEEE 802.1ah Provider Backbone Bridges standard, which encapsulates Ethernet frames with a Service Provider MAC header. MAC-in-MAC technology overcomes the inherent scalability limitations of VLAN and Q-in-Q networks that make them impractical for use in larger networks by enabling up to 4000 times as many service instances as supported by traditional VLAN and Q-in-Q networks.


In PBB and PBT switches at the edge of the carrier network **encapsulate customer traffic within an 802.1ah frame**. The carrier network core is only responsible for transporting frames from an edge device to another, alleviating Q-in-Q's issue of forwarding table scalability. The same feature – assigning a MAC address per edge device, not per service – creates a scalability issue for multipoint services. Multipoint services require full mesh connectivity between edge devices, a very inefficient method as all frames are duplicated at the root nodes, rather than at the optimal point as in VLAN connectivity. Moreover, the need to create forwarding entries for each unicast connection within the full mesh (as opposed to a single VLAN forwarding tree in the case of VLAN connectivity) will quickly become unsustainable as multipoint services become predominant in the near future.

Furthermore, the addition of a MAC header augments the frame size by about 128 bits, a significant overhead given the small size (64 byte) of real-time application (e.g. voice and video) packets.

Accordingly, there is a long felt need for leveraging the forwarding efficiency of VLAN networking, while at the same time solving its addressing space scalability issues described previously. Increasing the VLAN tag size would mean bigger forwarding table, longer forwarding table entries, and a modification of current mass-market Ethernet chips, requirements that are not vital to the enterprise world.


HVLAN introduces **hierarchy into the VLAN tag**, in a way somewhat similar to classless subnets in the Internet Protocol with Classless Inter-Domain Routing (CIDR). Consequently, forwarding at each node uses a “best match” approach that substantially reduces the number of forwarding entries in core switches. Additionally HVLAN removes the need for encapsulation in many cases, reducing the overall transport overhead. The proposed HVLAN frame format is as follows:

A full description of the HVLAN header can be found in [1], the most important field being the HVID. When traversing the carrier's Ethernet network, the HVLAN frames can be forwarded using HVID only, MAC address only or a combination of both. There is an explicit bit in an HVLAN frame that prevents the carrier's core switches from learning the HVLAN frames' MAC address when unnecessary. To understand HVLAN operation, consider a scenario (see diagram) which illustrates the provision of 3 point-to-point services (blue, green, and red) over an HVLAN network. The diagram shows all forwarding table entries needed to transport the 3 services. Only forwarding entries for one direction (left-to-right) are displayed, similar entries implement the other direction.



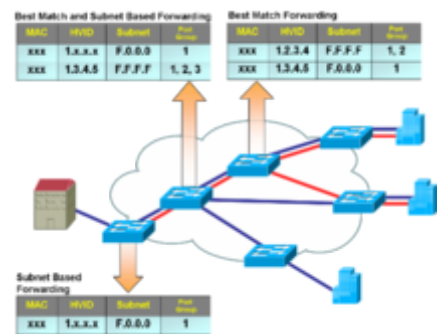
HVLAN Frame Format

Point-to-point services are provisioned using a unique HVID per service. Planning HVIDs wisely enables summarization (as shown at the leftmost edge device) and reduces the number of forwarding entries to a strict minimum; the network now scales to support millions of point-to-point services with minimum packet overhead (it can be noted that no encapsulation was used, frames were forwarded using HVID only).



Point-to-Point Scenario

A further example (see diagram) shows HVLAN operation in the case of **point-to-multipoint services** (e.g. IPTV). The diagram shows all forwarding table entries needed to transport the 2 multipoint services (red and blue) from a server (left) to 3 clients (right).



Point-to-Multipoint Scenario

As with point-to-point services, point-to-multipoint services are provisioned using a unique HVID per service. Encapsulation is not required and frames can be forwarded using HVID only. Summarization of HVIDs reduces the size of forwarding tables and creates scalability. Millions of point-to-multipoint services can be provided. The case of multipoint-to-multipoint is handled by HVLAN using encapsulation and provider MAC addresses. A full description of HVLAN multipoint-to-multipoint operation is provided in [1].

Conclusion

Hierarchical VLAN is a proposed extension to VLAN which, like PBB and PBT, turns cost-efficient Ethernet into a flexible, carrier-grade transport technology. Unlike other technologies, HVLAN uses the mature VLAN functionality to support all connectivity schemes: point-to-point, point-to-multipoint, and multipoint-to-multipoint. It uses a hierarchical VLAN allocation technique to achieve this. The technique allows summarization to reduce the number of forwarding table entries within the carrier network switches.

HVLAN is compatible with VLAN-related standards. It is currently being discussed by the ITU-T and the IEEE with the goal of standardization.

Links and References

[1] HVLAN White Paper coming soon

Retrieved from "<https://en.wikipedia.org/w/index.php?title=HVLAN&oldid=879276666>"

This page was last edited on 20 January 2019, at 05:21 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.