ECRA D2.4:

State of VLAN in the Automotive Domain

Patrik Thunström / TCN

21th November, 2018

State of VLAN in the Automotive domain

A look at current and emerging use cases

Table of contents

Original purpose and use case of VLANs

The 802.1Q Standard, originally named “IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks” [1] , is where the concept of a Virtual Local Area Network (VLAN) was initially standardized by the IEEE.

Since the 2014 version of the standard [2], the last part of the name has been changed to “Bridges and Bridged Networks”, as it has encompassed and integrated more than 25 standards not only related to VLANs but also traffic shaping, different versions of spanning tree algorithms, flow control, priority handling and stream reservation.

Some of these standards has been developed within the IEEE task group for Audio Video Bridging (AVB) which later on was renamed Time Sensitive Networking (TSN) to signify it’s expanded area of interest. This is worth noting, as work driven with this IEEE task group is intended to improve aspects of Ethernet for use within, among others, the automotive domain, something that will be covered later in this report.

At its time of conception [3], the purpose of VLANs was to allow multiple separate logical networks to communicate over a link shared between these separate networks. Since this work was done many years ahead of the initial version of the 802.1Q standard and predated the point to point based nature of today’s Media Access Controller (MAC) and Physical layer (PHY) standards, the communication was done as broadcasts on a single collision domain, very similar to a traditional, bus-based backbone in the automotive world. By the introduction of the VLAN concept, which was accomplished by adding an additional header to each frame, bridges could easily determine what frames were of interest to relay into each separate network by looking if they were members of the group denoted by the header.

This construct allowed for large networks consisting of sub networks to communicate efficiently, without the cost or penalties that would be needed to route traffic using IP, and for network sizes that would be impossible to fit in the addressing tables of the network equipment at the time.

VLAN Tag

The header that was introduced is what evolved into what today is known as the VLAN tag. The full length of the header is 32 bits, and consists of a Tag Protocol Indicator (TPI) with the 16 bit EtherType value that identifies the field as a VLAN tag, and the Tag Control Information (TCI) containing the actual VLAN information. The TCI in turn consists of the sub fields for VLAN ID (VID), Priority Code Point (PCP) and a Drop Eligibility Indicator (DEI). As the EtherType is the leading 16 bits, the remaining 16 bits are the TCI fields, divided as follows:

Size	3 bits	1 bit	12 bits
Field	Priority Code Point (PCP)	Drop Eligible Indicator (DEI)	VLAN ID (VID)

The 12 bits available for the VID means that VID values 0-4095 are possible, for a total of 4096 available values. Out of these values 0 and 4095 are reserved, leaving 4094 valid values for VLAN IDs. VID value 0 will be explained in section on Prioritization, while value 4095 is implementation specific, which for many vendors is used as a wildcard saying all VLAN IDs will be received on a port, and should not be used as a normal VLAN ID.

VLANs today

Since the introduction of VLANs in the late 80s much has changed within the ecosystem of networking. One of the largest differences, which itself has pushed forward development of new concepts and forced a lot of changes, is the transition into the point-to-point nature of today’s MACs and PHYs. This change has ensured that each link has its own collision domain, allowing for full duplex communications and guarantees on available bandwidth. It also means that the complexity of the bridges, or switches as will be used interchangeably for the rest of the document, has increased manifold.

The main use case for VLANs is still to transport multiple virtual networks over the same physical link, but for very different reasons than originally intended.

Isolation - Broadcast limitation and security considerations

As more critical and sensitive services has moved into the networking domain there has been more and more reason to limit the communication between end nodes. As a network grows larger, the number of broadcast messages also increases, which eventually will become a considerable number of messages to process.

For a network not utilizing VLANs both of these issues would be applicable, as all broadcasts and direct communication would be let through without any consideration. When utilizing VLANs this is one of the main reasons; Isolation between nodes.

For a switch using VLANs, each port would be configured to have a default Port VID (PVID), and to have a list of VIDs for which VLANs the port is a member of. For any traffic received on the switch port without a VLAN tag, the frames are automatically assigned the PVID, and any frame with a VLAN tag retains the VID from the tag. All frames are then only forwarded to ports that are members of the VLAN with the same VID the frame contains, and also meets the basic addressing based on MAC addresses. As such, any broadcast traffic would only be forwarded to other members of the same VID, introducing the concept of broadcast domains. In the same way any malicious node or attacker that would try to communicate with an end point would be unable to do so unless they were connected to a port that was also a member of the same VLAN.

Prioritization - 802.1P & Traffic Shaping

The mechanism used to enable prioritization of traffic is contained within the VLAN tags themselves. Part of the VLAN tag is the PCP (Priority Code Point) which by its length of 3 bits allows for assignment of 8 possible priority values to each frame.

Depending on the capabilities and configuration of the switch port the data flows can then be prioritized accordingly. The initial approach was through the work performed by the 802.1p Task Force, which allowed mapping of the priority levels into separate queues, which was then allowed to transmit in priority order. The work done by 802.1p was published in the 1998 revision of the 802.1D standard [4], and the additions to 802.1D was later integrated into the 2014 revision of the 802.1Q standard [2].

More complex means of prioritization for handling different types of dynamic or deterministic behavior also exists through the work of the AVB, and later TSN, group, with many types of traffic shapers. Credit-Based Shaper, 802.1Qav [5], is one of these, where the PCP field is used to categorise traffic into regular best effort traffic, and two classes of AVB traffic which has a higher priority, but that is restricted from overusing the link by an internal credit counter that is increased and decreased by a traffic shaping algorithm pre-configured to only allow traffic up to a specific limit for each traffic class.

Defined by the TSN Task Group as an amendment to the Credit-Based Shaper, after it was incorporated into the 2014 version of the 802.1Q standard, there is also Time-Aware Shaper, 802.1Qbv [6], which adds an additional traffic class to be used for communication where delays need to be deterministic. This is achieved by providing each switch in the path for the data flow with a schedule for which traffic classes can be forwarded at each point in time. For the dataflow supposed to have deterministic delay this means that as long as the frames are sent within this time window it will be forwarded without anything blocking. Synchronizing switches, using the 802.1AS standard [7] also introduced by the AVB Task Group, each switch along the path can then have their schedules configured to ensure determinism along the entire path. 802.1Qbv was recently incorporated into the 2018 version of the 802.1Q standard [8].

Prioritization in itself and the traffic shaping methods described briefly are not functionally a part of the VLAN concept itself, but as they are enabled by the inclusion of VLAN tags, they are worth mentioning. The PCP and VID fields themselves are entirely separate inside the VLAN tags, and as such VLANs and priority settings are fully independent from each other.

As mentioned earlier, VLAN ID 0 is reserved, and it is specifically for the case where a frame should not be tagged with a VID, but need to use the PCP fields for prioritization.

Contrasting VLAN to traditional Automotive architecture

Trying to bring the current functionality of VLANs into context of the in-vehicle architectures we can start with a high level example of a non-ethernet enabled architecture. In this theoretical scenario we have multiple functional domains, interconnected by gateways (GW) connected through a backbone allowing for ECUs outside a domain to still access data from an ECU in a different domain.

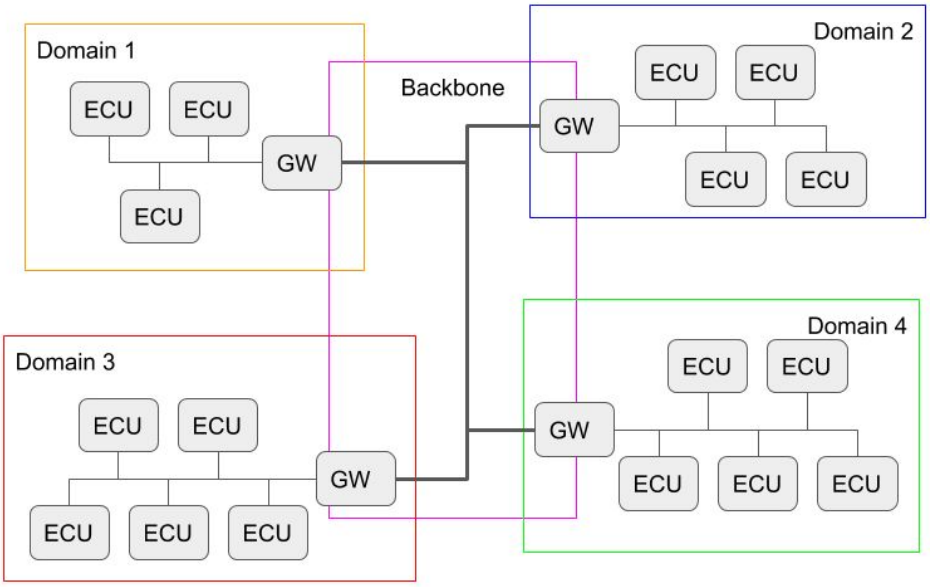

In the example above the Backbone is in recent years traditionally a FlexRay bus, while the communication inside of the domains are usually done over CAN or LIN buses, and, in some specific domains with higher requirements on bandwidth or synchronization, it can also be FlexRay or MOST.

Transitioning this into automotive Ethernet, the backbone could be replaced with an Ethernet switch, and each of the gateways connected to a port on this.

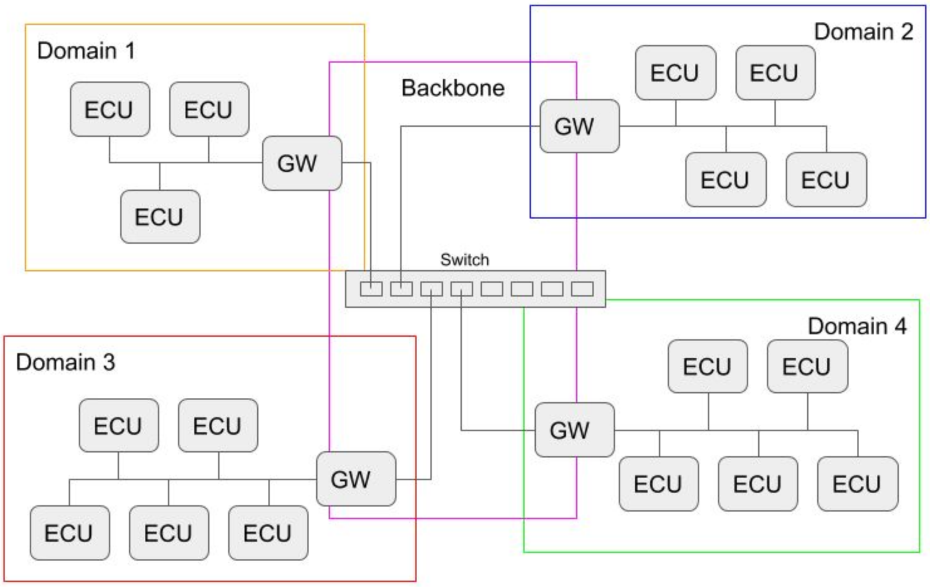

As long as the gateways remain, they would still be capable of filtering what data is allowed to go from one domain to another, and would not require VLAN capabilities. For security reasons, even with the traditional backbone approach over Ethernet, the isolation functionality enabled by VLANs could be a great advantage and simplify when trying to harden the network. If, for example, an ECU or a Gateway in domain 4 has external connectivity, restricting communication with it would greatly reduce the risks of external access.

Taking the most extreme approach, using a VLAN capable switch one could eliminate the need for Gateways, and connect each ECU to a central switch.

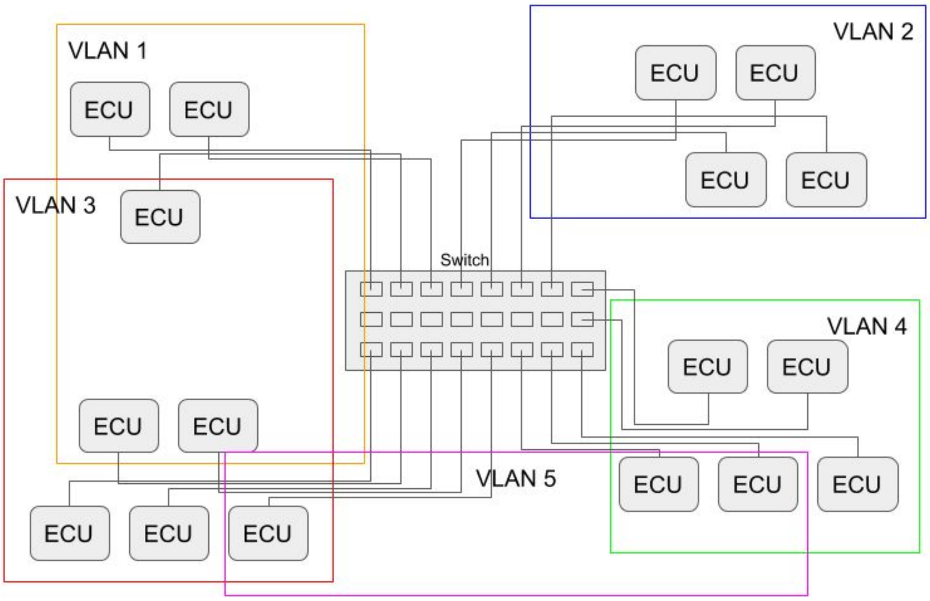

In this scenario the granularity for security and isolation is very fine, as each individual port can be included or excluded from every individual VLAN, allowing for very fine groupings between ECUs allowed to communicate with each other. To denote this, the domains shown in the two previous examples have been stretched and altered into new VLAN groupings.

This extreme example is not very likely in a foreseeable future, as there is a lot of commodity hardware that will still be connected through CAN and LIN, and because of that the need for Gateways will still be present. In addition, with regards to the topology there is a clear lack of redundancy, noticing that, if the central switch would stop operating, none of the ECUs would be able to communicate with each other. Also, the currently available and the trends in automotive grade Ethernet switches do not point to switches with such a high number of ports that a central switch would be possible, so a hierarchical topology or a topology with a few interconnected switches will be required at a minimum. For the use of VLAN and granularity in grouping this does not matter, however, as VLANs can exist over and cover multiple bridges.

Standards and topics of emerging interest

In addition to what is available today, what is being developed, enhanced and introduced by the TSN group, there are a few more areas that can be of interest for the automotive domain.

QinQ / Double tagging - 802.1ad

Defined in the 802.1ad standard [9], which was incorporated into the 2011 revision of 802.1Q [10], the concept of “Provider Bridges” was introduced. In essence this allows for two layers of VLANs, where a frame can have an “Outer” VLAN tag in addition to the previously existing “Inner” VLAN tag. The outer tag is discerned from the inner tag by having a different EtherType. This concept is also referred to as double tagging, or QinQ, as a “Q tag” is embedded in a outer “Q tag”.

By having two layers of VLANs it allows for simple routing inside the network, which is the original intention. This could for example allow a module with external connectivity to forward a broadcast message carried to a central switch, where the outer tag could be removed and all applicable VLANs addressed by the inner tag would receive the broadcasted frames.

A second use case for the two layer approach is for use by, and with, test and development hardware. A switch with port mirroring could be used to log when a frame passes through a specific link, but when trying to log, measure and analyse how the frame passes through a switch connecting both of the links to the same switch configured for port mirroring would normally create a loop in the network. Isolating each link forwarded through the switch by the use of a VLAN would let this setup continue to work, but if the frame itself is already VLAN tagged it would mean the original VLAN would be lost. With 802.1ad the mirroring switch could instead add an outer VLAN tag, which would allow for keeping the original VLAN tags unmodified.

VLAN as criteria for ACL/Firewalls

While VLANs themselves restrict communication between hosts, moving into Layer 3 and up the existence of VLANs can be used for further filtering and security restrictions. Layer 3 switches quite commonly has Access Control Lists (ACL), which allows configuring which IP addresses and ports, when using TCP or UDP for transport, frames will be let through for, based on various criterion. In case the network is divided into different subnets, this would enable the VLAN as a control for what traffic would be routed between them.

The more interesting use case is filtering on destination port based on VLAN. This would ensure that even though a node is allowed to communicate with another end node as members of the same VLAN, it could be configured to only be allowed to communicate with specific ports, ensuring no other services would be accessible.

There is currently a clear lack of automotive grade layer 3 switches, even if there are products for the industrial market. So, whether layer 3 switches and ACL functionality will be seen in the automotive domain still remains to be seen. The same isolation and filtering could, however, be achieved using most firewalls. The difference will be that processing is done by a CPU instead of dedicated hardware, leading to lower performance and possibly additional delays.

Hardware support and limitations

As the 802.1Q standard has been available since before Automotive Ethernet started to be widely introduced, many of the vendors of switch chips have had VLAN functionality in place from their very first products. Looking at the public product information and data sheets from Broadcom, Marvell, Microchip, NXP and Realtek, all of their automotive grade switch chips does provide VLAN functionality.

Where historic data was available, it was clear that some of the vendors did have limitations in the number of VLANs in use simultaneously, normally with the limit at 128 VIDs, but latest generation switch chips were all capable of the full range of 4096 VLAN IDs where numbers were provided.

Having a limit of 128 VIDs is not an issue in itself, as long as the data streams going through the switch do not need more VLANs than so. For a central switch this may be limiting, while for switches further out in the topology it would more rarely be a problem.

The support for QinQ/802.1ad is also very high, with products from Broadcom, Microchip and NXP listing support for it.

References

1. *IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks*, IEEE Standard 802.1Q-1998, 1998.
2. *IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks*, IEEE Standard 802.1Q-2014, 2014.
3. W. D. Sincoskie and C. J. Cotton, “Extended bridge algorithms for large networks,” *IEEE Network*, vol. 2, no. 1, pp. 16–24, Jan. 1988.
4. *IEEE Standard for Local Area Network MAC (Media Access Control) Bridges*, IEEE Standard 802.1D-1998, 1998.
5. *IEEE Standard for Local and metropolitan area networks-- Virtual Bridged Local Area Networks Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams*, IEEE Standard 802.1Qav-2009, 2009.
6. *IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks - Amendment 25: Enhancements for Scheduled Traffic*, IEEE Standard 802.1Qbv-2015, 2015.
7. *IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks*, IEEE Standard 802.1AS-2011, 2011.
8. *IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks*, IEEE Standard 802.1Q-2018, 2018.
9. *IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks---Amendment 4: Provider Bridges*, IEEE Standard 802.1ad-2005, 2005.
10. *IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks*, IEEE Standard 802.1Q-2011, 2011.