Power, Connect and Protect

Virtual LAN

Applications and Technology

A WHITE PAPER

Micrel, Incorporated 2180 Fortune Drive San Jose, CA 95131 USA TEL: +1 (408)-944-0800 FAX: +1 (408) 474-1000
www.micrel.com

CONTENTS

Introduction

Virtual LANs (VLANs) have recently developed into an integral feature of Ethernet switch solutions. One of the reasons for the importance placed on VLAN functionality today is the trend of using Ethernet LAN switches as a replacement for networking hubs.

Broadcast packets perform a number of behind-the-scenes, yet indispensable, network functions. For example, whenever an IP host or router needs to find the physical destination of an IP address, it generates an Address Resolution Protocol (ARP) broadcast packet, which asks something like, "Will the node with IP address x please send me your MAC address, so I can address this packet to you?" (These addresses are stored in an ARP cache for a couple of hours or so - subsequent traffic may not require broadcasts.)

On a single-segment Ethernet network (half duplex shared medium), the broadcast domain is the same as the collision domain. Multiple Ethernet segments connected via a switch form an extended broadcast domain made up of multiple collision domains. The switch forwards broadcast packets out to every port, and they also forward frames with unknown destination addresses to every port. The unicast traffic local to a collision domain remains in the same domain and is invisible to the rest of the broadcast domain, while broadcast frame and frames with unknown destinations are flooded to every other node on the broadcast domain. Thousands of nodes could be affected by this architecture.

With the emergence of low-cost switches, each network segment now contains only one user (desktop LAN switching), while broadcast domains can be still be as large as 1,000 users or more. These large, completely flat networks can become unworkable because of broadcast traffic. Segmenting networks with switches can extend indefinitely the bandwidth available for unicast messages; however, the broadcast traffic will eventually swamp the whole network. And in some cases, a broadcast storm could require shutting down part of the network in order to recover its operational state.

Therefore, in order for the network to operate properly, LAN broadcast domain cannot be very large. In a typical network implementation, Layer 2 switches and collapsed backbone routers are used to segment the network at Layer 3 and thus also contain broadcast traffic. VLANs represent an alternative solution to expensive routers for broadcast containment, since VLANs allow switches to contain broadcast traffic.

What is a VLAN?

A VLAN can be roughly equated to a broadcast domain. More specifically, VLANs can be seen as analogous to a group of end-stations, perhaps on multiple physical LAN segments that are not constrained by their physical locations and can communicate as if they were on a common LAN.

VLANs are set up between switches by inserting a tag into each Ethernet frame (see Figure 1). A tag field containing VLAN (and/or 802.1p priority) information can be inserted into an Ethernet frame. If a port has an 802.1q-compliant device attached (such as another switch), these tagged frames can carry VLAN membership information between switches, thus letting a VLAN span multiple switches.

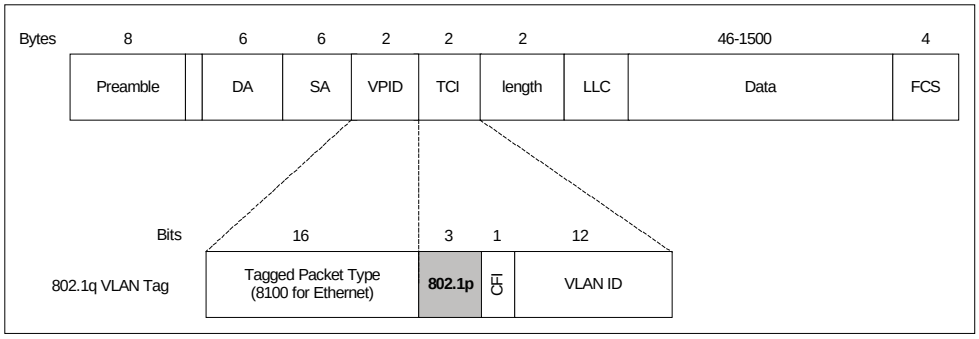

Figure 1. VLAN Tag in an Ethernet Frame Format

Port-based VLAN

VLAN membership can be defined in several different ways, which is one of the reasons they can be confusing. The simplest way is to assign specific ports on a switch to VLANs. For example, ports 1, 2, 7, and 8 on an 8-port switch make up VLAN A, while ports 3, 4, 5, and 6 make up VLAN B. If an assigned switch port is connected to a shared segment - one with multiple nodes - all those nodes will be part of the VLAN, along with the nodes on other ports assigned to that VLAN. This VLAN segmentation is basically a matter of electronically isolating the ports of each VLAN.

Port grouping is the most popular method of defining VLAN membership, and configuration is fairly straightforward. However, the primary limitation of defining VLANs by port is that the network manager may have to reconfigure VLAN membership when a user moves from one port to another.

MAC Address-based VLAN

MAC address-based VLAN has a different set of advantages and disadvantages. In this method, switches maintain tables of MAC addresses with their VLAN membership. With this form of membership, the nodes on a shared segment need not belong to the same VLAN. Because MAC addresses are hardwired into the Network Information Cards (NICs), this Layer 2 VLAN definition has the advantage of portability. Wherever a workstation or laptop is plugged in across the switched network, the switches will recognize it as a member of the assigned VLAN. This is especially useful for a WIFI connected switch port.

The disadvantage of MAC address-based VLAN solutions is the requirement that all users must initially be configured to be in at least one VLAN. Only after that initial manual configuration, the automatic VLAN grouping of users is possible. The manual VLAN configuration has to be repeated, if more than one switch is used. The disadvantage of having to configure VLAN membership becomes apparent in very large networks where thousands of users must each be explicitly assigned to a particular VLAN.

MAC address-based VLANs that are implemented in shared-media environments (e.g. WIFI) run into serious performance degradation as members of different VLANs coexist on a single switch port. This occurs because several broadcast domains coexist on the same shared-media that is to become one broadcast domain.

Layer 3-based VLANs

VLANs can also be defined by their network or Layer 3 addresses. Many network managers, frustrated by the administrative overhead that comes with highly changeable networks configured with IP, could happily replace their existing subnet structure with VLANs. Once the node's VLAN membership is defined, the VLAN handles everything, even if the node is moved to a port connected to a different subnet or if the IP addresses need to be reassigned.

Layer 3-based VLANs take into account protocol type (if multiple protocols are supported) or network-layer address (for example, subnet address for TCP/IP networks) in determining VLAN membership.

There are several advantages to defining VLANs at Layer 3. First, it enables partitioning by protocol type. This may be an attractive option for network managers who are dedicated to a service - or application-based VLAN strategy. Second, users can physically move their workstations without having to reconfigure each workstation’s network address—a benefit primarily for TCP/IP users. Third, defining VLANs at Layer 3 can eliminate the need for frame-tagging in order to communicate VLAN membership between switches, reducing transport overhead.

One of the disadvantages of defining VLANs at Layer 3 (versus MAC - or port-based VLANs) can be performance. Inspecting Layer 3 addresses in packets is more time consuming than looking at MAC addresses in frames. For this reason, switches that use Layer 3 information for VLAN definition are generally slower than those that use Layer 2 information.

IP Multicast Groups as VLANs

IP multicast groups represent a somewhat different approach to VLAN definition, although the fundamental concept of VLANs as broadcast domains still applies. When an IP packet is sent via multicast, it is sent to an address that is a proxy for an explicitly defined group of IP addresses that is established dynamically. Each workstation is given the opportunity to join a particular IP multicast group by responding affirmatively to a broadcast notification, which signals that group's existence. All workstations that join an IP multicast group can be seen as members of the same virtual LAN. However, they are only members of a particular multicast group for a certain period of time. Therefore, the dynamic nature of VLANs defined by IP multicast groups enables a very high degree of flexibility and application sensitivity. In addition, VLANs defined by IP multicast groups would inherently be able to span routers and thus WAN connections.

VLAN and QoS

In order to have a robust network for both multimedia and data services, Quality of Service (QoS) must be deployed so that delay sensitive packets such as voice and video can have the lowest delay. VLAN ID and User Priority can be used for such a purpose. Before network traffic receives differentiated treatment, it is necessary to first classify traffic and "mark" the packets in a way that indicates that these specific packets warrant different (or better) treatment than other packets. Packets are prioritized through the use of the user priority field of the VLAN tag.

The network operator may define up to eight classes of service using the bits in user priority bits in the VLAN tag (see Figure.1). Then it is possible to utilize other QoS features to assign appropriate traffic-handling policies, including congestion management, bandwidth allocation, and delay bounds for each traffic class.

As traffic comes into the switch, it is assigned to one of the output queues (traffic classes). Packets on the highest-priority queue are transmitted first. When that queue empties, traffic on the next highest-priority queue is transmitted, and so on.

The recommended traffic classifications are listed below.

IEEE 802.1D User Priority Class Descriptions –801.2D (1998) – Section H.2.2

User Priority	Application Class
7	**Network Control** Characterized by a "must get there" requirement to maintain and support the network infrastructure.
6	**"Voice"** Characterized by less than a 10ms delay, and hence maximum jitter (one-way transmission through the LAN infrastructure of a single campus).
5	**"Video" or "Audio"** Characterized by less than 100ms delay.
4	**Controlled Load** Important business applications subject to some form of "admission control," be that pre-planning of the network requirement at one extreme to bandwidth reservation per flow at the time the flow is started at the other.
3	**Excellent Effort** Or "CEO's best effort," the best-effort type services that an information services organization would deliver to its most important customers.
2	**Best Effort** LAN traffic, as we know it today.
1	**Background** Bulk transfers and other activities that are permitted on the network that should not impact the use of the network by other users and applications.

The mapping of the priority to traffic classes (number of transmit queues) is given as follows:

801.2D (1998) – Section H.2.4

		Number of Available Traffic Classes							
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
User Priority	**0 (Default)**	0	0	0	1	1	1	1	2
	1	0	0	0	0	0	0	0	0
	2	0	0	0	0	0	0	0	1
	3	0	0	0	1	1	2	2	3
	4	0	1	1	2	2	3	3	4
	5	0	1	1	2	3	4	4	5
	6	0	1	2	3	4	5	5	6
	7	0	1	2	3	4	5	6	7

VLAN in WAN/MAN (Provider VLAN)

Provider VLAN, (also known as Q-in-Q, double tagging and VLAN stacking) is a technology that allows carriers to add on a public VLAN tag onto packets, so that they can be transported over a metro network. The important thing about the technology is that it maintains the integrity of the traffic, and keeps one customer's traffic separated from another customer's traffic. The IEEE802.1ad "Provider Bridges" working group is working to fully define the Q-in-Q technology. Moving forward, 802.1ad will be a component of Virtual Private LAN Service (VPLS).

With stacked VLANs, customer's traffic is tagged with a service provider VLAN tag or P-VLAN Tag at the provider edge. The P-VLAN Tag is added after the Ethernet Source MAC address. The P-VLAN Tag includes a Provider VLAN ID (P-VID). The Provider Ethertype (P-Ethertype) often uses a value other than 8100h, indicating that this PVLAN Tag is not a standard IEEE 802.1Q VLAN Tag. The P-CFI is also set to zero for Ethernet. The subscriber's VLAN Tag (C-VLAN Tag) remains intact and is not altered by the service provider's anywhere within the provider's network.

For example, suppose a subscriber wants to use VLAN IDs 1, 2, and 10 over the network to another subscriber location. The provider could then assign P-VID 10 for this service. VLAN IDs 1, 2, and 10 would be mapped to P-VID 10. The VID 10 does not conflict with P-VID 10. As a result, subscribers are free to assign VLAN IDs and VLAN CoS values and extend the VLAN over the WAN. The customer can now integrate the networks at the remote sites as one single transparent network.

Provider VLAN is best suited for Internet Access Services and WAN connectivity, which typically uses a router as the subscriber CPE device. Details on VPLS, Bridge Protocol Data Unit (BPDU), and tunneling implementation are beyond the scope of this paper.

Security and Segregation

The ability of VLANs to create "boundaries" can also satisfy some of the security requirements that require LAN segregation. For example, in the MTU/MDU applications where the subscriber needs privacy in his network access, VLAN can be used as a logical firewall to prevent intrusion. The only broadcast traffic on this single-user segment would be from that user's VLAN (that is, traffic intended

for that user). It would be impossible to ease drop to any communication in that private LAN segment, because the traffic can not physically cross the VLAN boundary. In addition, the service provider can provision bandwidth, access and set priority through the use VLAN. This is especially important when VLANs are implemented in conjunction with private port switching in a public network.

A very similar requirement is needed in applications like a Web-hosting data center. Many independent customers share the servers in the data server. It would be prudent to segregate the traffic of each customer's server to prevent unlawful intrusion, virus infection, or attack launching from one server to the neighboring server.

Conclusion

With the deployment large numbers of switch ports, VLAN has become an indispensable tool for the network administration to segment the network to increase bandwidth per user, provide security, and provision multimedia service. The evolution of VLAN as a simple broadcast containment device to a necessary function in the network, propel VLAN to be the number 1 tool in an IT professional's bag of tricks. New standards involving VLAN are being created today. New function like Q-in Q and STP per VLAN should prove to be very helpful in IP VPN, VPLS and metro area network.

Micrel switch family (KS8993, KS8995 and KS8695) is fully compliant with IEEE802.1p and 802.1q. It provides the most comprehensive VLAN and priority-based QoS solutions today in the Ethernet switch market.