

Two Models For VLAN Tagging

Norman Finn **Cisco Systems**

Paul Frantz **Bay Networks**

John Wakerly **FORE Systems**

1.0 Introduction

Many of the issues regarding VLANs can be understood better if we first understand two different models for the operation of explicitly-tagged VLAN: the one-level model and the two-level model. Most of the confusion over hidden assumptions, trade-offs, and tagging formats are driven by the differences between these two models.

In the one-level explicitly-tagged model, the MAC addresses and VLAN associations of all endstations are visible to every packet switch through which a packet flows. To an untagged packet must be added at least a distinguishing marker identifying it as explicitly tagged (an Ethertype or LLC value) and a VLAN-ID.

In the two-level explicitly-tagged model, a tagged packet has two complete MAC layers. The original packet as transmitted by the endstation is enclosed in a wrapper consisting of at least a MAC header, a distinguishing marker (Ethertype or LLC), and a VLAN-ID. In this model, the outer MAC addresses are not endstation addresses, but addresses of packet switches and/or multi-cast addresses.

The two models have different operational characteristics, and thus different features and drawbacks. Either one or both may be required of an IEEE VLAN standard.

The purpose of this contribution is not to suggest specific formats for use in 802.1 VLANs. Its purpose is informational: to demonstrate that there are, in fact, two different models of VLAN operations underlying the two different styles of tagging presented.

2.0 Common Considerations

Certain issues are common to both models:

1. For any given LAN segment, there can be only one format for all packets belonging to a given VLAN: they must either all be transmitted in untagged format on that LAN segment, or they must all be transmitted with an explicit tag on that LAN segment. This restriction guarantees that existing endstations can send and receive packets, while ensuring that packet switches need not duplicate packets in different formats on a single LAN segment.

Note: *This may not turn out to be a hard-and-fast rule. One can imagine a need to accommodate both tagged-only switches and untagged-only endstations on the same LAN segment, and can imagine ways to accomplish this by violating the above rule. However, this rule greatly simplifies many of the issues with multiple encapsulations, translation bridging, and bridge loops. It is at least a useful simplifying assumption; exceptions for corner-case interoperability can be discussed at a later time.*

2. “Edge switches” will translate between tagged and untagged formats when transferring a packet between LAN segments which, for a given VLAN, the packet has different formats. This contribution does not specify how untagged packets are to be classified into one VLAN or another.
3. “Fabric switches” will transfer tagged packets from one LAN segment to another. The assignment of a tagged packet to a VLAN must be made solely on the basis of the VLAN-ID in the packet tag.
4. Any given packet switch may display both Edge and Fabric behaviors in different cases, when forwarding packets between different LAN segments.
5. The MAC layer of the LAN segment over which a tagged packet is transmitted is used unchanged. Tagging does not invent a new MAC layer. Thus, the tagging format is somewhat different between Ethernet, 802.3, FDDI, and Token Ring LANs.

3.0 One-Level Model

The Ethernet tagging format for the one-level model is:

6	6	2	2		4
Dest addr	Source addr	Ethertype = X	VLAN-ID	Remainder of original frame	FCS

The Source and Destination addresses from the original, untagged, frame have been split apart from the rest of the frame, two fields have been inserted, and a new FCS computed. The two added fields are an Ethertype to identify this as a tagged packet, and a VLAN-ID identifying the VLAN to which the packet is assigned.

In order to support the following potential VLAN features, a fabric switch must pay attention to the Ethertype and VLAN-ID fields of a tagged packet when making its forwarding decisions. If none of these features is required, then a fabric switch may handle all packets in the same way, regardless of VLAN membership.

1. Separation of duplicate MAC addresses: If two different MAC interfaces on two different VLANs can have the same MAC address, then the VLAN-ID is required to distinguish them from each other and forward packets to the correct destination.
2. Multiple spanning trees: If two different VLANs have two separate spanning trees, then a given port on a packet switch may be in the “blocked” state for one VLAN and in the “forwarding” state for another VLAN. The VLAN-ID of each packet must then be examined in order to handle the packet properly.

If neither of these considerations is relevant to a given installation, then a fabric switch may operate in the same manner as an 802.1D bridge.

If separation of duplicate MAC addresses is required, then for the purposes of learning and forwarding packets, the fabric bridge must utilize the VLAN-ID information. For the purposes of learning addresses and forwarding frames, the fabric bridge must use the tuple, {VLAN-ID, MAC address}, in place of the MAC address alone. With this enhancement, a fabric switch may otherwise operate in the same manner as an 802.1D bridge.

The issues with multiple spanning trees are not addressed in this document.

4.0 Two-Level Model

The Ethernet tagging format for the two-level model is:

6	6	2	2	4	Entire original frame with FCS
VDest	VSource	Ethertype = Y	VLAN-ID	CRC fix	

The VDest and VSource are entirely new addresses added to the frame. The Ethertype must be a value different from the one-level model's Ethertype. The CRC fix is described in [1].

The original frame may have a different MAC type (e.g. FDDI or Token Ring) from the MAC type of the tagged frame. A tagged packet is therefore not considered to belong to any one VLAN, but to the "Fabric LAN". Packet switches have MAC interfaces to the Fabric LAN. The MAC address space of the Fabric LAN is distinct from the MAC address space of any VLAN.

4.1 Unicast Frame Forwarding

Let us assume for a moment that one universal spanning tree is used to open any physical forwarding loops in the network topology. The forwarding rules for a VLAN packet switch are then:

1. If the packet is tagged (arrived on the Fabric LAN), examine the VDest. If that VDest is not any unicast MAC address belonging to this packet switch, or is a multicast MAC address identifying a multicast group to which this packet switch does not belong, then the packet is forwarded on the Fabric LAN based on its VDest and VSource addresses. In this case, the Destination and Source MAC addresses of the original packet may be safely ignored by the packet switch. They may be examined or not, and source MAC address information learned or not, according to implementation/configuration trade-offs between memory, processing, bandwidth, etc.
2. If is tagged, and either the VDest is a unicast MAC address belonging to this packet switch, or is a multicast MAC address identifying a multicast group to which this packet switch belongs, then the packet is assigned the VLAN-ID from the tag, and the packet's Incoming Tag is the tag received from the Fabric LAN.
3. If the packet is untagged, classify (by means not specified in this document) it and assign it a VLAN-ID. The packet's Incoming Tag is set to VDest == VSource == the packet switch's own MAC address on the Fabric LAN. As far as the two-level model is concerned, packets tagged with the one-level tagging scheme are reassembled and treated like untagged packets, using the VLAN-ID from the one-level tag (or some algorithmic mapping between VLAN-ID spaces).
4. For the purposes of MAC address learning, the source tuple, {VLAN-ID, Source MAC address, VSource} must be formed and learned. (VSource is equal to the packet switch's own Fabric LAN MAC address for packets received without tags.)
5. The tuple {VLAN-ID, Destination MAC address} is used as the key into the packet switch's forwarding table. From the forwarding table, the outgoing port and the learned VSource are obtained.
6. When forwarding a packet on a LAN segment on which the packet's VLAN is not tagged, the Incoming Tag is ignored, and the packet is transmitted without a tag.

7. When forwarding a packet on a LAN segment which is tagged, the learned VSource is used as the VDest of the packet. The packet switch's own MAC address on the Fabric LAN is used as the VSource.
8. If a tagged packet's VDest is a unicast MAC address belonging to this packet switch, or is a multicast MAC address identifying a multicast group to which this packet switch belongs, then that packet must not be transmitted back to the fabric LAN on which it was received. Otherwise, inconsistencies in different switches' bridge tables can cause forwarding loops.

4.2 Multicast Frame Forwarding

Multicast frames follow the same general rules as unicast frames, with the following alterations:

1. For each multicast group on each VLAN, a multicast group may be created and a multicast MAC address assigned on the Fabric VLAN. These are called "specific multicast" groups and addresses. Specific multicast groups are not discussed further in this document.
2. A specific multicast MAC address is assigned on the Fabric LAN for each VLAN to be used for all broadcasts and non-specific multicast frames on that VLAN. This is the "VLAN-specific multicast group" for that VLAN.
3. If a given VLAN is used in the untagged (or not two-level tagged) form on any port on a packet switch, then that packet switch must belong to that VLAN's VLAN-specific multicast group.
4. A multicast frame may need to be forwarded in different formats on different ports, according to whether its VLAN's format is the same or different on the different ports.
5. A tagged frame with a VDest to which a packet switch does not belong is forwarded along the Fabric LAN.

5.0 Issues

1. Perhaps any given port either *is* or *is not* a Fabric Port and thus a part of the Fabric LAN. This is independent of whether a given VLAN's form is tagged or untagged on that port. I think this clarifies the multicast situation.
2. Multiple spanning trees still are not folded into this plan.

6.0 References

- [1] John Wakerly, "An Efficient Frame-Tagging Format for VLANs", presentation at 802.1 interim meeting, January 24-25, 1996.