Application Note

VLAN Aggregation with the
Model 6081RC EdgeRoute

Applies to the following product

- ForeFront Model 6081RC

Application Overview

This application note discusses the network design to aggregate VLANs in the Model 6081RC EdgeRoute card for transport to the local MPLS/Access router. One design goal is to maintain a simplified topology while keeping each VLAN's tag intact without compromising data throughput on the incoming channels. Each channel connects a remote site to the 6081 via a BCP link.

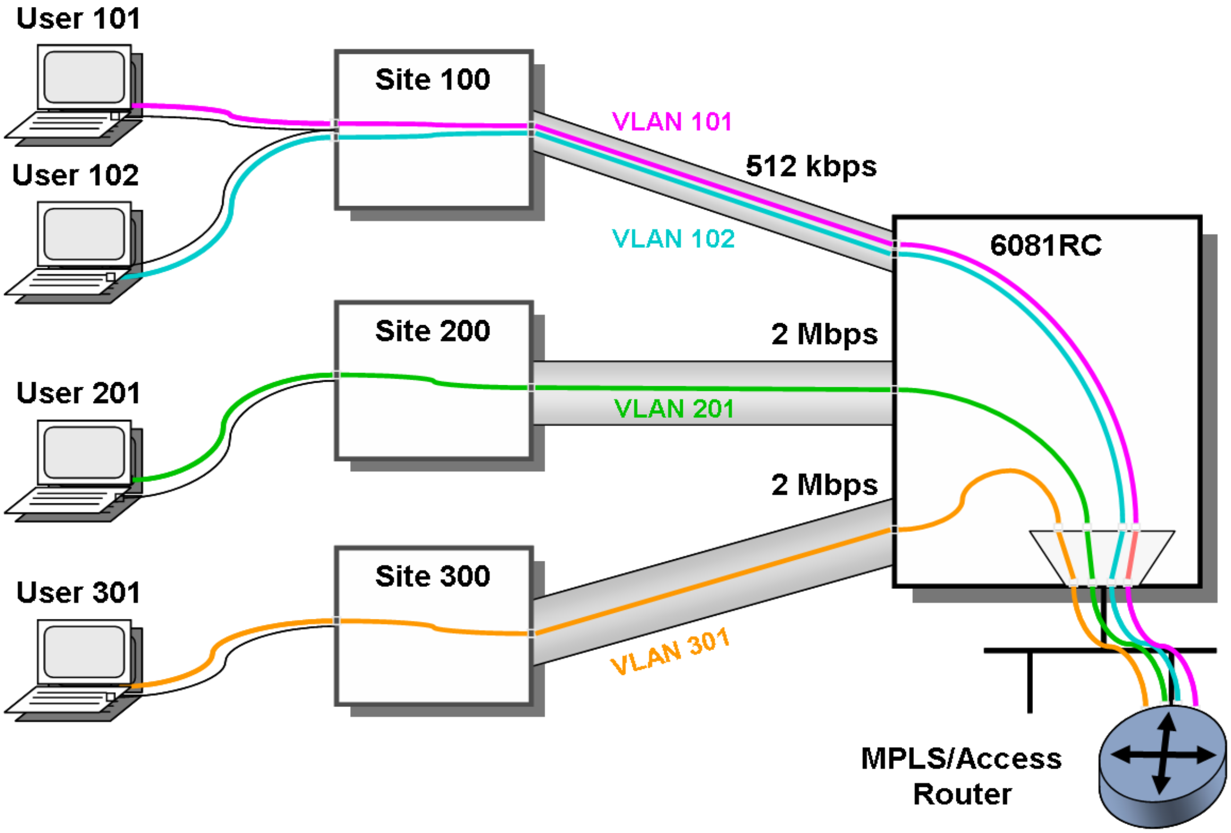

Figure 1. High level network design goal

Referring to figure 1, there are three channels for connecting the remote sites to the 6081. Two sites, with a bandwidth of 2 Mbps, each carry VLAN traffic, one VLAN per channel. The other site, with bandwidth of 512 kbps, carries two VLANs. The four VLANs partition each user's traffic end-to-end. All the VLANs are aggregated in the 6081 and forwarded through the Ethernet port to the local MPLS/Access router. From the local router, the VLAN traffic is routed to its final destination. (The final destination is not shown on figure 1.)

The VLANs maintain separation of each user's traffic. One intended consequence is to prevent one VLAN's traffic from being transmitted on another channel. The 512 kbps pipe carries two VLANs, one for

User 101, the other for User 102. Users 201 and 301 each have exclusive rights to its own 2 Mbps channel. Ideally no traffic for VLAN 201 should occur on the channels for Sites 100 and 300, only on Site 200's channel. Similarly, no traffic for VLANs 101 and 102 should occur on the channels for Sites 200 and 300. The only exception is an "arp" broadcast. Broadcast traffic is forwarded on all the VLANs. Normal unicast traffic remains on its designated VLAN. This provides the maximum bandwidth proper traffic.

Simpler but poorer

"The art of simplicity is a puzzle of complexity." (Doug Horton)

Simplicity lends itself to elegance, but oversimplification may conceal hidden demons. The solution in figure 2 appears simple, but does not meet all the design goals discussed in the "Application Overview".

Note Refer to the Appendix for the definitions of the various symbols used in these diagrams.

Note figure 2 and figure 3 are simplified by eliminating the third VLAN (User/VLAN 301). It will be re-introduced in the section with the actual 6081 configuration.

First consider some basic details in figure 2's solution. The three key components are the VLANs, the bridge group, and the BCP links.

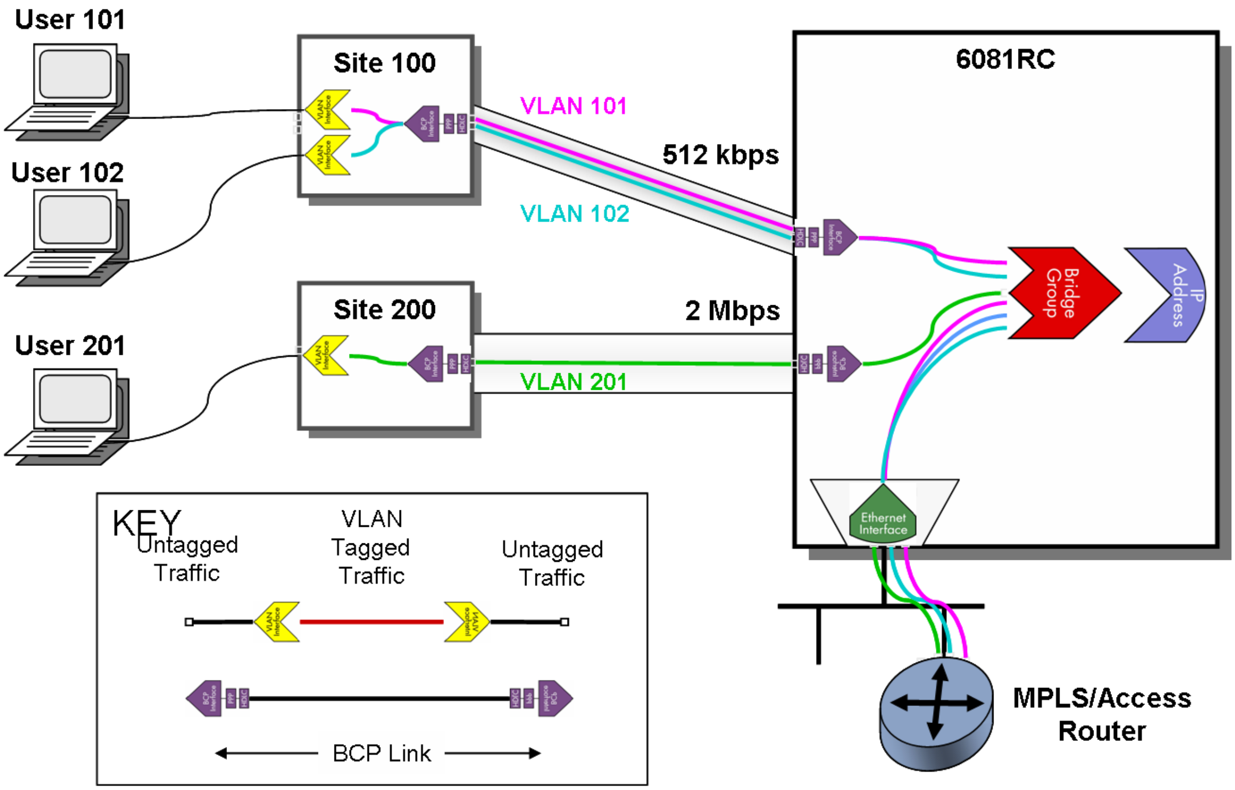

Figure 2. Simpler solution with degraded throughput

Each channel between a remote site and the 6081 uses BCP as the link-layer protocol. This is a bridged protocol, so all the VLANs are connected to the bridge group where frames are forwarded between the channels and the Ethernet port. Let’s examine the behavior of broadcast and unicast traffic.

Broadcast traffic, such as “arps”, clearly will be forwarded on all the VLANs, meaning that an “arp” originated from User 101 is forwarded on all VLAN segments and all non-VLAN segments. Therefore the “arp” occurs on both the 512 kbps and 2 Mbps channels in figure 2.

Unicast traffic is slightly different. If unicast traffic originates from User 102, it occurs on the following links:

- the non-tagged frames from User 102 to Site 100,
- VLAN 102 since this is the VLAN associated with User 102,

But since the 6081’s bridge group (the red symbol) does not filter VLANs, the unicast frame is forwarded on all VLANs. Consequently the frame is forwarded also to the 2 Mbps channel. We do not want this traffic on the 2 Mbps channel, so it is stealing bandwidth from legitimate traffic on VLAN 201. This can be very troublesome when the 2 Mbps channel is running near capacity because all its unicast traffic be on the 512 kbps channel. Legitimate VLAN 101 and VLAN 102 unicast traffic is competing with traffic from VLAN 201. The 512 kbps channel becomes a bottleneckl; it is being choked, and there can be severe loss of data. Considering they may be retransmission schemes (for example, with TCP sessions), the traffic load is increased and throughput suffers even more.

In summary, this appears to be a simple implementation. There is only one bridge to configure. But the hidden demon is bandwidth stealing.

Next let’s consider an implementation that *appears* more complex but is superior in terms of network behavior and performance.

More optimal implementation

This implementation is configured in two basic steps. Firstly, create a separate bridge group in the 6081 for each VLAN. Secondly, break the VLAN into two segments. One segment exists between the remote site and the associated bridge group. The second segment is between the same associated bridge group and the other termination of the VLAN. (The termination is not shown in figure 3.) So the bridge group for each VLAN has two VLAN endpoints associated with it. Thes two steps precludes bandwidth stealing for unicast VLAN traffic.

Detailed frame analysis

In figure 4 unicast traffic is shown for three of the VLANs in the improved implementation. Refer to the Symbol Key in figure 4. Notice the difference in the symbols for raw untagged Ethernet frames (which carry the IP unicast packets) and the VLAN tagged frames. The diagram shows the location of untagged and VLAN tagged frames for all three of the VLANs to understand why bandwidth stealing is precluded for unicast traffic.

Let’s follow the case of traffic originating from User 101. The untagged Ethernet frame from User 101 enters Site 100 where it acquires the VLAN tag “101.” After traveling over the 512 kbps channel to the 6081, it arrives at its associated bridge group which removes the VLAN tag. (See the “User101” symbol inside the bridge group symbol.) The bridge group reassigns the VLAN tag “101” before forwarding it to the Ethernet interface. There is nothing special here.

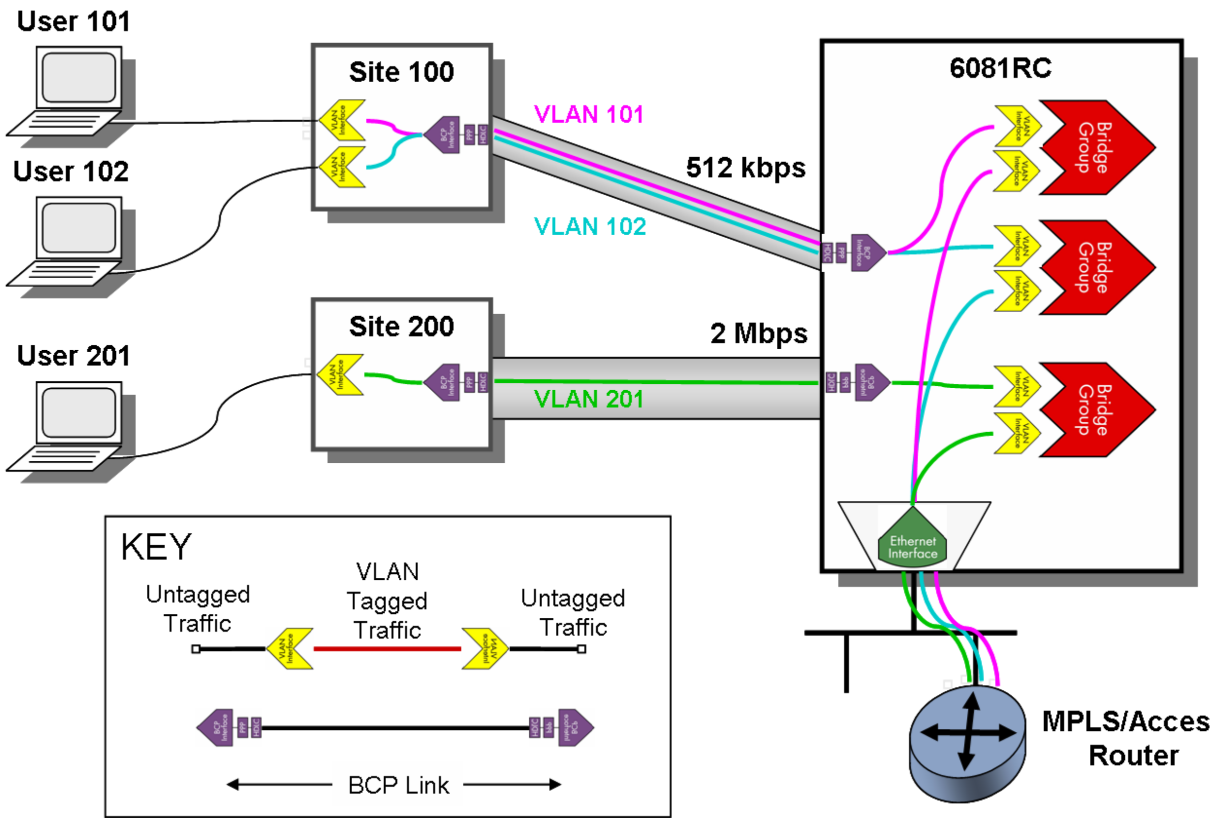

Figure 3. VLAN aggregation with proper VLAN partitioning

However a tagged VLAN frame (for example, VLAN 101) arriving from the MPLS/Access router would be on the same path to all three bridge groups. Only the VLAN termination on the associated bridge group untags it so the bridge group can forward it. The other VLAN terminations (that is, those for VLANs 102 and 201) reject the frame because there is not a match with the VLAN tag. It does not get forwarded into that bridge. Consequently all unicast traffic will travel only on the channel between the 6081 and remote site which carries that specific VLAN. Bandwidth stealing is eliminated and data throughput is maximized.

Advanced (Optional) - Adding IP addresses to bridge groups

The assignment of IP addresses can be difficult without understanding how bridge groups control the flow of data.

Note We will assign a term to each side of an interface. Refer to figure 5 to identify the Source and Sink sides of an interface. Source and sink do *not* imply the direction of the data flow.

Understanding this topic is crucial if you are going to connect one of the Ethernet interfaces to a bridge group. Refer to figure 6 for the subsequent discussion.

- Case (a): In figure 6 (a), an IP address is assigned to the Ethernet interface. By connecting a PC to the Ethernet interface on the same subnet, you can ping the IP address from the PC. Notice that you can connect other interfaces to the Ethernet interface. The traffic will also flow to and from the VLAN interface.

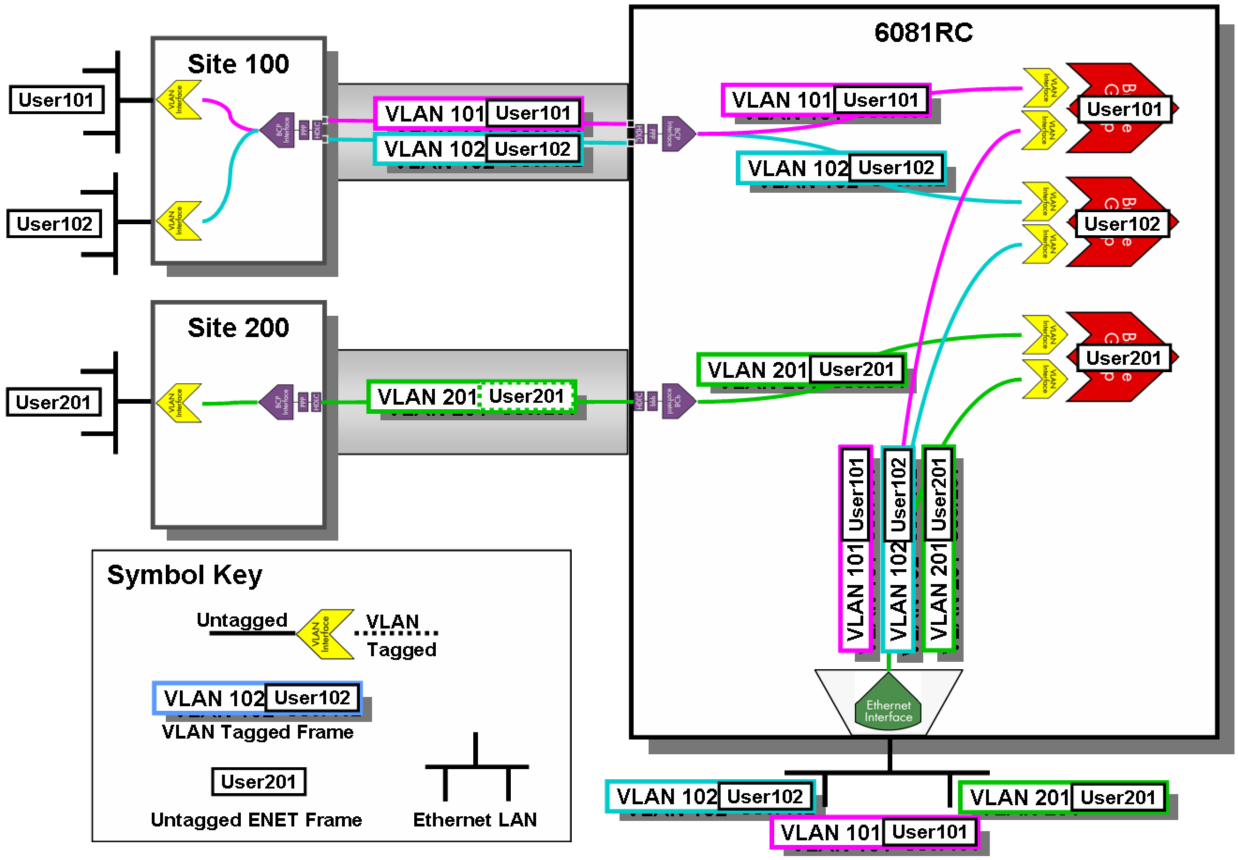

Figure 4. Detailed analysis - Unicast traffic

- Case (b): In figure 6 (b), the example graphically shows that multiple interfaces can be connected to the Sink side of the bridge group.
- Case (c): In figure 6 (c), the Ethernet interface connects directly to the Sink side of the bridge group. Although multiple interfaces can connect to the bridge group (as in Case (c)), the interface connected to the bridge group *cannot* connect to any other interface. In this case, the Ethernet interface cannot connect to any other interface. The reason is that the bridge group attracts all traffic from all of the interfaces connected to it. Continue to Case (d) to understand the consequences.
- Case (d): In this instance, the Ethernet interface connects to both the bridge group and the IP address. Due to the behavior of the bridge group (Sink side), all Ethernet interface traffic goes to the bridge group. None of the traffic arrives at the IP address. So if we ping the IP address still attached to the Ethernet, we will not see the ping response.

  Here is a metaphorical way to describe this behavior. It applies only to the Bridge Group. The Sink side of the bridge group relentlessly demands all data from the connected interface. The connected interface is not permitted to send data to both a bridge group's sink and another interface (figure 6 (d)). Even if multiple interfaces are connected to a bridge group's sink side, all data from the multiple interfaces must go to the bridge group. Therefore the IP address in (d) cannot respond to a ping, since it never receives the ping request from the Ethernet interface.

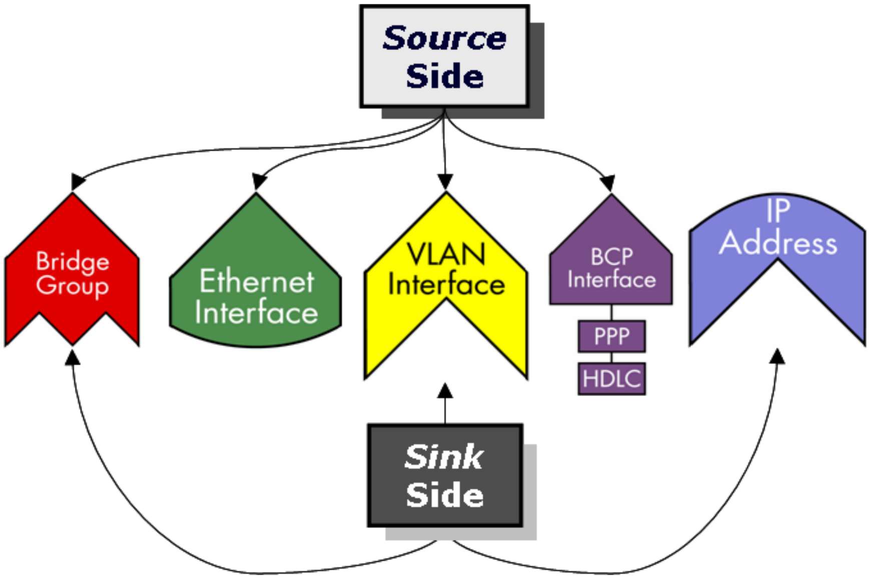

Figure 5. Source and Sink sides of an interface

To repeat, only the bridge group interface has this personality. No other interface “sinks” all data from an attached interface.

- Case (e): The solution is to assign another IP address in the same subnet and attach it to the bridge group. As soon as you enable the bridge group connected to the Ethernet interface, you can ping the IP address with the bridge group, whereas the previous IP address attached to the Ethernet interface disappears, that is, it does not function.

Even though our example considers the most crucial case in considering the Ethernet interface with the bridge group, this applies whenever an interface is connected to the bridge group’s Sink side.

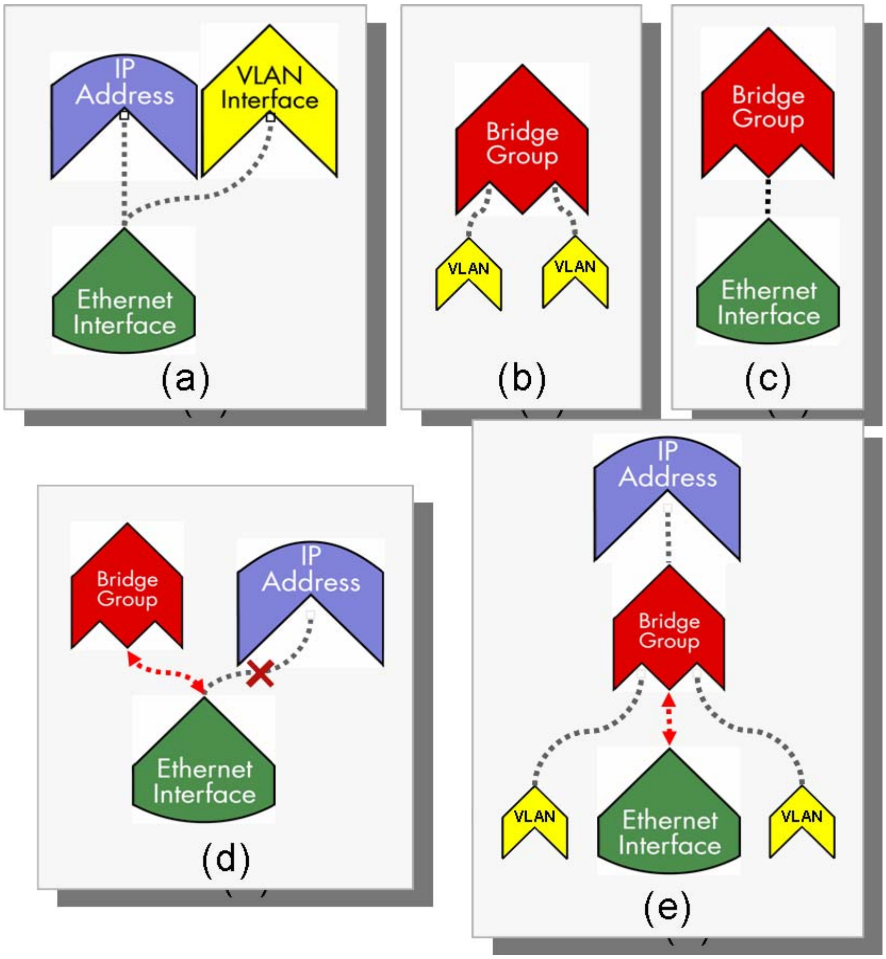

Figure 6. Interaction of data flow between IP addresses and bridge groups

Appendix

Refer to the Symbol Key (figure 7) to understand better the diagrams in the body of this Application Note. Although this is not a tutorial with the complete set of rules for combining these symbols, here are a few of the rules as a guideline.

1. The pointed convex side (Source side) of a symbol can connect to a pointed concave side (Sink side) of another symbol.

2. Bridge group:

– multiple symbols or multiple instances of the same symbol can connect to the Sink side of the bridge group.

– whenever a symbol connects to the Sink side of the bridge group, the Source side of that symbol cannot connect to to any other symbol.

3. Ethernet Interface: the rounded side is the physical layer interface of the Ethernet, e.g., a 10BaseT twisted pair cable.

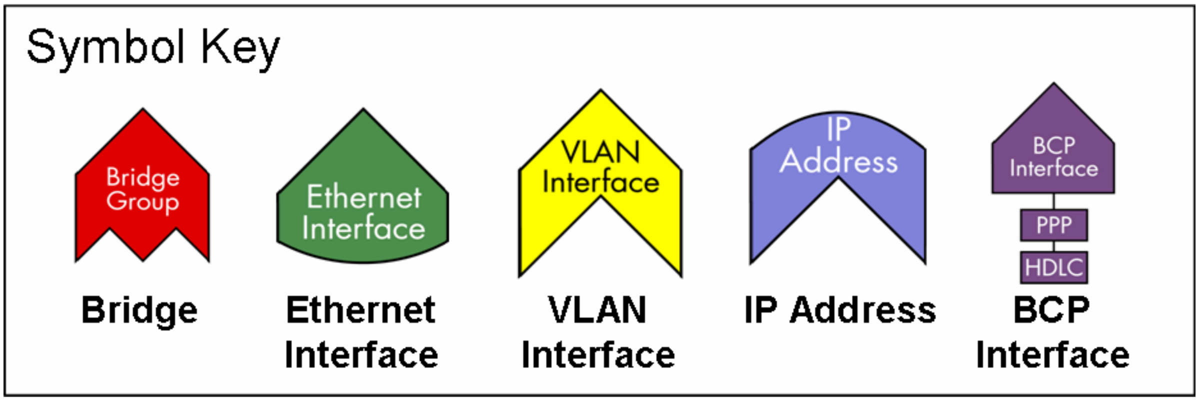

Figure 7. Symbol Key

4. VLAN:

– The Source side of the VLAN symbol is non-tagged traffic.

– The Sink side transmits and receives VLAN-tagged traffic.

5. IP Address: This symbol makes a connection only on the Sink side.

6. BCP Interface:

– This interface terminates one end of a BCP link.

– The Source side of this interface connects to the Sink side of any other symbol.

– The rectangle denoted as "HDLC" is a physical layer connection, e.g., DSL, E1, etc.

First Created

May 23, 2006

Last Updated
June 6, 2006 1:08 pm

PATTON
Electronics Co.

7622 Rickenbacker Drive
Gaithersburg, MD 20879
Tel: +1 301.975.1000
Fax: +1 301.869.9293