Handout #
C. Noronha
Winter 98-99

EE384A: Network Protocols and Standards
Homework #3 - Solutions
IEEE 802.1Q/D9 VLANs

1) a. Describe what VLANs are.

A VLAN is logical group of stations that retains all the properties of a physical LAN (i.e. single broadcast domain) even though stations may be located anywhere in an extended LAN. Traffic in a given VLAN is restricted from going to another VLAN.

b. What does the concept of VLANs bring to the user from a functionality point of view?

i) *VLANs facilitate easy administration of logical groups of stations that can communication as if they were on the same LAN. They also facilitate easier administration of moves, adds and changes in members of these groups.*
ii) Traffic between VLANs is firewalled. This limits the propagation of multicast and broadcast traffic between VLANs.
iii) VLANs are supported over all 802 media, and over shared media as well as point-to-point LANs. This means that end stations on a Token Ring and be on the same VLAN as a station on an Ethernet.

c. Give three scenarios that explicitly illustrate the usefulness and advantages of VLANs.

a) *In one large corporation, engineers, sales and marketing people, and customer service are usually located in different buildings and have their own LANs. Using VLANs it is possible to create a LAN for a task group that includes employees from several departments without moving individual stations.*
b) *VLANs can be used to restrict access to some critical servers by specific users. To dedicate server A to a group GA of people, then the server and the group of stations are placed on the same VLAN that does not include others.*
c) *If there were users that are interested in video traffic, then you would prevent such traffic from reaching these users by placing them on different VLANs.*

2) What technological advances are making VLANs possible?
    - *Advanced switch architecture, larger and higher bandwidth switches*
    - *Introduction of high speed LANs.*
    - *Advances in VLSI, memory and software design*

3) What are the implications of VLANs on the use of GMRP for dynamic multicast filtering?

In the absence of VLANs, GMRP data units are propagated throughout the entire spanning tree: this is referred to as the Base Spanning Tree context. With VLANs, it is possible to allow GMRP registrations to be made that are specific to VLANs.
This is simply accomplished by:

- *Considering that there is an applicant and a registrar per VLAN, identified by the VID of the VLAN*
- *Tagging GMRP PDUs with the VID corresponding to the VLAN to which they apply*
- *Apply the same Ingress Rule to received GMRP PDUs as to VLAN tagged frames*
- *Apply the same Egress Rule to GMRP PDUs to be transmitted on a port as to VLAN tagged frames*

The main implication of the above is that the registration information is not allowed to reach outside the subtree corresponding to the VLAN. All VLAN members hear sources of multicast in that subtree. Sources outside the VLAN subtree however may or may not be heard by VLAN members depending on the default group-filtering behavior set at ports in the portion of the LAN topology outside the VLAN.

Question 4: IP Subnetting

An organization with 200 IP stations is given a class C address **X.Y.Z.-** for its network. The organization uses subnetting and organizes it network into multiple subnetworks as shown in Figure 1 below. The 200 stations are divided into three groups, one group consisting of 100 machines and the other two 50 machines each. Each group forms a single bridged LAN. A company backbone and a divisional backbone are also created as shown in the figure. It is required that all machines and all router interfaces be given IP addresses. You are asked to assign IP addresses to the stations and router interfaces, and to indicate the entries in each router's table needed to route IP datagrams destined to stations within the company network.

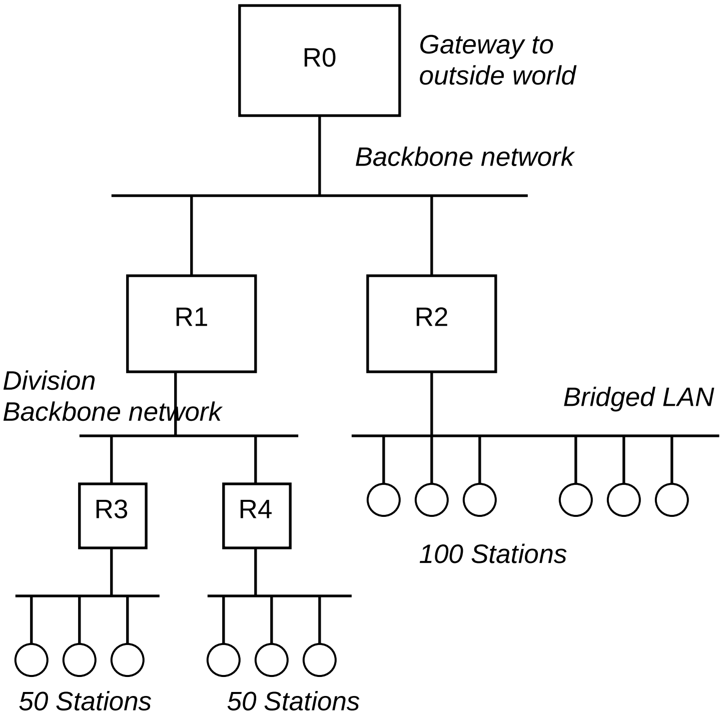

Figure 1

Suggested Solution:

In the class C address X.Y.Z.- we have 256 IP addresses, two of which are reserved (X.Y.Z.0 reserved to refer to the class C network itself and X.Y.Z.255 reserved for broadcast on this network). The remaining 254 addresses starting with X.Y.Z.1 and going all the way to X.Y.Z.254 are available for assignment. Below is a table which lists the subnetworks, the number of addresses needed by each subnetwork, the number of addresses assigned for each, and the range of the assigned addresses.

		Addresses	
Subnetwork	***Needed***	***Assigned***	***Range***
A. Backbone subnetwork	*3*	*7*	*X.Y.Z.1 to X.Y.Z.7*
B. LAN under R2	*101*	*120*	*X.Y.Z.8 to X.Y.Z.127*
C. Division subnetwork	*3*	*4*	*X.Y.Z.128 to X.Y.Z.131*
D. LAN under R3	*51*	*60*	*X.Y.Z.132 to X.Y.Z.191*
E. LAN under R4	*51*	*63*	*X.Y.Z.192 to X.Y.Z.254*

1. *IP address assignment:*

- *Subnet on backbone network:*

We chose to provide for a 7 IP address subnet on the backbone segment (the minimum is 3, one for each interface). This would allow the company to change its network topology more easily, if needed in the future.

Mask: 255.255.255.1111-1000 (the last byte is shown in binary for clarity).
Subnet number: X.Y.Z.0000-0000
R0's interface down: X.Y.Z.1
R1's interface up: X.Y.Z.2
R1's interface up: X.Y.Z.3
For later use: X.Y.Z.4 to X.Y.Z.7

Note that the special address X.Y.Z.0 was left unused.

Now, we have to divide the original IP address range we had minus the 8 addresses already assigned almost evenly into two parts, one for the 3 networks under R1 and one for the network under R2. The most natural way is to divide it by assigning the range 0000-0100 to 0111-1111 to one part and the range 1000-0000 to 1111-1110 to the other.

We chose to assign the first (which contains 120 addresses) to the network under R2 and the second (which contains 127 addresses) to the part under R1. Therefore, at the backbone segment level, the same mask (i.e. 255.255.255.1000-0000) is used to distinguish between the two ranges.

- *Subnet under R2:*

We have 120 addresses available and we need 100 for the stations and one for the router interface, therefore the assignment can be easily done as follows:

Mask: 255.255.255.1000-0000
Subnet number: X.Y.Z.0000-0000
R2's interface down: X.Y.Z.8
IP stations address range: X.Y.Z.9 to X.Y.Z.108
Unused: X.Y.Z.109 to X.Y.Z.127

- *Subnet on division backbone network:*

Here, we provided for 4 IP addresses as follows:

Mask: 255.255.255.1111-1100
Subnet number: X.Y.Z.1000-0000
R1's interface down: X.Y.Z.128
R3's interface up: X.Y.Z.129
R4's interface up: X.Y.Z.130
For later use: X.Y.Z.131

Note that these addresses should be in the range of addresses used by the two 50-station subnets, otherwise routing at the backbone is complicated.

- *Subnets under R3 and R4:*

We need 50 IP addresses on this LAN. In the range 128-255 we have already assigned the 4 IP addresses above and the IP address X.Y.Z.255 is reserved for broadcast. We can still divide this range into two parts that are both larger than 50, namely 60 and 63 valid addresses (a more even distribution is possible but it makes things much more complicated). Therefore:

- *Subnet under R3*
  Mask: 255.255.255.1100-0000
  Subnet number: X.Y.Z.1000-0000
  R3's interface down: X.Y.Z.132
  IP stations address range: X.Y.Z.133 to 182
  Unused: X.Y.Z.183-191

- *Subnet under R4*
  Mask: 255.255.255.1100-0000
  Subnet number: X.Y.Z.1100-0000
  R3's interface down: X.Y.Z.192
  IP stations address range: X.Y.Z.193 to X.Y.Z.242
  Unused: X.Y.Z.242 to X.Y.Z.254

2. Routing Tables:
Given this assignment, it is now straightforward to route all IP datagrams destined for addresses in the range X.Y.Z.1 to X.Y.Z.254. Below are the routing table entries that are needed in R0, R1, R2, R3 and R4's tables in order to achieve that. We assume that the router interfaces are numbered such that the network interface down is #1 and the one up is #2.

Note that the last entry in all routing tables is the default entry, which will be used to route packets generated inside the company network that are destined to hosts outside in the Internet. The statement of this problem requested the routing tables for intra-company packets only, so you did not need to have this in your solution to receive full credit. However, this entry is shown here for completeness; in a real-world situation, it is required.

Mask	Address	Next Hop
Router R0 Table		
255.255.255.1111-1000	X.Y.Z.0000-0000	Local, interface 1
255.255.255.1000-0000	X.Y.Z.1000-0000	X.Y.Z.2 // R1
255.255.255.1000-0000	X.Y.Z.0000-0000	X.Y.Z.3 // R2
0.0.0.0	0.0.0.0	ISP Router[1]
Router R1 Table		
255.255.255.1111-1100	X.Y.Z.1000-0000	Local, interface 1
255.255.255.1111-1000	X.Y.Z.0000-0000	Local, interface 2
255.255.255.1100-0000	X.Y.Z.1000-0000	X.Y.Z.129 // R3
255.255.255.1100-0000	X.Y.Z.1100-0000	X.Y.Z.130 // R4
255.255.255.1000-0000	X.Y.Z.0000-0000	X.Y.Z.3 // R2
0.0.0.0	0.0.0.0	X.Y.Z.1 // R0

[1] This would be interface 2 on router R0, which connects the company to the outside world. This interface is not shown in the picture.

Router R2 Table

255.255.255.1111-1000	X.Y.Z.0000-0000	Local, interface 2	
255.255.255.1000-0000	X.Y.Z.0000-0000	Local, interface 1	
255.255.255.1000-0000	X.Y.Z.1000-0000	X.Y.Z.2	// R1
0.0.0.0	0.0.0.0	X.Y.Z.1	// R0

Router R3 Table

255.255.255.1111-1100	X.Y.Z.1000-0000	Local, interface 2	
255.255.255.1100-0000	X.Y.Z.1000-0000	Local, interface 1	
255.255.255.1100-0000	X.Y.Z.1100-0000	X.Y.Z.130	// R4
255.255.255.1000-0000	X.Y.Z.0000-0000	X.Y.Z.128	// R1
0.0.0.0	0.0.0.0	X.Y.Z.128	// R1

Router R4 Table

255.255.255.1111-1100	X.Y.Z.1000-0000	Local, interface 2	
255.255.255.1100-0000	X.Y.Z.1100-0000	Local, interface 1	
255.255.255.1100-0000	X.Y.Z.1000-0000	X.Y.Z.129	// R3
255.255.255.1000-0000	X.Y.Z.0000-0000	X.Y.Z.128	// R1
0.0.0.0	0.0.0.0	X.Y.Z.128	// R1

Sometimes it is possible to have less specific entries (i.e. with shorter masks). For example, note that the last entry in both R3 and R4's tables can be replaced by: (however, it is a better practice to have it as above)

255.255.255.0000-0000	X.Y.Z.0000-0000	X.Y.Z.128	// R1

and this will be ok since routing is done on a longest matching prefix basis and the other three entries in both tables are more specific.