ANC4 Tutorial 3: Solutions

1. In 802.1D, a bridge will forward a frame on all ports except the one it arrives on if it does not have a filter table entry for the destination. In the diagram below if *A* cannot identify the destination address in a frame arriving from LAN 2 it forwards it on all outputs including that to LAN 1 and thus to bridge B. If B also has no table entry for the frame it will send it back to A (as well as on all other outputs) creating a loop.

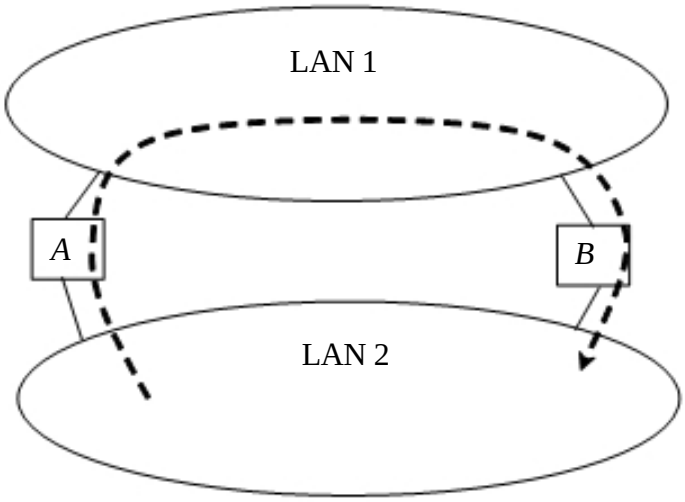

2. If all 24 stations are sending and receiving at maximum rate the switch must sustain 24 streams of 100Mbps requiring a bandwidth of 2.4Gbps.

3. A bridge may not be used if it is not included in the spanning tree (only one designated bridge is allowed for each LAN). An IP router has multiple IP addresses which can only be reached by routing to it. An unused bridge may be useful if a bridge currently in use fails and a new spanning tree includes it.

4. a) The most straightforward way to do this would be to allocate a different subnet to each LAN, using private addressing. There are numerous ways to do this. For example,

   10.0.0.0/24 for LAN A and 10.0.1.0/24 for LAN B.

   b) Suppose machine 10.0.0.1 sends to 10.0.1.1. The UDP segment will be packaged into an IP packet with destination address 10.0.1.1. The routing table on the source will identify this as a different subnet and will forward it to a tunnelling router, *X*, interfacing to the Internet. This router must tunnel the IP packet to its remote counterpart, *Y*, encapsulating it in another IP packet with the destination address of *Y*. At Y the original IP packet is extracted and routed according to its destination address 10.0.1.1.

   An alternative approach is to use IP tunnelling to carry Ethernet frames via remote bridging. In this case subnetting is not necessary.

   c) Alternatives would be to use: a dial-up network such as ISDN or ATM; a leased point-point line; a frame relay connection.

5. This can be done using overlapping VLANs. A port based VLAN normally allocates a fixedVLAN ID (VID) to a frame originating from a station attached to a given port (the port's *PVID*). This frame is then forwarded to all other ports with the same VID. However, when overlapping VLANs are supported, a port may be configured to receive frames with VIDs different from that port's PVID.

6. A station cannot join a VLAN until the switch is aware of which subnet it belongs to. The easiest way to do this is for a switch to wait until a station tries to send an IP packet. The switch must be configured with the appropriate subnet mask and can then immediately

allocate the station to its VLAN. Two such VLANs need an IP router attached to the switch to communicate.

7. Multicasting at the LAN level provides a simple mechanism whereby frames can be sent across a system of switched or bridged LANs to multiple destinations. Such multicasting is compatible with 802.1D and (after learning) frames are forwarded only to necessary branches of the spanning tree to reach the set of recipients in the multicast group. However, LAn multicasts cannot go further than the nearest IP router. IP multicasts on the other hand are intended for potentially much wider ranges across several subnets, networks or even AS's. IP multicasting uses LAN multicasting locally but requires specially aware IP mulicast routers for non-local transmission.

8. Multicasting is useful in routing protocols because it allows groups of routers operating the same routing protocol to exchange data as often required by such protocols.