1 True or False

1. **[2]** IP addresses can be divided into two parts: network prefix part and the host part.
   True

2. **[2]** ICMP Route Redirect packets update the routing table of hosts.
   False, ICMP Route Redirect packets update the routing cache.

3. **[2]** ARP can be used to detect and avoid duplicate IP addresses in the subnet.
   True

4. **[2]** In OSPF, all Area Border Routers are part of the backbone area.
   True

5. **[2]** BGP is a path vector protocol.
   True

6. **[2]** If we use the filter 'src host 10.0.1.12 and tcp dst port 80', tcpdump will show only HTTP packets that are from 10.0.1.12.
   True

7. **[2]** Cisco IOS commands are divided into hierarchical modes, and the mode with the least priviledge is the global configuration mode.
   False, user exec mode has the least priviledge.

2 The Network Administrator

Kahn is a network administrator for a small but fast growing company. He has configured the network as shown below. For all the questions below, assume that there are no physical connection problems and that all hosts have the correct default gateway configured.

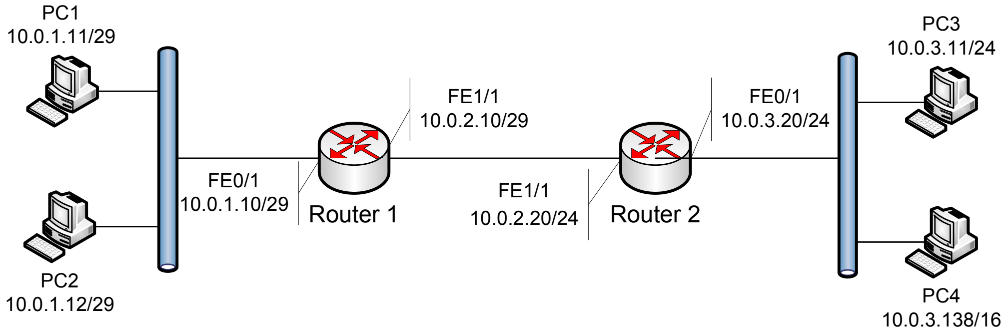

1. **[2]** How many subnets did he setup?
   3 subnets

2. **[3]** Kahn used static routes to configure the routing tables. What would the final routing table look like? Draw and fill two-column routing tables for Router 1 and Router 2.

Table 1: **Router 1**

Prefix	**Interface or address**
00001010 00000000 00000001 00001 (10.0.1.8/29)	**FE0/1**
00001010 00000000 00000010 00001 (10.0.2.8/29)	**FE1/1**
00001010 00000000 00000011 (10.0.3.0/24)	**10.0.2.20**

Table 2: **Router 2**

Prefix	**Interface or address**
00001010 00000000 00000011 (10.0.3.0/24)	**FE0/1**
00001010 00000000 00000010 (10.0.2.0/24)	**FE1/1**
00001010 00000000 00000001 00001 (10.0.1.8/29)	**10.0.2.10**

3. [**8**] Kahn pings from PC3 to PC1 but it's not working. He thinks that maybe there is something wrong with the addresses and netmasks he assigned to the hosts and routers. Is his hypothesis correct? If so, how should he fix it?
ICMP Echo Request packets have no trouble traveling from PC3 to PC1. The problem is when the ICMP Echo Reply packet travels from PC1 to PC3, specifically the path from Router1 to Router2. Since Router1's FE1/1 network prefix is 29 bits, the interface assumes that any address that starts with 00001010 00000000 00000010 00001 (29 bits) is within the same subnet. The first 29 bits of Router2's FE1/1 network address is 00001010 00000000 00000010 00010. The last two bits are different, so Router1 will not think Router2 is in the same subnet. Router1 also doesn't have any routing information on how to forward packets to this type of address. So the packet is discarded. To fix this problem, Kahn can reconfigure FE1/1's network prefix to 24 bits.

4. [**5**] Kahn notices that when PC4 pings Router 1's FE1/1 interface, PC4 first sends an ARP request packet asking for the MAC address of 10.0.2.10. He thinks that no one should answer since Router 1 is not in the same subnet. But to Kahn's surprise, Router 2 answers with an ARP reply packet with its own MAC address inside! What is going on? How can Kahn fix Router 2's behavior?
This happens when proxy ARP is configured. Kahn should disabled proxy ARP on Router2.

5. [**5**] The company is growing fast and Kahn needs to add another subnet to the network. The new subnet needs to support at most 50 hosts. Kahn wants to reduce unused IP addresses in the new subnet. What is the optimal length of network prefix that Kahn should assign to the new subnet?
To accomodate 50 hosts in the subnet, we need to have at least 6 bits in the host part of IP address since $2^6 = 64$. Anything shorter will not be able to accomodate 50 hosts. Anything longer will result in waste of addresses. So, the optimal prefix length is $32 - 6 = 26$ bits.

3 Count-To-Infinity

The count-to-infinity problem is a convergence problem that arises in distance vector protocols such as RIP.

1. [**5**] What is the main symptom of the count-to-infinity problem in terms of convergence time? Explain how the count-to-infinity problem affects the convergence time.
When there is count-to-infinity problem in the network, it takes much longer for routing tables to converge. This is because when count-to-infinity problem appears in the network due to a link failure, routers keep exchanging increasing distance metrics until the distance finally reaches 'infinity'.

2. [**5**] Why does count-to-infinity problem occur in distance vector algorithms?
   The problem occurs because distance vector algorithms exchange only the distance information with their immediate neighbours. Let's assume router X must send a packet to router Y to reach subnet S. When the link to subnet S breaks, router Y may calculate a new route to S via X since X has a better metric value to S. This happens because router Y thinks that X can forward a packet to subnet S. It does not know that router X is using router Y to get to subnet S. This forms a routing loop between router X and router Y. Router Y thinks it can reach subnet S via X. Router X thinks it can reach subnet S via Y. Since there is one hop between router X and router Y, they will keep exchanging increasing metric values until it reaches 'infinity' where both realize that the link is broken.

3. [**8**] Name two methods used to avoid the count-to-infinity problem and describe them.
   Potential answers are: triggered updates, split horizon, hold-down timer, and poisoned reverse.

4 File Transfer Using TCP

Vinton wants to send a 292 kilobyte picture from Host A to Host B using TCP. Both hosts are physically connected to the same Ethernet. The Maximum Transmission Unit (MTU) on the Ethernet is 1500 bytes. The IP header is 20 bytes and the TCP header is 20 bytes. Assume that there is no packet loss or corruption within this Ethernet.

1. [**2**] What is the Maximum Segment Size (MSS) of the hosts?
   MSS is the size of maximum TCP payload. $1500 - 20 - 20 = 1460$ **bytes.**

2. [**3**] How many data packets will Host A send to Host B?
   Each data packet will contain 1460 bytes of data. We need to send 292 kilobytes. $\frac{292*1024}{1460} = 204.8$**. So 205 packets will be sent.**

3. [**5**] Based on your answers above, fill in the sequence numbers and acknowledgement numbers in the ladder diagram.
   Answer is in the ladder diagram.

4. [**5**] Which phase of the congestion control algorithm is the ladder diagram showing? Explain your answer.
   The number of packets sent in one Round Trip Time increases exponentially: that is, the number of packets increase from 1 to 2 to 4. This is because congestion window is increased by 1 MSS whenever an ACK is received. This is the Slow Start phase.

5 Extra Credit

[**+2**] In the real world, Dr. Vinton Cerf and Dr. Robert Kahn are famous for inventing what?
TCP/IP

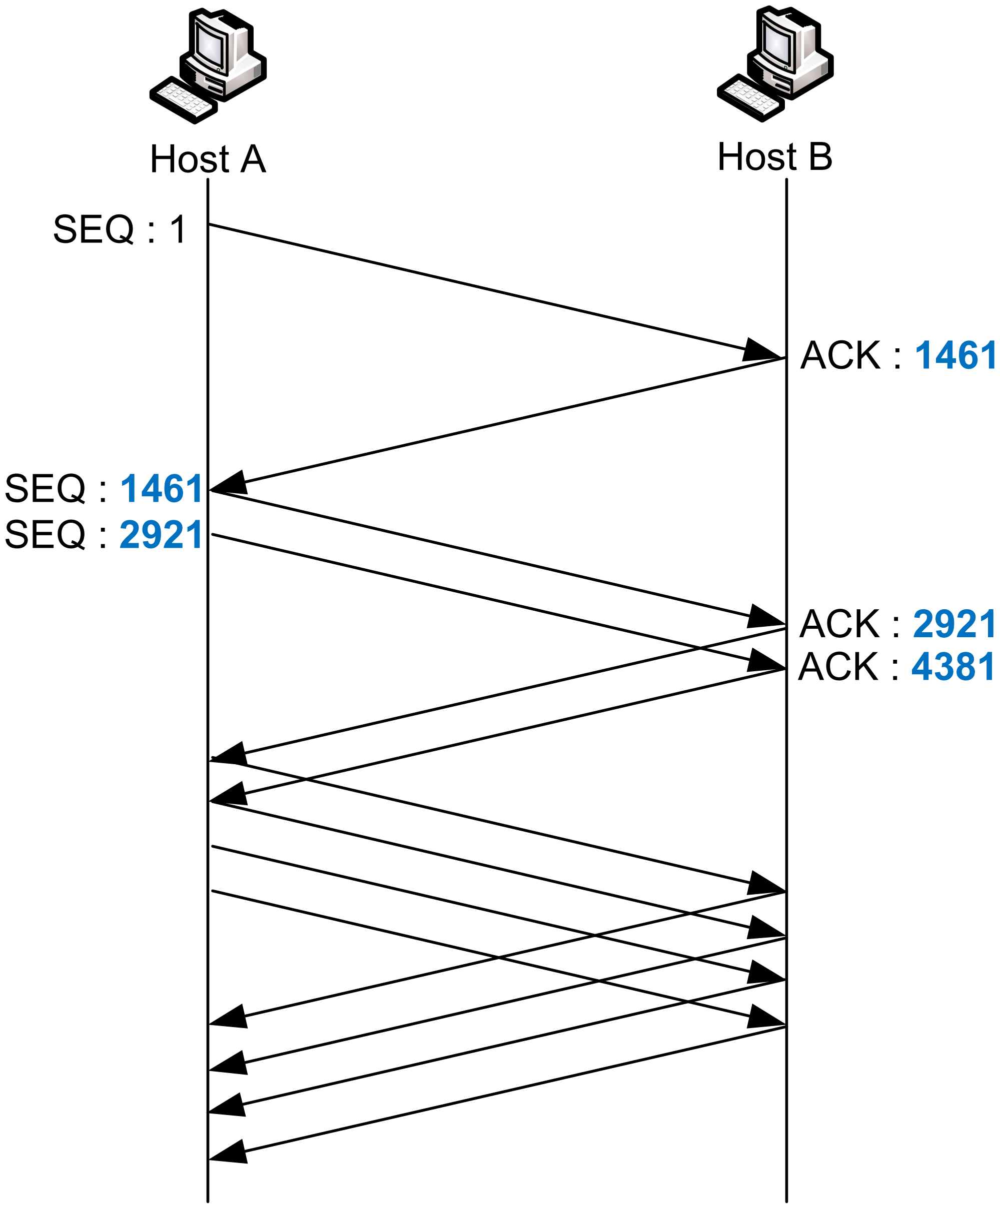
Host A
Host B
SEQ : 1
ACK : 1461
SEQ : 1461
SEQ : 2921
ACK : 2921
ACK : 4381