AHRAM CANADIAN UNIVERSITY

Lecture (08)
Internet Protocol (II)

<Internetwork @ tcp/ip> <Network @ OSI>

Dr. Ahmed M. ElShafee

Agenda

- Introduction
- Network Layer Interaction with the Data Link Layer
- Network Layer (Layer 3) Addressing
- IP Routing and Routing Protocols

Network Layer (Layer 3) Addressing (cont,..)

In the previous lecture we talked about ,…

- **IP addressing structure (layer 3)**
- **Routing protocols**
- **IP addresses grouping, classes and subnetting.**

Network Layer (Layer 3) Addressing (cont,..) (2)

Network Layer Utilities

The TCP/IP network layer uses several utility protocols to help it complete its task

1. Address Resolution Protocol (ARP) and the Domain Name System
2. ICMP Echo and the ping Command
3. RARP, BOOTP, and DHCP

Network Layer (Layer 3) Addressing (cont,..) (3)

1. Address Resolution Protocol and the Domain Name System

- TCP/IP needs a way to find MAC addresses associated with other computers on the same LAN subnet.

Hannah Knows Jessie's Name, Needs IP Address and MAC Address

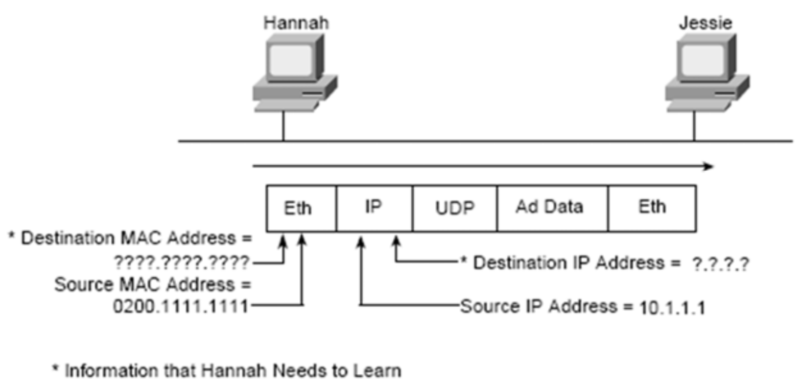

Network Layer (Layer 3) Addressing (cont,..) (4)

- Hannah knows her own name, IP address, and MAC address because those things are configured in advance.
- *What Hannah does not know are Jessie's IP and MAC addresses.*
- *To* find the two missing facts, Hannah uses the Domain Name System (DNS) and the Address Resolution Protocol (ARP).

DNS Request and Reply

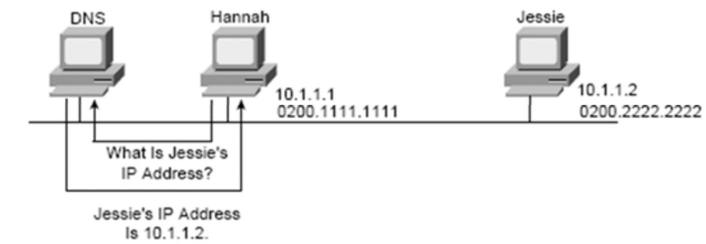

Network Layer (Layer 3) Addressing (cont,..) (5)

- Hannah knows the IP address of a DNS server because the address was preconfigured on Hannah's machine or learned using Dynamic Host Configuration Protocol (DHCP) (will be discussed in a later section)
- Hannah now sends a *DNS request to the* DNS, asking for Jessie's IP address.
- The DNS replies with the address, 10.1.1.2.

DNS Request and Reply

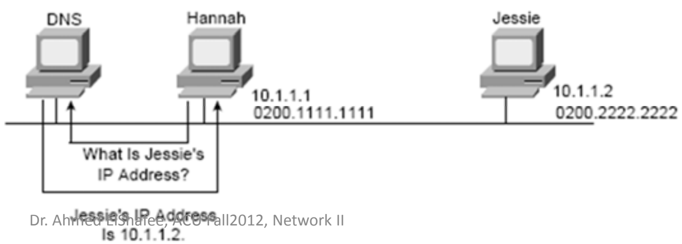

Network Layer (Layer 3) Addressing (cont,..) (6)

- Hannah still needs to know the Ethernet MAC address used by 10.1.1.2
- Hannah issues something called an *ARP broadcast.*
- An ARP broadcast is sent to a broadcast Ethernet address, so everyone on the LAN receives it.

Sample ARP Process

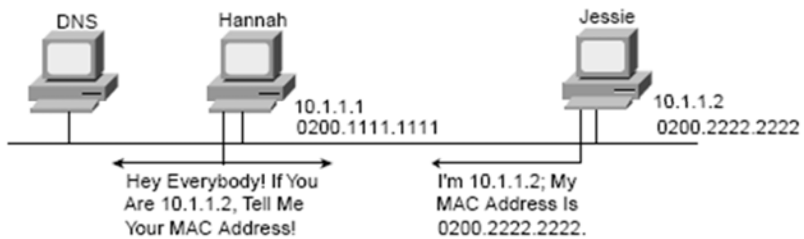

Network Layer (Layer 3) Addressing (cont,..) (7)

- Because Jessie is on the LAN, Jessie receives the ARP broadcast.
- Because Jessie's IP address is 10.1.1.2 and the ARP broadcast is looking for the MAC address associated with 10.1.1.2, Jessie replies with her own MAC address.

Sample ARP Process

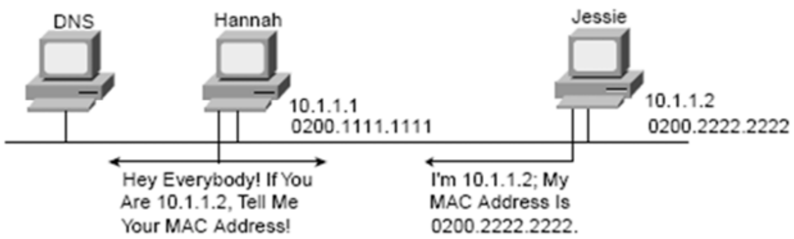

Network Layer (Layer 3) Addressing (cont,..) (8)

- Now Hannah knows the destination IP and Ethernet addresses that she should use when sending frames to Jessie,

Hannah Knows Jessie's Name, Needs IP Address and MAC Address

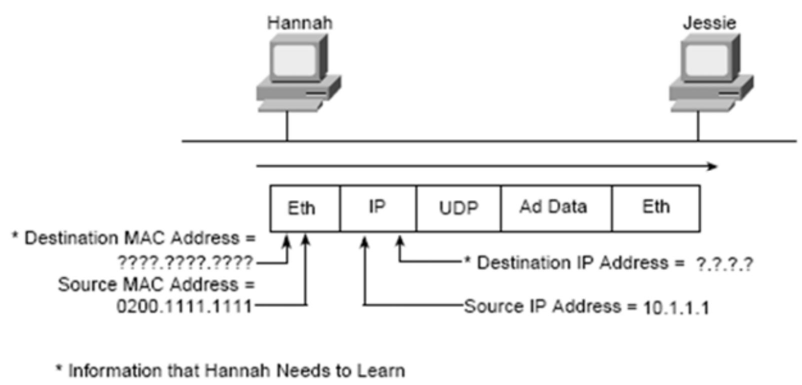

Network Layer (Layer 3) Addressing (cont,..) (9)

2. **ICMP Echo and the ping Command**

- IP needs to have a way to test basic IP connectivity, without relying on any applications to be working.
- can test basic network connectivity using the **ping command.**

Sample Network, ping Command

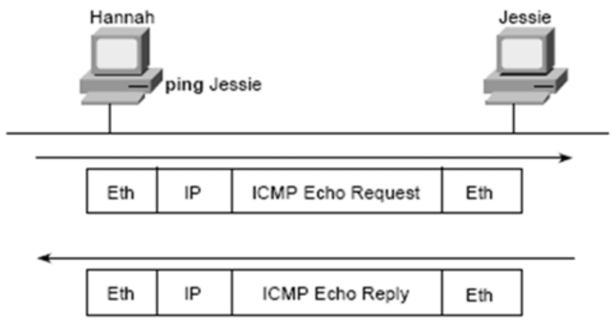

Network Layer (Layer 3) Addressing (cont,..) (10)

- **ping (Packet** INternet Groper) uses the *Internet Control Message Protocol (ICMP), sending a message* called an *ICMP echo request to another IP address.*
- The computer with that IP address should reply with an *ICMP echo reply.*

Sample Network, ping Command

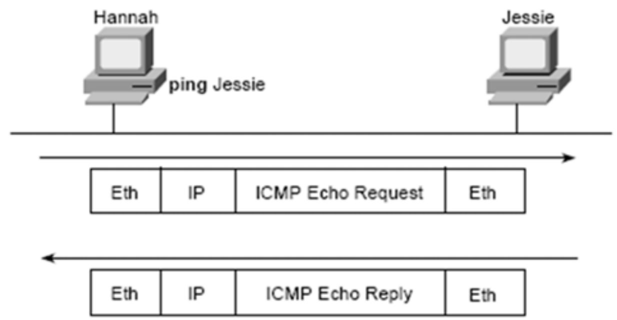

Network Layer (Layer 3) Addressing (cont,..) (11)

- *If that works, you successfully have tested the IP* network.
- ICMP does not rely on any application, so it really just tests basic IP connectivity— Layers 1, 2, and 3 of the OSI model.

Sample Network, ping Command

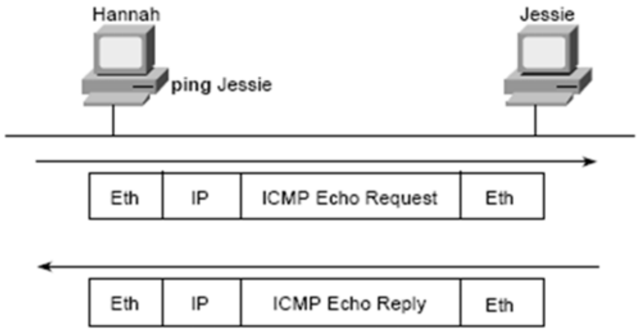

Network Layer (Layer 3) Addressing (cont,..) (12)

3. **RARP, BOOTP, and DHCP**

- Over the years, three protocols have been popular to allow a host computer to discover the IP address it should use:

1. Reverse ARP (RARP)
2. Boot Protocol (BOOTP)
3. Dynamic Host Configuration Protocol (DHCP)

Network Layer (Layer 3) Addressing (cont,..) (13)

- RARP and BOOTP work using the same basic process. To use either protocol, a PC needs a LAN interface card.
- The computer sends a LAN broadcast frame announcing its own MAC address and requests that someone assign it an IP address.

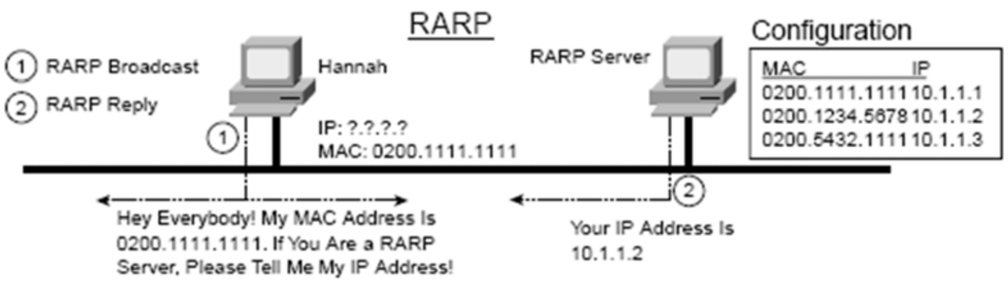

Network Layer (Layer 3) Addressing (cont,..) (14)

- Both protocols allow for IP address assignment, but that is all that RARP can ask for—it can't even ask for the subnet mask used on the LAN.

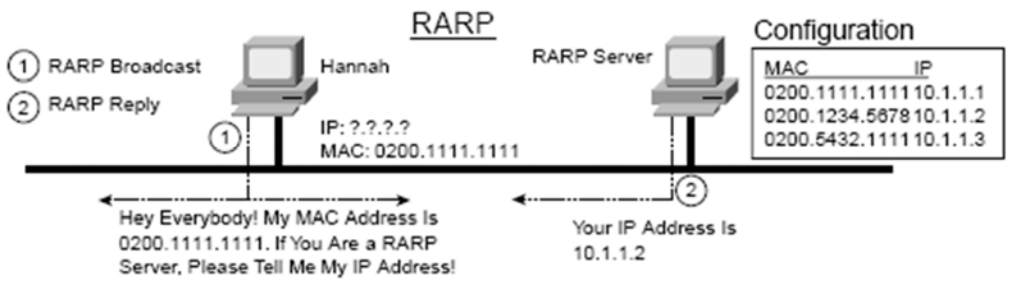

Network Layer (Layer 3) Addressing (cont,..) (15)

- BOOTP allows many more tidbits of information to be announced to a BOOTP client—its IP address, its subnet mask, its default gateway IP addresses, its other server IP addresses, and the name of a file that the computer should download.

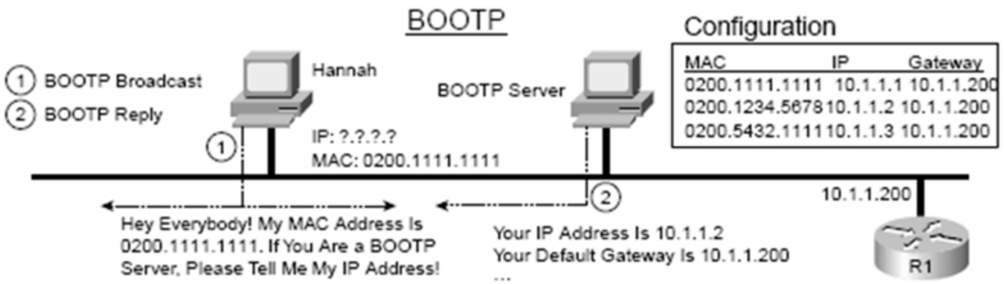

Network Layer (Layer 3) Addressing (cont,..) (16)

Why?

- Both RARP and BOOTP were created with the motivation to allow a diskless workstation to come up and start operating.

RARP?

With RARP, the creators of the protocol just wanted to get the machine an IP address so that a knowledgeable user could type in commands and copy the correct files from a server onto the diskless computer's RAM memory so that they could be used

Network Layer (Layer 3) Addressing (cont,..) (17)

BOOTP?

The creators of, anticipating a less sophisticated user in the future, wanted to automate as much of the process as possible—including the dynamic assignment of a default gateway (router) IP address.

- BOOTP's name really comes from the feature in which BOOTP supplies the name of a file to the BOOTP client.
- Typically, the diskless workstations had enough permanent memory to boot a very simple operating system, with the expectation that the computer would use a simple protocol, such as the Trivial File Transfer Protocol (TFTP), to transfer a file containing a more sophisticated operating system into RAM.

Network Layer (Layer 3) Addressing (cont,..) (18)

- So, with the ultimate goal being to let a diskless computer complete the processing of initializing, or *booting, a full operating system,* BOOTP was aptly named.
- Neither RARP nor BOOTP is used much to
- One of the **problems** with both RARP and BOOTP is that they required a computer to act as a server, and the server was required to know the MAC address of every computer and the corresponding configuration parameters that each computer should be told.
- So, administration in a network of any size was painful.

Network Layer (Layer 3) Addressing (cont,..) (19)

DHCP

- Like BOOTP, DHCP uses the concept of the client making a request and the server supplying the IP address to the client, plus other information such as the default gateway, subnet mask, DNS IP address, and other information.
- The biggest advantage of DHCP compared to BOOTP and RARP is that DHCP does not require that the DHCP server be configured with all MAC addresses of all clients.
- DHCP defines a process by which the server knows the IP subnet in which the DHCP client resides, and it can assign an IP address from a pool of valid IP addresses in that subnet.

Network Layer (Layer 3) Addressing (cont,..) (20)

- So, the DHCP server does not need to know the MAC address ahead of time.
- Also, most of the other information that DHCP might supply, such as the default router IP address, is the same for all hosts in the same subnet, so DHCP servers simply can configure information per subnet rather than per host and save a lot of administrative hassle compared to BOOTP.

Network Layer (Layer 3) Addressing (cont,..) (21)

- DHCP has become a very prolific protocol, with most end-user hosts on LANs in corporate networks getting their IP addresses and other basic configuration via DHCP.

DHCP Messages to Acquire an IP Address

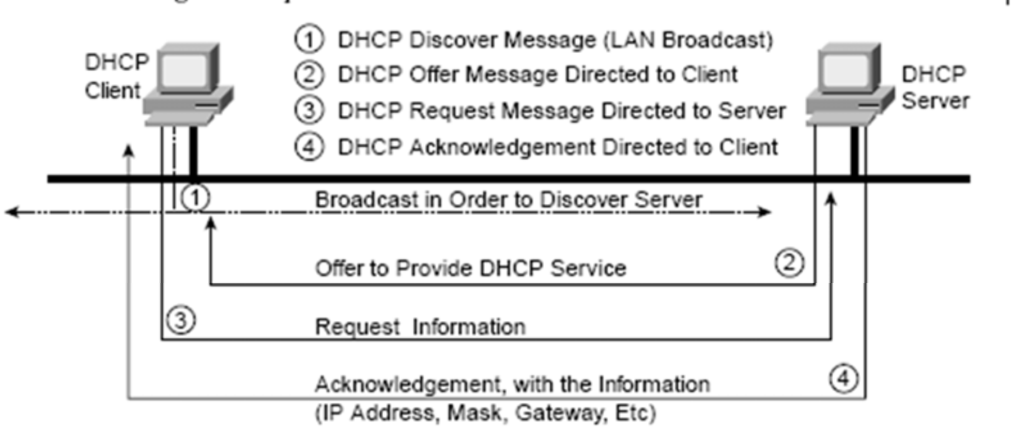

Simple Routing Example, with IP Subnets

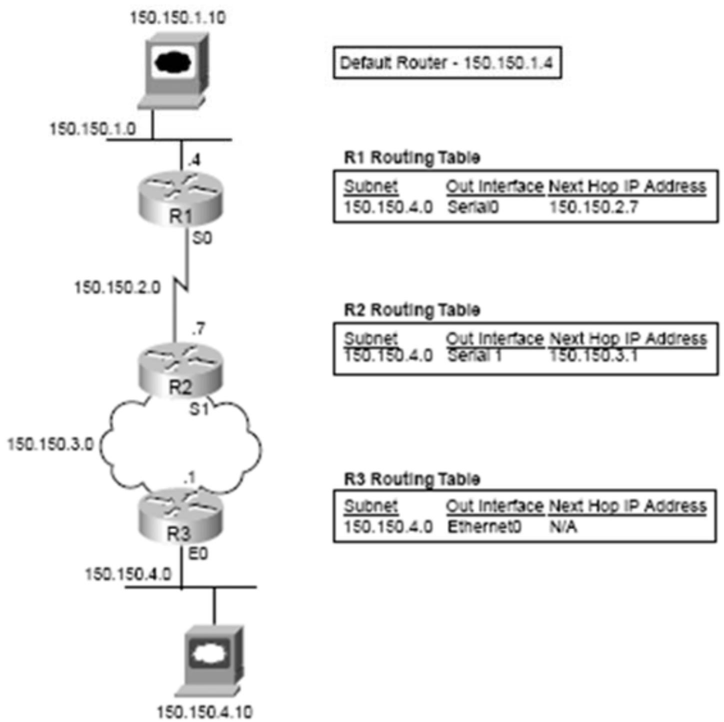

IP Routing and Routing Protocols

- First, a few detail about the figure need to be explained.
- The subnet numbers are shown, with the whole third octet used for the subnet part of the addresses.
- The actual IP addressed for PC1 and PC2 are shown.
- the full IP addresses of the routers are not shown
- only the host part of the address is listed

Simple Routing Example, with IP Subnets

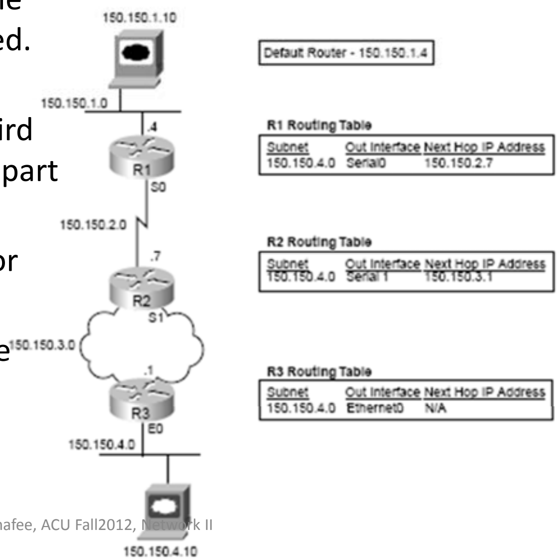

IP Routing and Routing Protocols (2)

- For instance, R2's IP address on the serial link to R1 is 150.150.2.7.
- The subnet is 150.150.2.0, and the .7 shown beside R2 in the figure represents the host part of the address, which is the fourth octet in this case.

Simple Routing Example, with IP Subnets

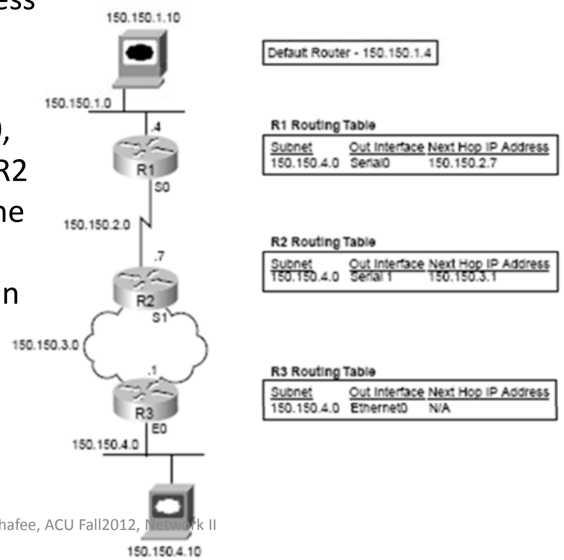

IP Routing and Routing Protocols (3)

Step 1

- PC1 sends the packet to R1—PC1 first builds the IP packet, with a destination address of PC2's IP address (150.150.4.10). PC1 needs to send the packet to R1 because it knows that its default router is 150.150.1.4.
- PC1 first checks its ARP cache, hoping to find R1's Ethernet MAC address.

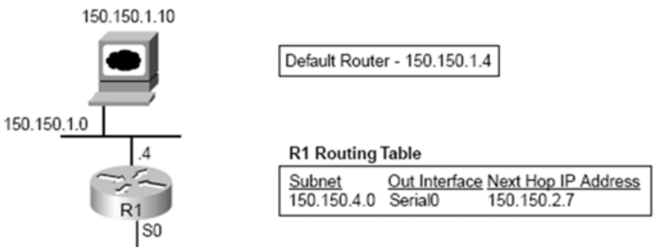

IP Routing and Routing Protocols (4)

- If it is not found, PC1 ARPs to learn R1's Ethernet MAC address.
- Then PC1 places the IP packet into an Ethernet frame, with a destination Ethernet address of R1's Ethernet address.
- PC1 sends the frame onto the Ethernet.

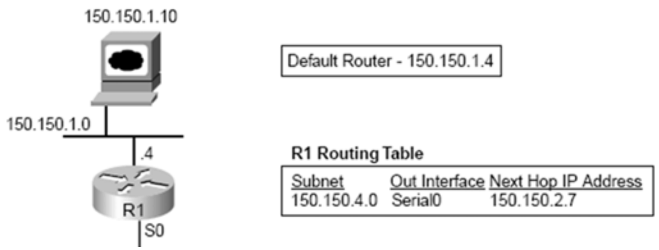

IP Routing and Routing Protocols (5)

Step 2

- **R1 processes the incoming frame and forwards the packet to R2—** Because the incoming Ethernet frame has a destination MAC of R1's Ethernet MAC, R1 copies the frame off the Ethernet for processing.
- If the FCS passes, meaning that the Ethernet frame did not have any errors in it, R1 looks at the Protocol Type field to discover that the packet inside the frame is an IP packet.

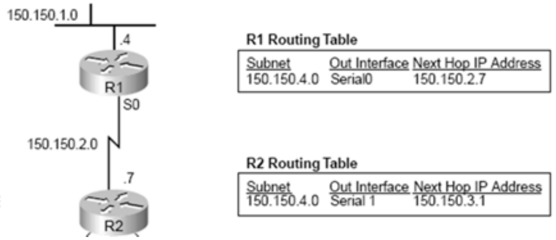

IP Routing and Routing Protocols (6)

- R1 then discards the Ethernet header and trailer.
- Next, R1 looks for the routing table entry that matches the destination address in the packet, 150.150.4.10.
- The routing table entry is listed in the figure—a route to subnet 150.150.4.0, with outgoing interface Serial0 to next-hop router R2 (150.150.2.7).

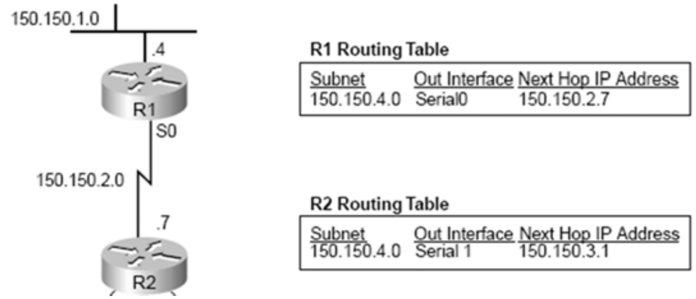

IP Routing and Routing Protocols (7)

- Now R1 just needs to build an HDLC frame and send it out its Serial0 interface to R2. As mentioned earlier, ARP is not needed on a point-topoint HDLC WAN link.
- R1 knows all the information necessary to out the packet inside an HDLC frame and send the frame.

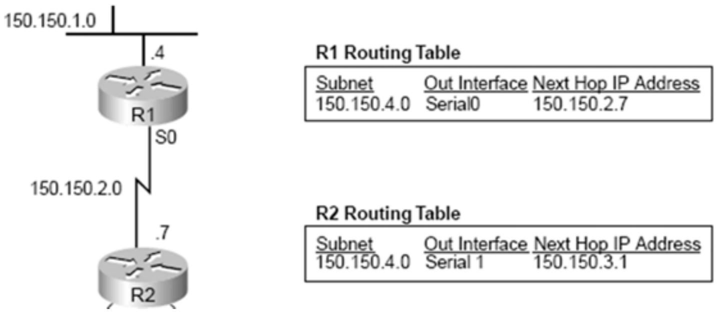

IP Routing and Routing Protocols (8)

Step 3

- R2 processes the incoming frame and forwards the packet to R3—R2 repeats the same general process as R1 when it receives the HDLC frame.
- After stripping the HDLC header and trailer, R2 also needs to find the routing table entry that matches destination 150.150.4.10.

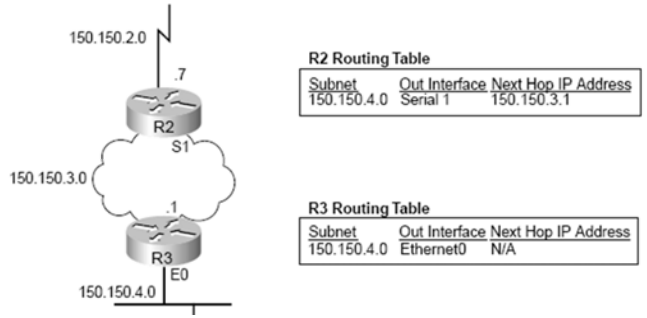

IP Routing and Routing Protocols (9)

- R2's routing table has an entry for 150.150.4.0, outgoing interface serial1, to next-hop router 150.150.3.1, which is R3.
- Before R2 can complete the task, the correct DLCI for the VC to R3 must be decided.

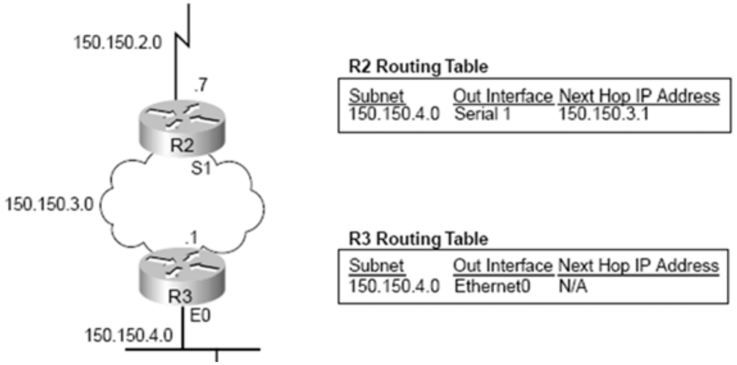

IP Routing and Routing Protocols (10)

Step 4

- R3 processes the incoming frame and forwards the packet to PC2— Like R1 and R2 before it, R3 checks the FCS in the data-link trailer, looks at the type field to decide whether the packet inside the frame is an IP packet, and then discards the Frame Relay header and trailer.

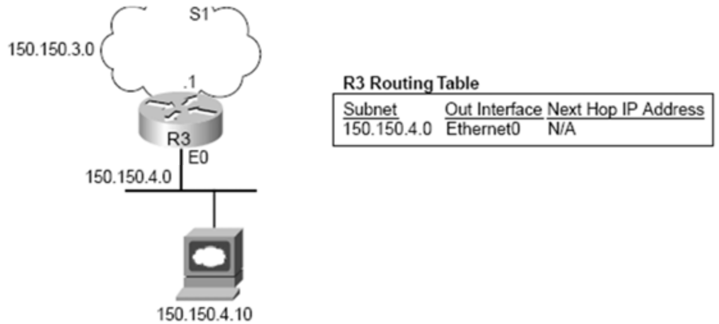

IP Routing and Routing Protocols (11)

- The routing table entry for 150.150.4.0 shows that the outgoing interface is R3's Ethernet interface, but there is no next-hop router because R3 is connected directly to subnet 150.150.4.0.

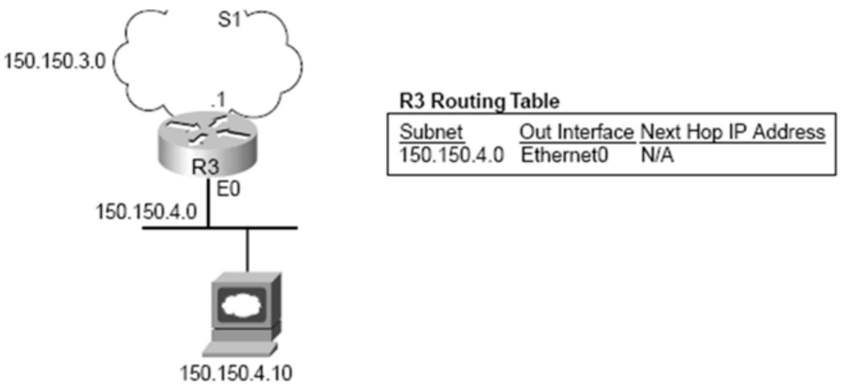

IP Routing and Routing Protocols (12)

- All R3 has to do is encapsulate the packet inside a Ethernet header and trailer, and forward the frame.
- Before R3 can finish building the Ethernet header, an IP ARP broadcast must be used to find PC2's MAC address (assuming that R3 doesn't already have that information in its IP ARP cache).

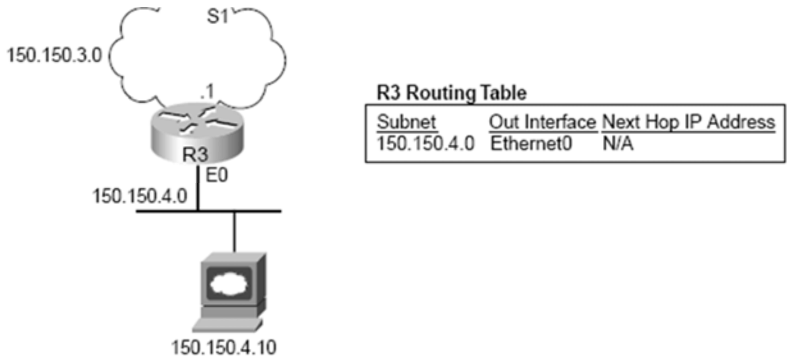

IP Routing and Routing Protocols (13)

- The routing process relies on the rules relating to IP addressing. For instance, why did 150.150.1.10 (PC1) assume that 150.150.4.10 (PC2) was not on the same Ethernet?
- Well, because 150.150.4.0, PC2's subnet, is different than 150.150.1.0, which is PC1's subnet.
- Because IP addresses in different subnets must be separated by some router, PC1 needed to send the packet to some router—and it did.
- Similarly, all three routers list a route to subnet 150.150.4.0, which, in this example, includes IP addresses 150.150.4.1 to 150.150.4.254.

IP Routing and Routing Protocols (14)

- What if someone tried to put PC2 somewhere else in the network, but still using 150.150.4.10? The routers then would forward packets to the wrong place.
- So, Layer 3 routing relies on the structure of Layer 3 addressing to route more efficiently.

IP Routing and Routing Protocols (15)

IP Routing Protocols

- IP routing protocols fill the IP routing table with valid, (hopefully) loop-free routes.
- Each route includes a subnet number, the interface out which to forward packets so that they are delivered to that subnet, and the IP address of the next router that should receive packets destined for that subnet (if needed).

IP Routing and Routing Protocols (16)

the goals of a routing protocol

- To dynamically learn and fill the routing table with a route to all subnets in the network.
- If more than one route to a subnet is available, to place the best route in the routing table.
- To notice when routes in the table are no longer valid, and to remove those routes from
- the routing table.
- If a route is removed from the routing table and another route through another neighboring router is available, to add the route to the routing table.

IP Routing and Routing Protocols (17)

- To add new routes, or to replace lost routes, with the best currently available route as quickly as possible. The time between losing the route and finding a working replacement route is called *convergence time.*
- To prevent routing loops.

IP Routing and Routing Protocols (18)

The logic

- basic logic that they use is relatively simple.
- Routing protocols take the routes in a routing table and send a message to their neighbors telling them about the routes.
- After a while, everyone has heard about all the routes.

IP Routing and Routing Protocols (19)

- Figure shows a sample network, with routing updates shown.

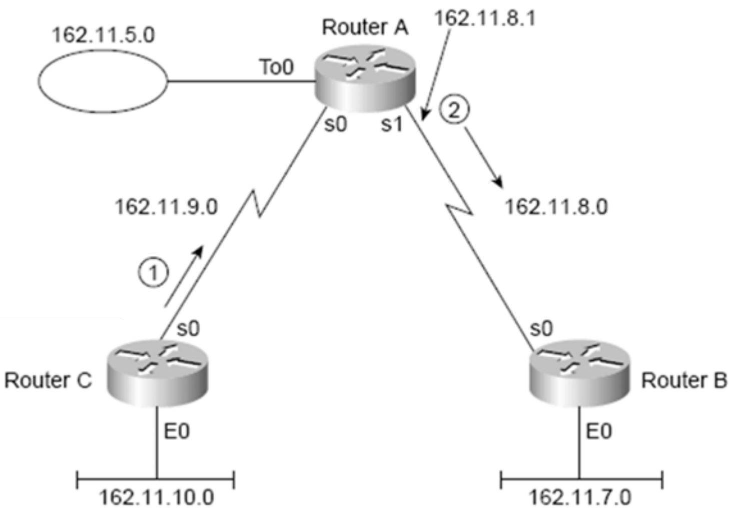

IP Routing and Routing Protocols (20)

- Table shows Router B's routing table before receiving the routing updates,

Group	Outgoing Interface	Next-Hop Router	Metric	Comments
162.11.7.0	E0	—	0	This is a directly connected route.
162.11.8.0	S0	—	0	This is a directly connected route.

- Router B adds routes for directly connected subnets when the interfaces first initialize.
- So, before Router B receives any routing updates, it knows about only two routes—the two connected routes—as listed

IP Routing and Routing Protocols (21)

- After receiving the update from Router A, Router B has learned three more routes.
- Because Router B learned those routes from Router A, all three of B's routes point back to Router A as the next hop router.

Group	Outgoing Interface	Next-Hop Router	Metric	Comments
162.11.5.0	S0	162.11.8.1	1	Learned from Router A, so next-hop is Router A.
162.11.7.0	E0	—	0	This is a directly connected route.
162.11.8.0	S0	—	0	This is a directly connected route.
162.11.9.0	S0	162.11.8.1	1	Learned from Router A, so next-hop is Router A.
162.11.10.0	S0	162.11.8.1	2	This one was learned from Router A, which learned it from Router C.

IP Routing and Routing Protocols (22)

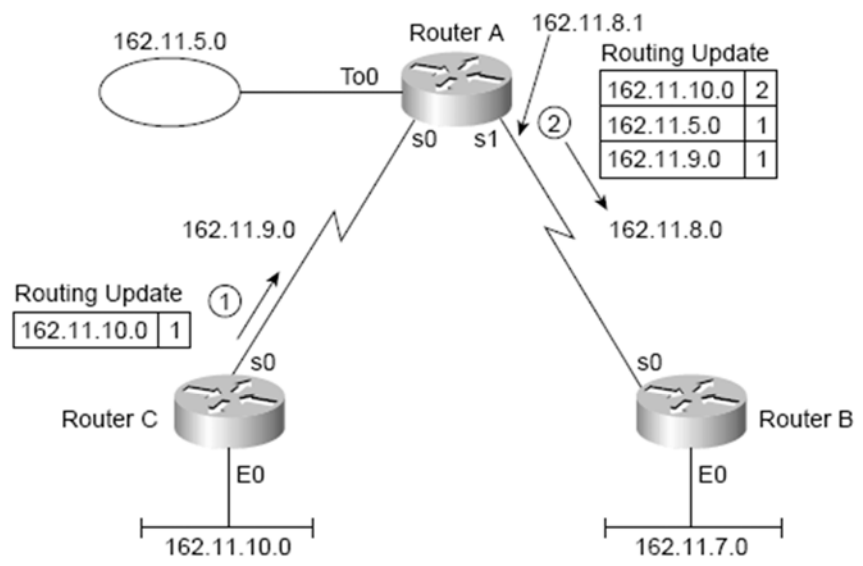

IP Routing and Routing Protocols (23)

- Router A learned about subnets 162.11.5.0 and 162.11.9.0 because A is connected directly to those subnets.
- Router A, in turn, learned about subnet 162.11.10.0, the subnet off Router C's Ethernet, from routing updates sent by Router C

Group	Outgoing Interface	Next-Hop Router	Metric	Comments
162.11.5.0	S0	162.11.8.1	1	Learned from Router A, so next-hop is Router A.
162.11.7.0	E0	—	0	This is a directly connected route.
162.11.8.0	S0	—	0	This is a directly connected route.
162.11.9.0	S0	162.11.8.1	1	Learned from Router A, so next-hop is Router A.
162.11.10.0	S0	162.11.8.1	2	This one was learned from Router A, which learned it from Router C.

Thanks,..

See you next week, isA,…