Mr. Ali bin Tahir **Email: ali@biit.edu.pk**

Week # 16 – Lecture # 31 – Application Layer

Recommended Reading:

Book: Computer Networking, A Top-Down Approach 6th edition, Authors: James F. Kurose, K W. Ross

Application Layer

In last lecture, we have completed our discussion on DNS. Today, we will discuss yet another popular network application i.e., electronic mail.

Here is the outline for today's lecture:

- **Electronic Mail in the Internet**
  - **Simple Mail Transfer Protocol (SMTP)**
  - **Mail Access Protocols**
- **Dynamic Host Configuration Protocol (DHCP)**
- **DHCP Client-Server Scenario**

Electronic Mail in the Internet

Electronic mail has been around since the beginning of the Internet. It remains one of the Internet's most important and utilized applications. As with ordinary postal mail, e-mail is an asynchronous communication medium—people send and read messages when it is convenient for them, without having to coordinate with other people's schedules. In contrast with postal mail, electronic mail is fast, easy to distribute, and inexpensive. Modern e-mail has many powerful features, including messages with attachments, hyperlinks, HTML-formatted text, and embedded photos.

Let's take a high-level view of the Internet mail system and its **key components**. Figure 2.16 presents a high-level view of the Internet mail system. We see from this diagram that it has **three major components**: **user agents**, **mail servers**, and the **Simple Mail Transfer Protocol (SMTP)**. We now describe each of these components in the context of a sender, Alice, sending an e-mail message to a recipient, Bob. User agents allow users to read, reply to, forward, save, and compose messages. Microsoft Outlook and Apple Mail are examples of user agents for e-mail. When Alice is finished composing her message, her user agent sends the message to her mail server, where the message is placed in the mail server's outgoing message queue. When Bob wants to read a message, his user agent retrieves the message from his mailbox in his mail server.

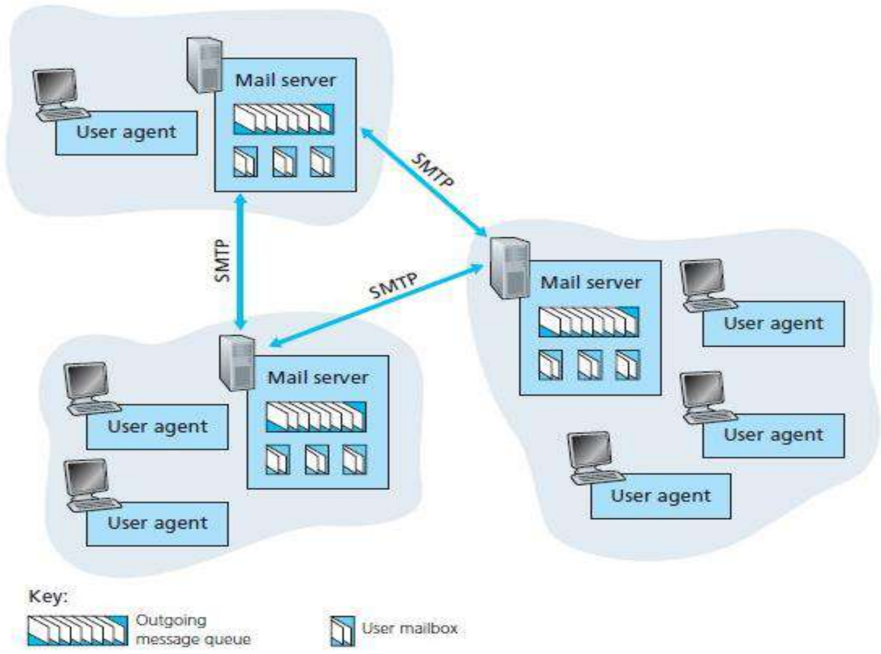

Figure 2.16 ♦ A high-level view of the Internet e-mail system

Mail servers form the core of the e-mail infrastructure. Each recipient, such as Bob, has a mailbox located in one of the mail servers. Bob's mailbox manages and maintains the messages that have been sent to him. A typical message starts its journey in the sender's user agent, travels to the sender's mail server, and travels to the recipient's mail server, where it is deposited in the recipient's mailbox.

When Bob wants to access the messages in his mailbox, the mail server containing his mailbox authenticates Bob (with usernames and passwords). Alice's mail server must also deal with failures in Bob's mail server. If Alice's server cannot deliver mail to Bob's server, Alice's server holds the message in a message queue and attempts to transfer the message later. Reattempts are often done every 30 minutes or so; if there is no success after several days, the server removes the message and notifies the sender (Alice) with an e-mail message.

Simple Mail Transfer Protocol (SMTP)

SMTP is the principal application-layer protocol for Internet electronic mail. It uses the reliable data transfer service of TCP to **transfer mail from the sender's mail server to the recipient's mail server**. As with most application-layer protocols, **SMTP has two sides: a client side, which executes on the sender's mail server, and a server side, which executes on the recipient's mail server**. Both the client and server sides of SMTP run on every mail server. When a mail server sends mail to other mail

Mr. Ali bin Tahir Email: ali@biit.edu.pk

servers, it acts as an SMTP client. When a mail server receives mail from other mail servers, it acts as an SMTP server.

To illustrate the basic operation of SMTP, let’s walk through a common scenario. Suppose Alice wants to send Bob a simple ASCII message. The scenario is summarized in Figure 2.17.

1. Alice invokes her user agent for e-mail, provides Bob’s e-mail address (for example, bob@someschool.edu), composes a message, and instructs the user agent to send the message.
2. Alice’s user agent sends the message to her mail server, where it is placed in a message queue.
3. The client side of SMTP, running on Alice’s mail server, sees the message in the message queue. It opens a TCP connection to an SMTP server, running on Bob’s mail server.
4. After some initial SMTP handshaking, the SMTP client sends Alice’s message into the TCP connection.
5. At Bob’s mail server, the server side of SMTP receives the message. Bob’s mail server then places the message in Bob’s mailbox.
6. Bob invokes his user agent to read the message at his convenience.

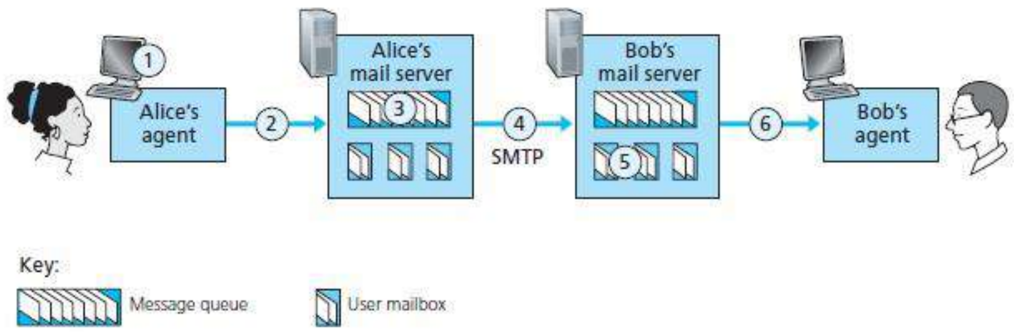

Figure 2.17 ♦ Alice sends a message to Bob

Let’s now take a closer look at how SMTP transfers a message from a sending mail server to a receiving mail server. First, the client SMTP (running on the sending mail server host) has TCP establish a connection to **port 25** at the server SMTP (running on the receiving mail server host). If the server is down, the client tries again later. Once this connection is established, the server and client perform some application-layer handshaking—just as humans often introduce themselves before transferring information from one to another, SMTP clients and servers introduce themselves before transferring information. During this SMTP handshaking phase, the SMTP client indicates the e-mail address of the sender (the person who generated the message) and the e-mail address of the recipient. Once the SMTP client and server have introduced themselves to each other, the client sends the message. SMTP relies on the reliable data transfer service of TCP to get the message to the server without errors.

SMTP Mail Transfer Example

Let's next take a look at an **example of messages exchanged between an SMTP client (C) and an SMTP server (S)**. The **hostname of the client is crepes.fr** and the hostname of the **server is hamburger.edu**. The ASCII text lines starts after C: are the lines the client sends into its TCP socket. The ASCII text lines starts after S: are the lines the server sends into its TCP socket. The following conversation begins as soon as the TCP connection is established.

S: 220 hamburger.edu

C: HELO crepes.fr

S: 250 Hello crepes.fr, pleased to meet you

C: MAIL FROM: <alice@crepes.fr>

S: 250 alice@crepes.fr ... Sender ok

C: RCPT TO: <bob@hamburger.edu>

S: 250 bob@hamburger.edu ... Recipient ok

C: DATA

S: 354 Enter mail, end with "." on a line by itself

C: Do you like ketchup?

C: How about pickles?

C: .

S: 250 Message accepted for delivery

C: QUIT

S: 221 hamburger.edu closing connection

In the example above, the client sends a message ("Do you like ketchup? How about pickles?") from mail server crepes.fr to mail server hamburger.edu. As part of the dialogue, the client issued **five commands**: **HELO** (an abbreviation for HELLO), **MAIL FROM, RCPT TO, DATA, and QUIT**. These commands are self-explanatory. The client also sends a line consisting of a single period, which indicates the end of the message to the server. The server issues reply to each command, with each reply having a reply code and some (optional) English-language explanation.

Mail Access Protocols

Once SMTP delivers the message from Alice's mail server to Bob's mail server, the message is placed in Bob's mailbox. But there is still one missing piece to the puzzle! How does a recipient like Bob, running a user agent on his local PC, obtain his messages, which are sitting in a mail server within Bob's ISP? Note that Bob's user agent can't use SMTP to obtain the messages because obtaining the messages is a pull operation (since bob user agent needs to pull email from server), whereas SMTP is a push protocol since it pushes email from one server to another. The puzzle is completed by introducing a special mail access protocol that transfers messages from Bob's mail server to his local PC.

There are currently a number of popular **mail access protocols**, including **Post Office Protocol—Version 3** (**POP3**), **Internet Mail Access Protocol (IMAP)**, and **HTTP**. Figure 2.18 provides a summary of the protocols that are used for Internet mail: SMTP is used to transfer mail from the sender's mail server to the recipient's mail server; SMTP is also used to transfer mail from the sender's user agent to the sender's mail server. A mail access protocol, such as POP3, is used to transfer mail from the recipient's mail server to the recipient's user agent.

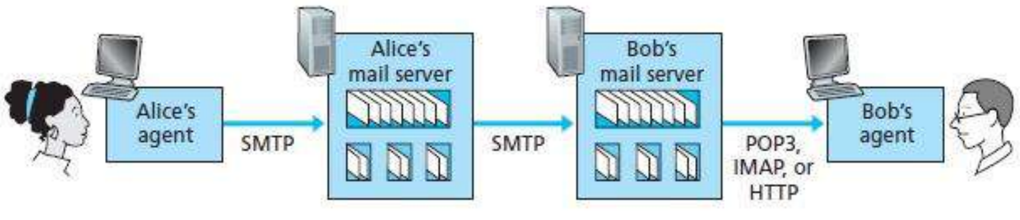

Figure 2.18 ♦ E-mail protocols and their communicating entities

Dynamic Host Configuration Protocol (DHCP)

Once an organization has obtained a block of addresses, it can assign individual IP addresses to the host and router interfaces in its organization. A system administrator will typically manually configure the IP addresses into the router (often remotely, with a network management tool). Host addresses can also be configured manually, but more often this task is now done using the Dynamic Host Configuration Protocol (DHCP). DHCP allows a host to obtain (be allocated) an IP address automatically. In addition to host IP address assignment, DHCP also allows a host to learn additional information, such as its **subnet mask**, the address of its first-hop **router** (often called the **default gateway**), and the address of its **local DNS server**.

Usage Scenarios

Because of DHCP's ability to automate the network-related aspects of connecting a host into a network, it is often referred to as a plug-and-play protocol. This capability makes it *very* attractive to the network administrator who would otherwise have to perform these tasks manually!

DHCP is also enjoying widespread use in residential Internet access networks and in wireless LANs, where hosts join and leave the network frequently. Consider, for example, the student who carries a laptop from a dormitory room to a library to a classroom. It is likely that in each location, the student will be connecting into a new subnet and hence will need a new IP address at each location. DHCP is ideally suited to this situation, as there are many users coming and going, and addresses are needed for only a limited amount of time.

DHCP is similarly useful in residential ISP access networks. Consider, for example, a residential ISP that has 2,000 customers, but no more than 400 customers are ever online at the same time. In this case, rather than needing a block of 2,048 addresses, a DHCP server that assigns addresses dynamically needs only a block of 512 addresses (for example, a block of the form a.b.c.d/23). As the hosts join and leave, the DHCP server needs to update its list of available IP addresses. Each time a host joins, the DHCP server allocates an arbitrary address from its current pool of available addresses; each time a host leaves, its address is returned to the pool.

Mr. Ali bin Tahir Email: ali@biit.edu.pk

Week # 16 – Lecture # 32 – Application Layer

Recommended Reading:

Book: Computer Networking, A Top-Down Approach 6th edition, Authors: James F. Kurose, K W. Ross

In previous lectures we have started our discussion on DHCP. In this lecture, we will continue the same topic. After completing DHCP, we will discuss firewalls which is an important component of today's computer networks from security perspective.

Here is the outline for today's lecture:

- **DHCP Client-Server Scenario**

DHCP Client-Server Scenario

DHCP is a client-server protocol. A client is typically a newly arriving host wanting to obtain network configuration information, including an IP address for itself. In the simplest case, each subnet will have a DHCP server. If no server is present on the subnet, a **DHCP relay agent** (typically a router) that knows the address of a DHCP server for that network is needed. Figure 4.20 shows a DHCP server attached to subnet 223.1.2/24, with the router serving as the relay agent for arriving clients attached to subnets 223.1.1/24 and 223.1.3/24. In our discussion below, we'll assume that a DHCP server is available on the subnet.

For a newly arriving host, the DHCP protocol is a **four-step process**, as shown in Figure 4.21 for the network setting shown in Figure 4.20. In this figure, yiaddr (as in "your Internet address") indicates the address being allocated to the newly arriving client. The **four steps** are:

- ***DHCP server discovery.*** The first task of a newly arriving host is to find a DHCP server with which to interact. This is done using a DHCP discover message, which a **client sends within a UDP packet to port 67**. The UDP packet is encapsulated in an IP datagram. But to whom should this datagram be sent? The host doesn't even know the IP address of the network to which it is attaching. Given this, the DHCP client creates an IP datagram containing its DHCP discover message along with the broadcast destination IP address of

255.255.255.255 and a “this host” source IP address of 0.0.0.0. The DHCP client passes the IP datagram to the link layer, which then broadcasts this frame to all nodes attached to the subnet.

- ***DHCP server offer(s).*** A DHCP server receiving a DHCP discover message responds to the client with a DHCP offer message that is broadcast to all nodes on the subnet, again using the IP broadcast address of 255.255.255.255. (You might want to think about why this server reply must also be broadcast). Since several DHCP servers can be present on the subnet, the client may find itself in the position of being able to choose from among several offers. Each server offer message contains the transaction ID of the received discover message, the proposed IP address for the client, the network mask, and an IP address lease time—the amount of time for which the IP address will be valid. It is common for the server to set the lease time to several hours or days.

- ***DHCP request.*** The newly arriving client will choose from among one or more server offers and respond to its selected offer with a DHCP request message, echoing back the configuration parameters.

- ***DHCP ACK.*** The server responds to the DHCP request message with a DHCP ACK message, confirming the requested parameters.

Mr. Ali bin Tahir Email: ali@biit.edu.pk

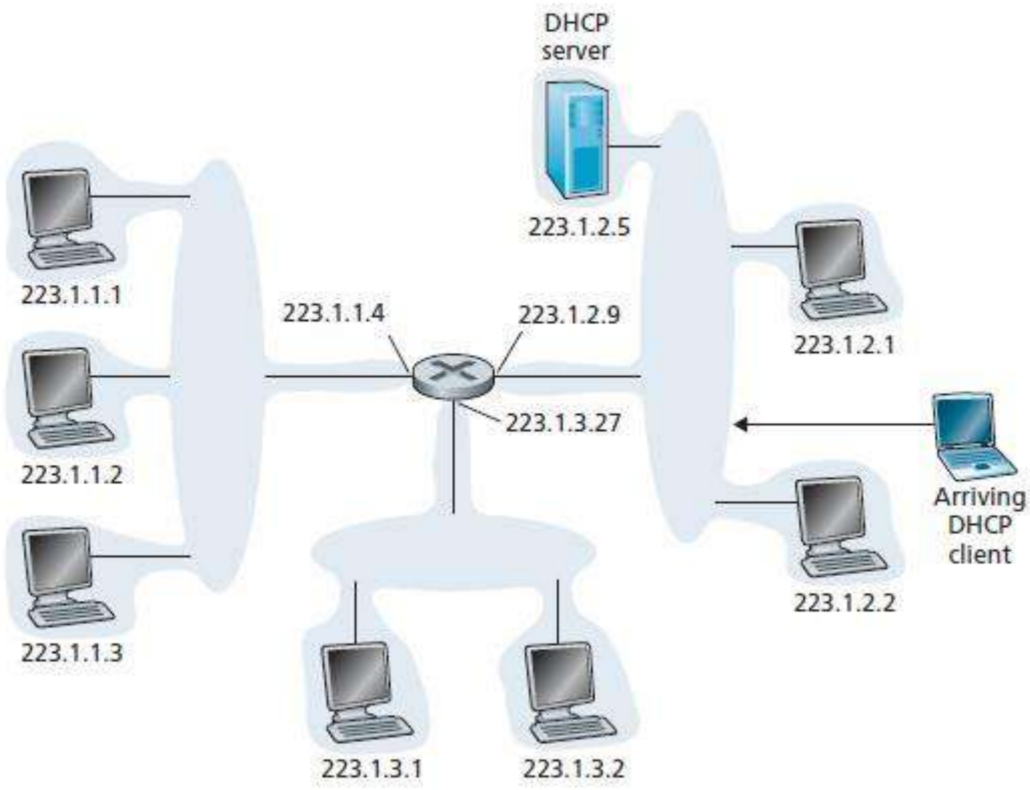

Figure 4.20 ♦ DHCP client-server scenario

Once the client receives the DHCP ACK, the interaction is complete and the client can use the DHCP-allocated IP address for the lease duration. Since a client may want to use its address beyond the lease's expiration, DHCP also provides a mechanism that allows a client to renew its lease on an IP address.

Mr. Ali bin Tahir
Email: ali@biit.edu.pk

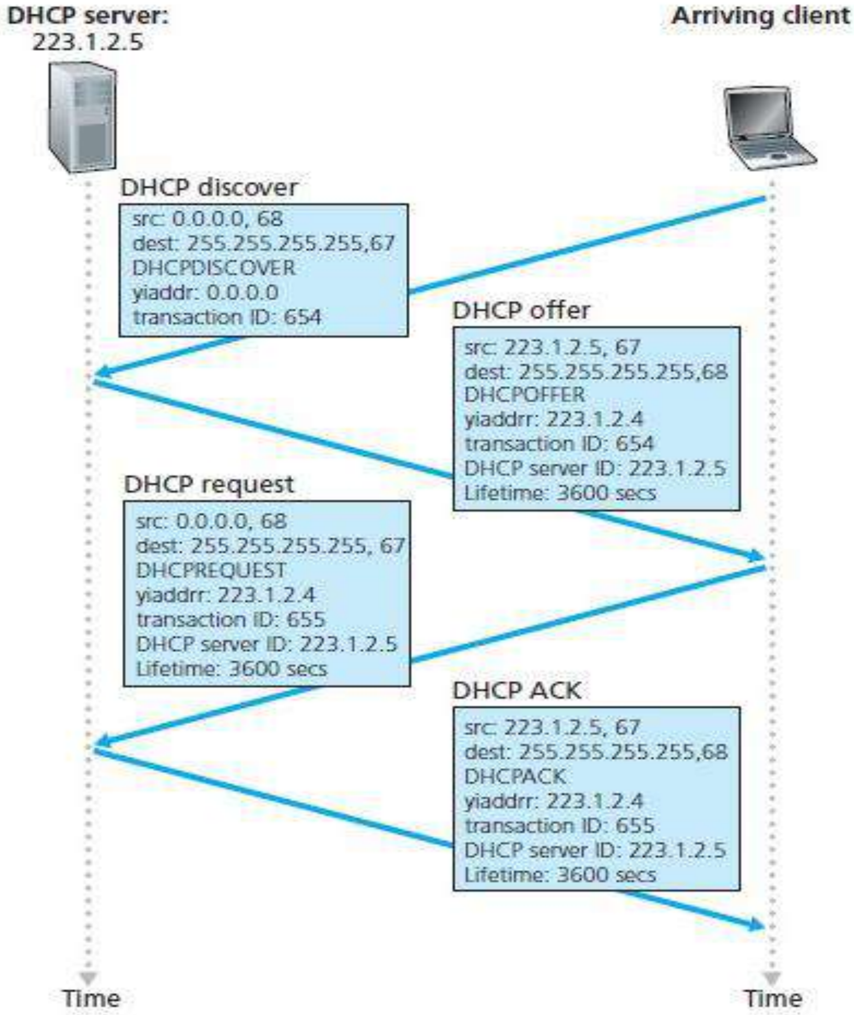

Figure 4.21 ♦ DHCP client-server interaction

The End!