E-mail on the Internet

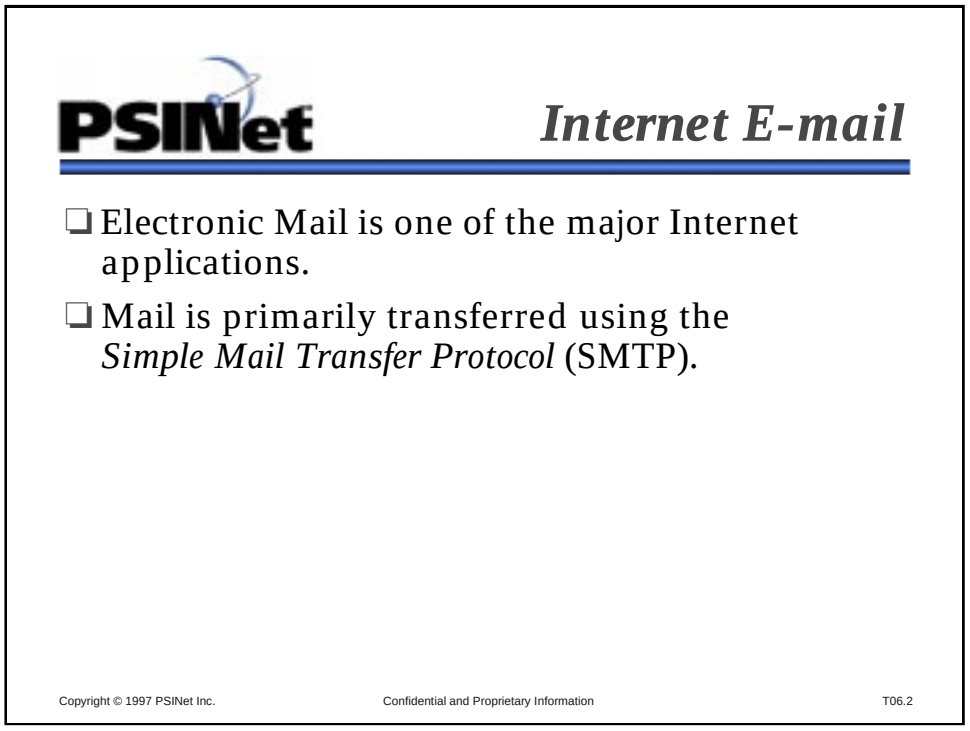

One of the primary uses of the Internet is **e-mail**. E-mail has become a major means of communication. Although there are many languages used to transfer e-mail, the e-mail that is sent across the Internet primarily uses the SMTP language or protocol. There are other techniques and protocols for e-mail, such as POP and UUCP, but this presentation will concentrate on the SMTP standard.

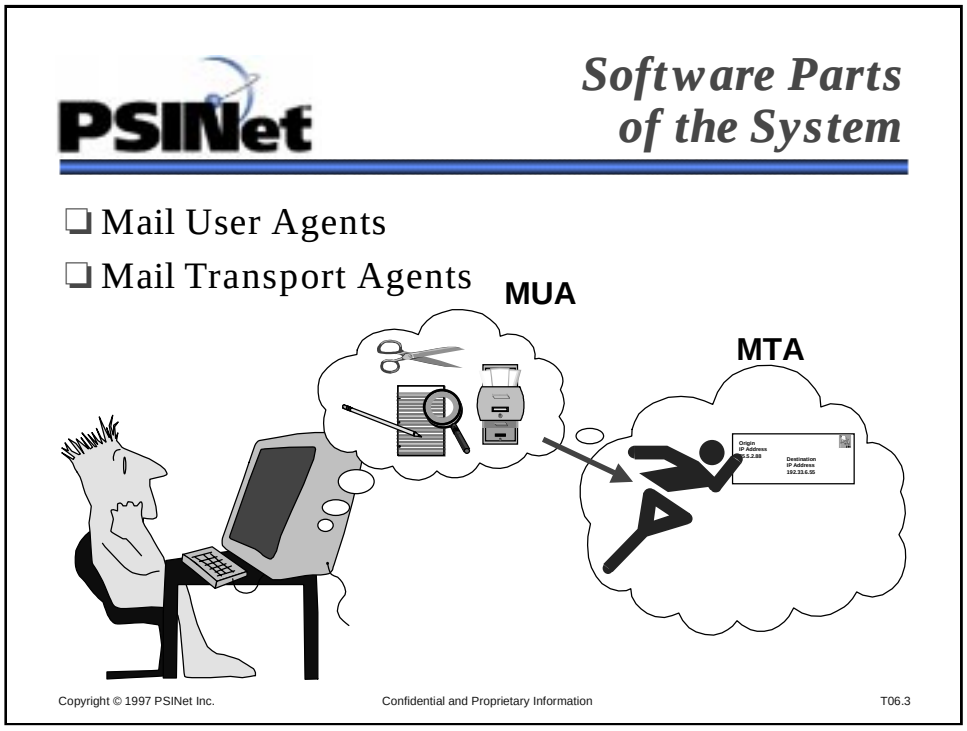

A **Mail User Agent**, or **MUA**, is the program that is used to create, edit, read, or manage e-mail messages. The MUA hands off the message to the **MTA** for delivery.

A **Mail Transport Agent**, or **MTA**, is the program that is responsible for the delivery of a message from one server to another.

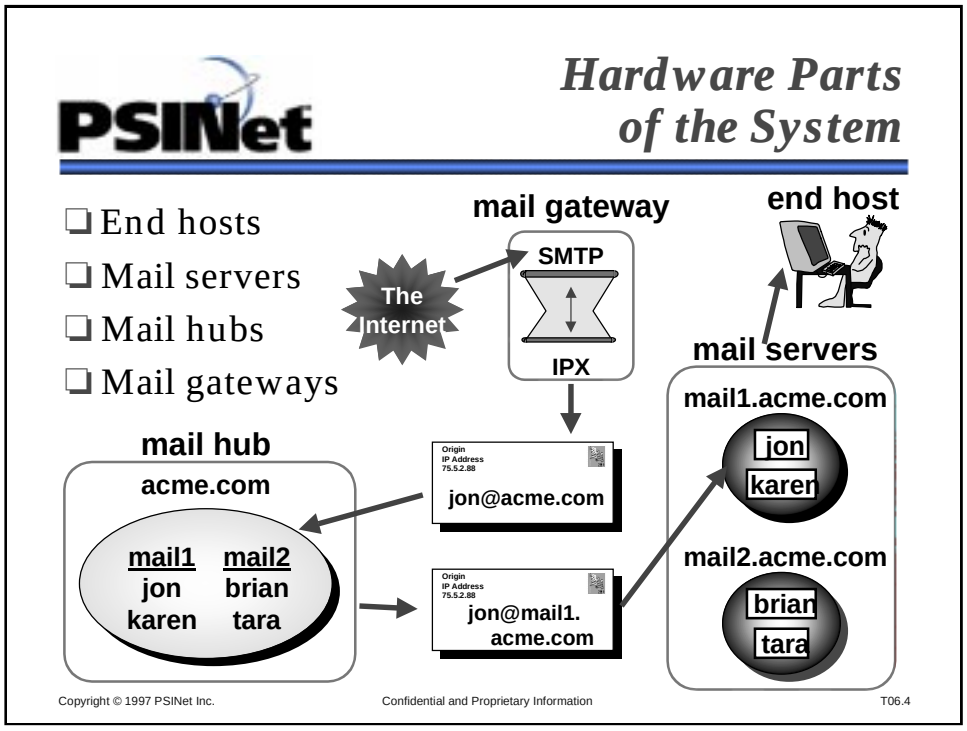

Generally, user e-mail is stored in large central spools called **mailboxes**. The user runs a MUA on an end host to access the mailbox on a mail server. An **end host** is a user's workstation, and a **mail server** is a generic term for all mail machines that contain mailboxes or **mail spools**. The end host and the mail server could be the same machine, or they can be on different networks.

Some sites run multiple mail servers. In this situation steps must be taken to ensure that user e-mail is directed to the correct machine. There are two approaches that may be used to accomplish this. One is to use distributed e-mail addresses that specify the correct mail server. Outside hosts would deliver directly to a specific machine such as "mail1.company.com" or "mail2.company.com." This configuration reveals internal structural information. The second approach is to create a mail hub that knows which users are served by which servers. All inbound mail must go through the hub, and the hub re-addresses the mail appropriately for each user.

Some sites use other networking protocols in addition to TCP/IP, such as IPX, DECNet, AppleTalk, etc. These sites need to run a mail gateway. Outgoing mail must be translated from the native networking protocol to SMTP over TCP/IP. Incoming mail must also be translated from SMTP over TCP/IP into the appropriate native mail and networking protocols.

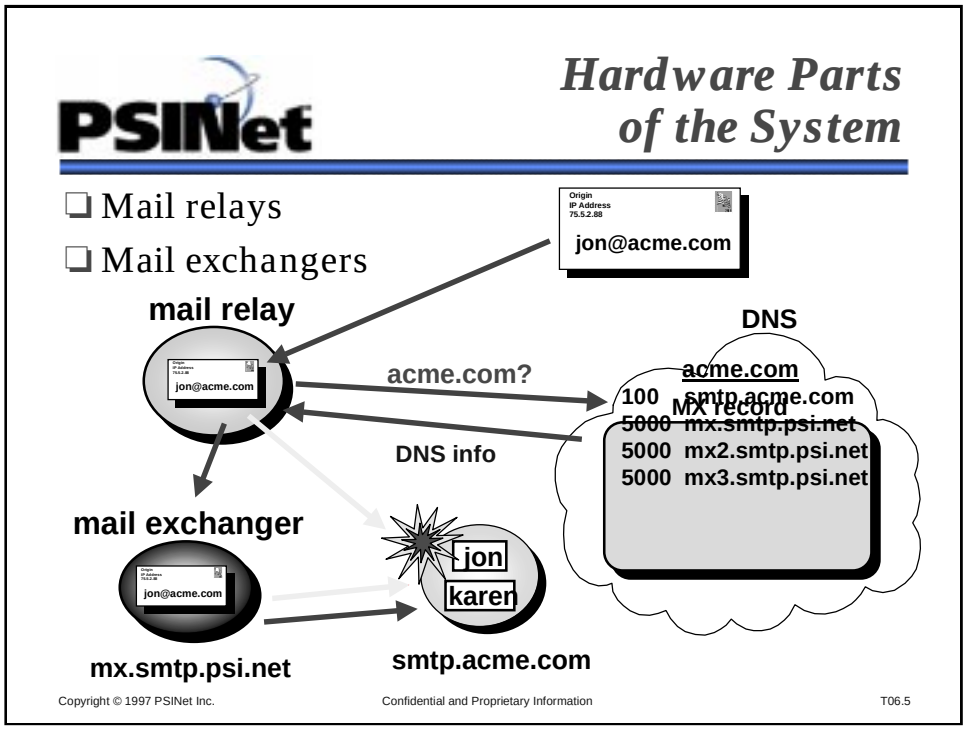

Mail gateways sometimes don't have the ability to execute DNS queries. Such machines send all outbound mail to a **mail relay**, or **smarterhost**, for delivery. The mail relay machine has the ability to resolve DNS name queries.

Mail destinations are assigned to domain names using **DNS mail exchanger (MX) records**. Multiple MX records can be assigned and ranked in order of preference, with zero being the highest preference. An MX record with a low preference (often indicated by a high number) is a backup to higher preferenced MX records. These lower preference records contain DNS information of backup machines called **mail exchangers**.

Mail exchangers receive messages that the mail relay is not able to forward to the destination mail server. Once the message is received, the mail exchanger will continue to attempt to deliver the message to the destination mail server until it either succeeds or gives up after a predetermined time period. PSINet mail exchangers will attempt to deliver an e-mail message for a period of seven days.

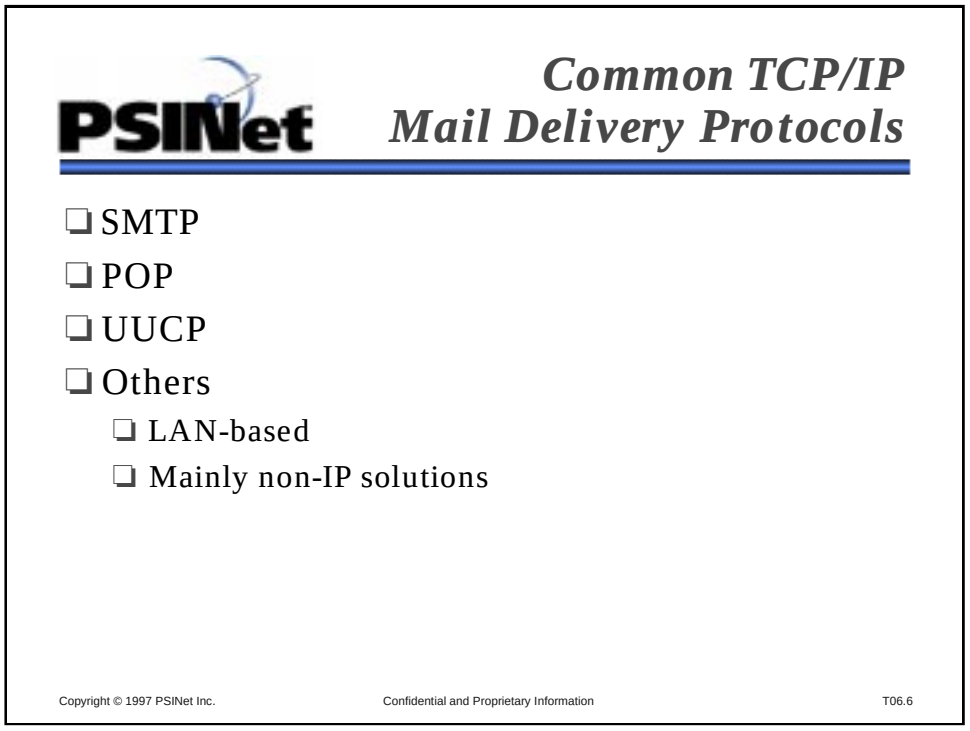

A protocol is a language computers use to exchange mail with each other. Computers must be able to speak the same protocol to exchange mail.

Simple Mail Transport Protocol (SMTP) is the Internet standard. It is required for the delivery of mail. It delivers mail between servers, but not for final delivery.

POP is a common client-server mail protocol. POP is a store and forward delivery paradigm. A client connects to a local POP server and downloads the user's mailbox. The client then processes the mail locally.

UUCP is common for sites that don't have a TCP/IP connection to the Internet, and also is a store and forward delivery paradigm. A client connects to the specified UUCP server and downloads the site's mail to the local spool. Once the mail is in the local spool, the mail is processed and delivered to the user's mailbox.

Other LAN-based protocols provide e-mail for local area networks. They are often based on non-IP network protocols such as IPX or AppleTalk. Some examples are Lotus Notes, cc:Mail, MS-Mail, and QuickMail. All of these programs have their own protocols which are proprietary to the software vendor, thus there is a need for an interpreter. Since SMTP is the universal language (protocol), there must be a proprietary protocol to SMTP translator or **SMTP gateway**.

Mail Protocols - SMTP

- SMTP is defined by RFC 821.
- Message format is defined by RFC 822.
- All sites that are connected to the Internet must have the ability to send SMTP mail.
- The recipient has no control over when the sender is going to send; therefore, it must always be ready to receive messages.

Local Mail Server Remote Mail Server

SMTP is a TCP-based, client-server protocol. Its operation is simple: A port greater than 1023 is opened on a client. The connection is established to port 25 on a server. The client sends the server a list of the recipients of the message as well as the sender's address. This is called the "envelope." The client then sends one or more messages to the server and the connection is closed. This conversation takes place in a formal language of four-character commands and three-digit reply codes.

Here is an example of a SMTP conversation:

jpublic@domain.com... Connecting to smtp.domain.com via smtp...

220 smtp.domain.com SMTP Sendmail 8.7/8.7/0806 Sun, 15 Oct 1995 22:47:52 -0700

>>> HELO foo.edu

250 foo.edu Hello user@foo.edu, pleased to meet you

>>> MAIL From: user@foo.edu

250 ... Sender ok

>>> RCPT To: jpublic@domain.com

250 Recipient ok

>>> DATA

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

>>> .

250 WAA12161 Message accepted for delivery jpublic@domain.com... Sent (WAA12161 Message accepted for delivery) Closing connection to smtp.domain.com.

>>> QUIT

221 smtp.domain.com closing connection

Mail Protocols - POP

- Post Office Protocol uses TCP / IP.
- Defined by RFC 1725.
- POP server receives mail via SMTP.

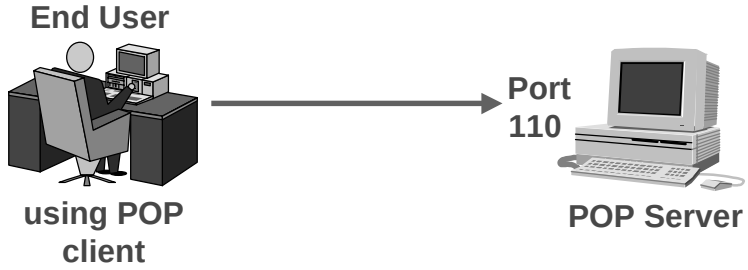

POP requires both server and client software. A client opens a port greater than 1023 and connects to port 110 on the POP server. The client authenticates, requests mail, and cleans up the spool. When the client requests mail, the messages are downloaded to the client where the user can read the messages. Note that there are two version of POP mail in existence, POP2 and POP3. POP2 connects to port 109 on the POP server and POP3 connects to port 110.

Some examples of POP servers:

MailShare, PD Macintosh version

NTMail, commercial Windows NT version

MH, PD version for UNIX

Popper, PD version for UNIX

Some examples of POP Clients:

Chameleon (PC)

Eudora (PC/MAC)

TCP/Connect II (PC/MAC)

Internet Valet (MAC)

Microphone (MAC)

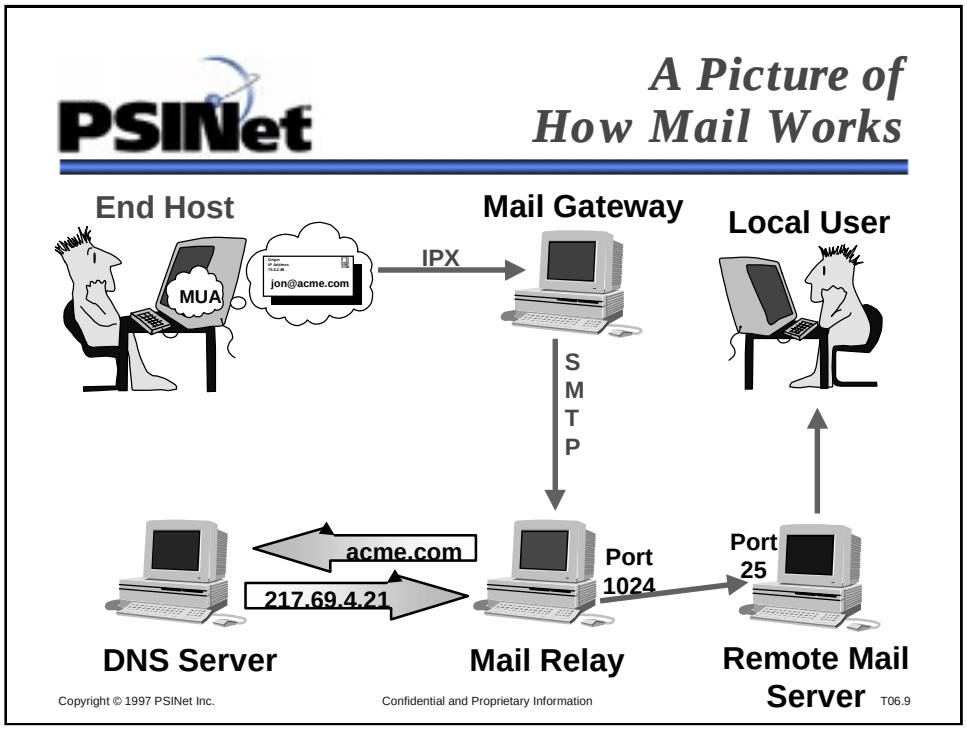

An end user uses a MUA to create a message on an end host.

The message gets to a mail gateway which converts the message to SMTP and sends it to a mail relay.

The mail relay looks up who it's going to by translating the domain name to an IP address.

Once the domain name has been resolved, the mail relay delivers the mail to the remote mail server.

The remote mail server then accepts the message and delivers it to the local user.

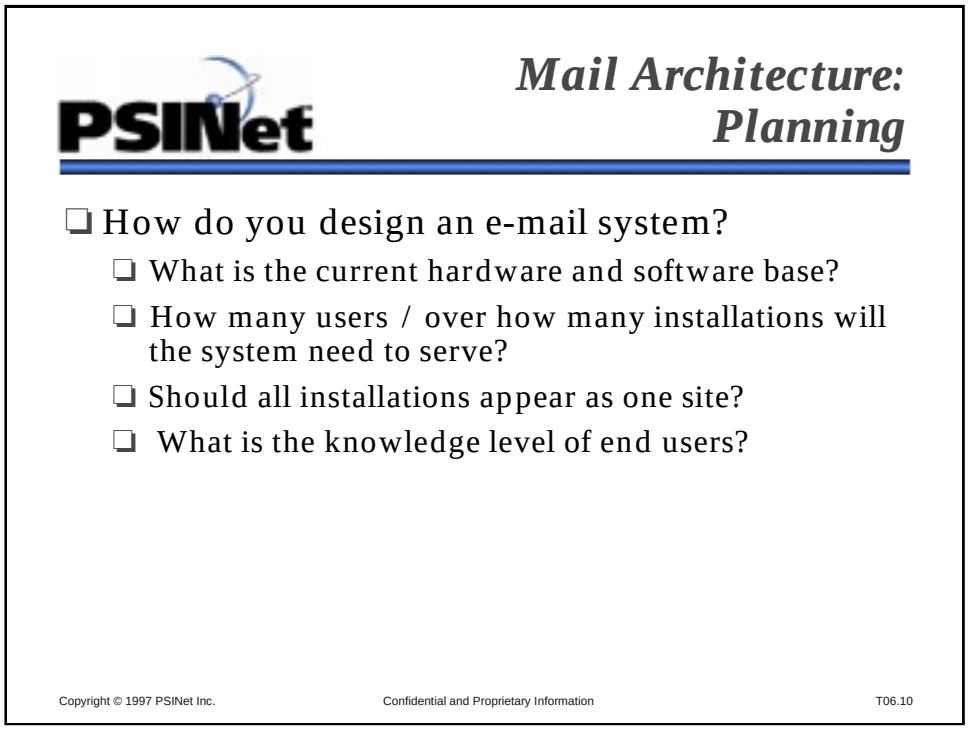

There are lots of things to consider when planning an e-mail system.

The first would be, **what is the current hardware and software base?** If there is a solution that could use the current hardware and software, it would be preferred.

How many users / over how many installations will the system need to serve? This would influence the need of resources. If there are only going to be two or three users, there would be no need to have multiple mail servers.

Should all installations appear as one site? This would identify the need for a mail hub or mail server.

What is the knowledge level of end users? If all the users have knowledge of UNIX, then it would be natural to use a UNIX host instead of a cc:Mail machine.

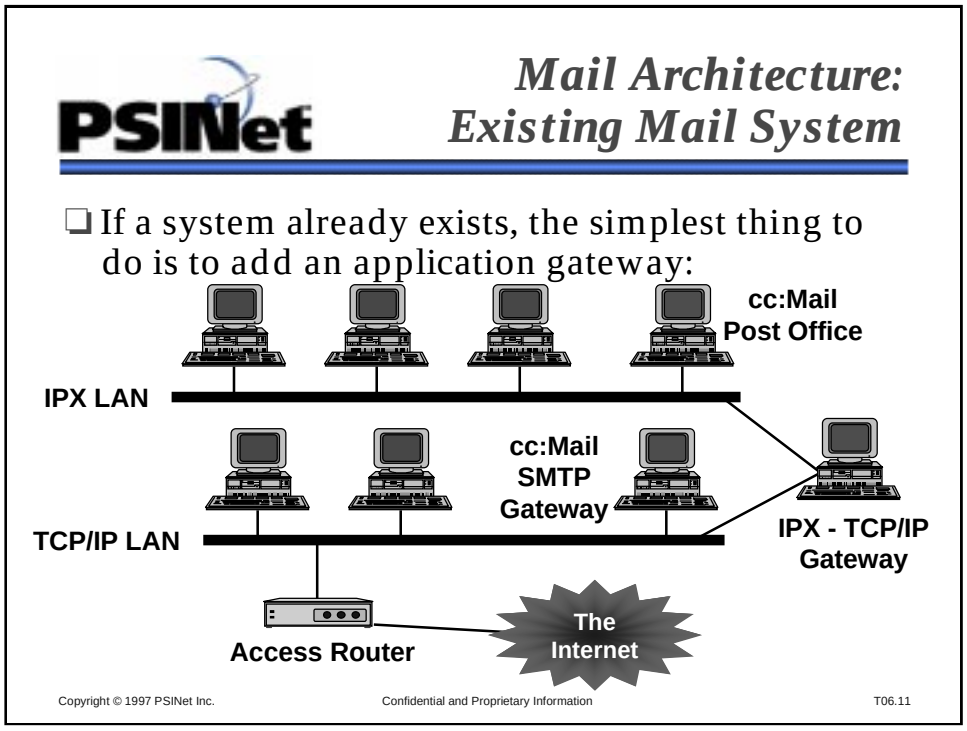

If a system already exists, the simplest thing to do is to add an SMTP gateway. In this example, a cc:Mail system is in existence and the LAN has become connected to the Internet. The best solution here is to purchase a cc:Mail SMTP gateway. It would be wasteful to replace the current system with a new SMTP system.

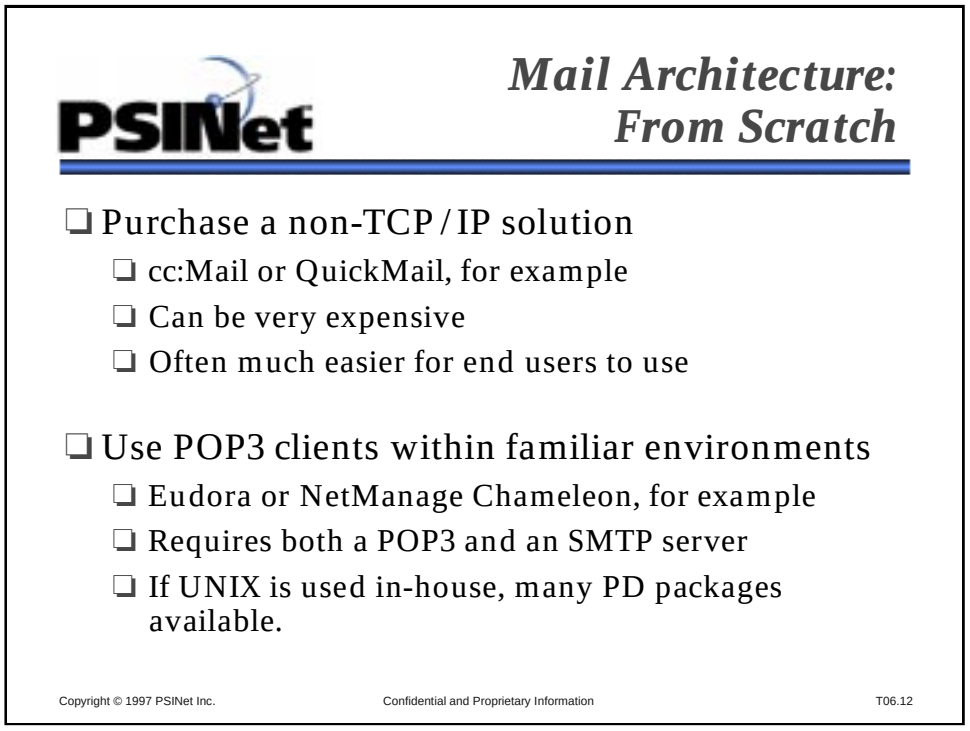

If there is nothing in place, then the whole system will have to be purchased. Purchasing a non-TCP/IP solution may be expensive, but it may be easy for the end-user to adapt to. Non-TCP/IP solutions tend to be very user friendly, hence the higher price.

Another solution would be to obtain a POP server and some POP clients. There are a some UNIX-based POP server software packages that are public domain and therefore are free. However, most POP clients are part of a TCP/IP software package which is usually not free. If the TCP/IP software package that is installed on the end hosts has the ability to use POP, then it would be best to obtain a POP server and use the POP client to read the mail from the POP server.

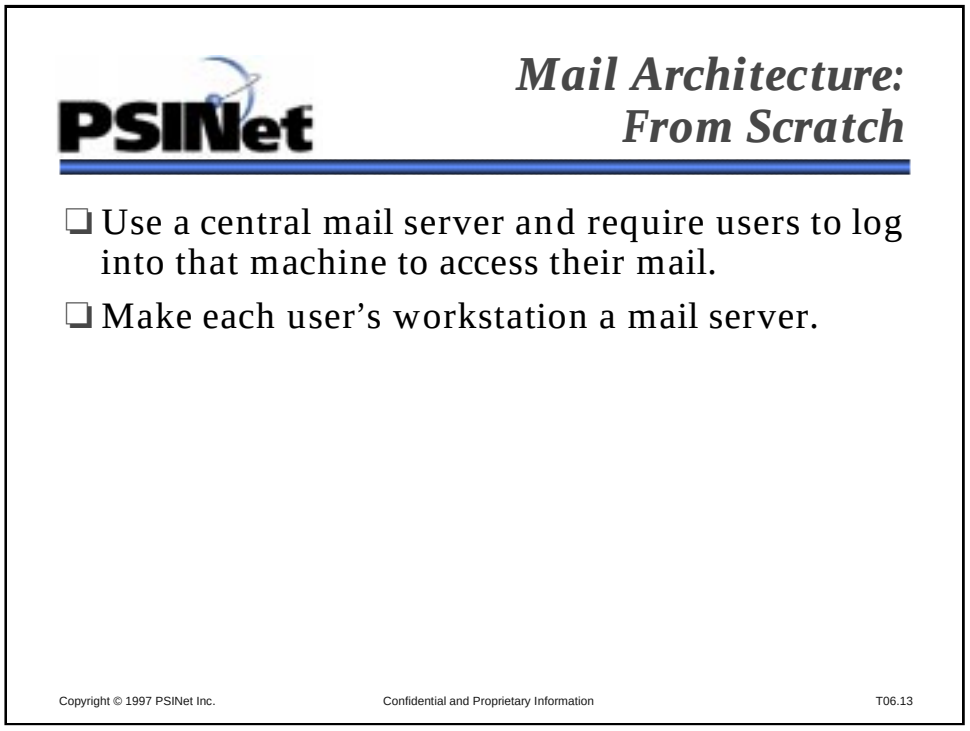

One way to create a mail architecture is to set up a central mail server that contains all user mailboxes. This would force everyone to use a specific platform (UNIX, VMS, NT, etc.). Although this system is somewhat simple to implement, the interface through which users access their e-mail may be complex. Additional time and resources may be required to simplify the interface and/or train all users on the new system.

Making each user's workstation an SMTP server would be a system administrator's nightmare. This would be a great waste of resources and effort. It would also reveal the internal structure of the LAN. In order to conceal this particular architecture, a mail hub would have to be added.

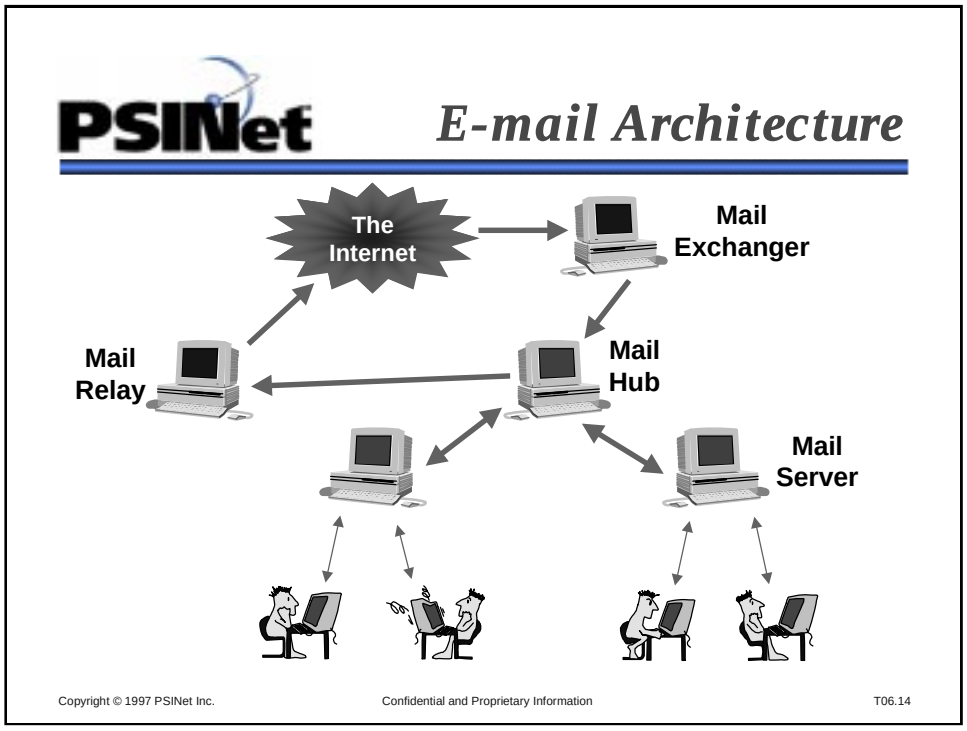

In general, it is best to have one mail server for each group that needs different access privileges, each group that should have a different e-mail address, each group that is connected by WAN or other slow link, and each distinct network protocol (e.g., IPX vs. IP).

Mail Exchangers may be useful if the mail hub has trouble resolving domain names or if the machine has high loads. Avoid setting up mail exchangers for hosts who will be receiving sensitive documents. The mail exchangers are just one more possible interception point.

Avoid forcing all mail through an exchanger. It will likely to cause bursty flow patterns and possible unnecessary outages.

For many mail gateways, a relay is necessary (cc:Mail and QuickMail gateways both need relays as do many others). Local relays are far more effective than remote because WAN bandwidth is more scarce than LAN bandwidth.

Avoid using relays for sensitive documents for the same reason as mail exchangers; it’s just one more potential interception point. Consider encrypting the mail if a relay is unavoidable.

PSINet

Questions?