Mail System Background

Mail System (1)

- Major components
  - **Mail User Agent (MUA)**
    - Help user read and compose mails
  - **Mail Transport Agent (MTA)**
    - Route mails among machines
  - **Delivery Agent (DA)**
    - Place mails in users' mail boxes
  - **Access Agent (AA)**
    - Connects the user agent to the mail box using POP or IMAP protocols
  - **Submission Agent (SA)**
    - Route mails to local MTA

Mail System (2)

Exhibit A Mail system components

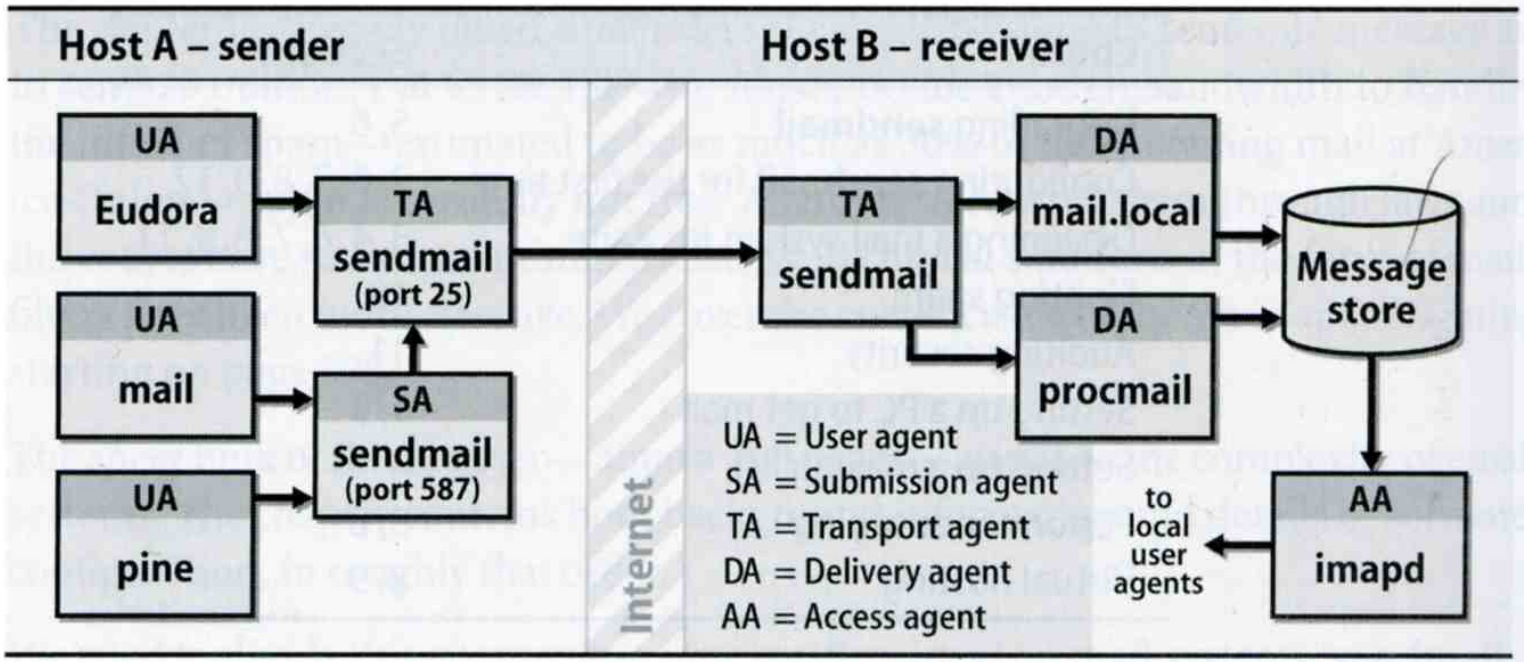

Mail System – The Message Stores

- The place on the local machine where email is stored
  - **Usually the directory: /var/mail or /var/spool/mail**
    - Users' mails are stored in files named with each user's login name
      - Such as /var/mail/tytsai
    - Permission "775" and root:mail as the owner and group owner
      - drwxrwxr-x 2 root mail 512 Dec 16 15:51 mail/
  - **Using database**
    - When the organization is large or for ISP with millions of customers

Mail System – The User Agent (1)

- Help user read and compose mails
  - **UA must know mail format**
    - Originally: Text only
    - Now: MIME
  - **MIME (Multipurpose Internet Mail Extensions)**
    - Include several types of content that can be encoded in the mail, such as image, video, …

Mail System – The User Agent (2)

- **Popular Mail User Agents**

User Agent	System Config.	User Config.	MIME	POP	IMAP	SMTP
bin/mail	mail.rc	.mailrc				
pine	pine.conf	.pinerc	✓	✓	✓	✓
elm	lib/elm.rc	.elm/elmrc	✓	✓	✓	
mutt	/etc/Muttrc	.muttrc	✓	✓	✓	
Netscape	-	-	✓	✓	✓	✓
Eudora	-	-	✓	✓	✓	✓
Outlook Ep.	-	-	✓	✓	✓	✓

Mail System – The Transport Agent (1)

- Route mails among machines
  - **Accept mail from UA, examine the recipients' addresses, and delivery the mail to the correct host**
  - **Protocols**
    - SMTP (Simple Mail Transport Protocol)
      - RFC 821
    - ESMTP (Extended SMTP)
      - RFC 1869, 1870, 1891, 1985
  - **Popular transport agents**
    - sendmail
    - Portfix

Mail System – The Transport Agent (2)

> Conversation between TAs

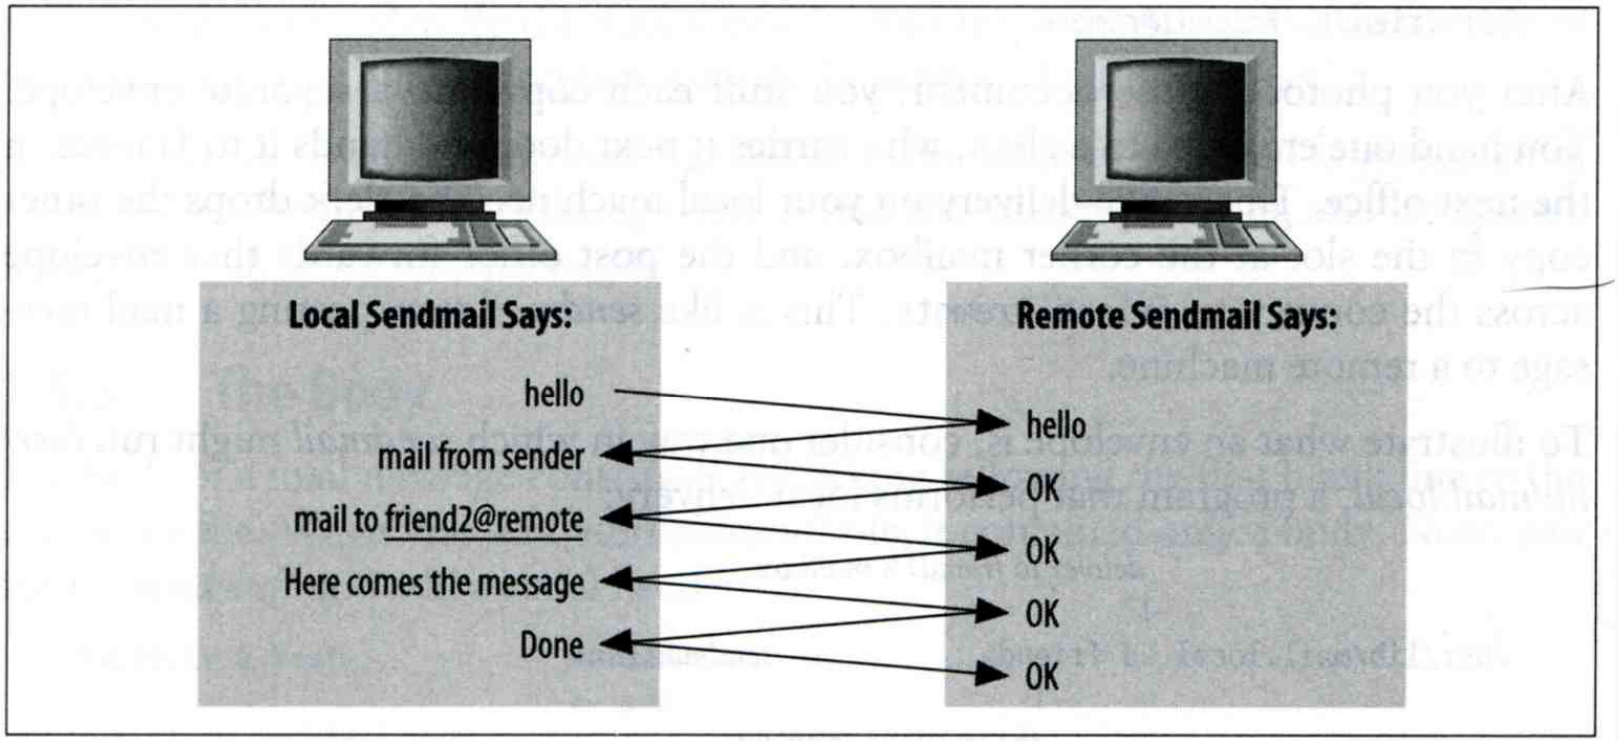

Figure 1-2. A simplified conversation

Mail System – The Transport Agent (3)

> Taking in SMTP

```
tytsai@tybsd:/etc> telnet localhost 25
Trying 127.0.0.1...
Connected to localhost.csie.nctu.edu.tw.
Escape character is '^]'.
220 tybsd.csie.nctu.edu.tw ESMTP Sendmail 8.13.1/8.13.1; Mon, 20 Dec 2004 15:49:00 GMT
HELP
214-2.0.0 This is sendmail version 8.13.1
214-2.0.0 Topics:
214-2.0.0       HELO    EHLO    MAIL    RCPT    DATA
214-2.0.0       RSET    NOOP    QUIT    HELP    VRFY
214-2.0.0       EXPN    VERB    ETRN    DSN     AUTH
214-2.0.0       STARTTLS
214-2.0.0 For more info use "HELP <topic>".
214-2.0.0 To report bugs in the implementation send email to
214-2.0.0       sendmail-bugs@sendmail.org.
214-2.0.0 For local information send email to Postmaster at your site.
214 2.0.0 End of HELP info
EHLO tybsd
250-tybsd.csie.nctu.edu.tw Hello localhost [127.0.0.1], pleased to meet you
...
quit
221 2.0.0 tybsd.csie.nctu.edu.tw closing connection
Connection closed by foreign host.
```

Mail System – The Delivery Agent

- Place mails in users' mail boxes
  - **Accept mail from MTA and deliver the mail to the local recipients**
  - **Type of recipients**
    - User
    - Program, such as
      - mail.local
      - procmail
  - **mail.local**
    - Read the stdin up to an EOF and appends it to each user's mail file
  - **procmail**
    - Do something between mail coming in and stored in mail box
    - CSIE: Help → 5 → 2 → 8

Mail System – The Access Agent

- Help user download mail from server
  - **Protocols**
    - IMAP (Internet Message Access Protocol)
    - POP (Post Office Protocol)

Mail System – The Submission Agent

- Route mails to local MTA
  - **Typical works that a MTA must do:**
    - Ensuring that all hostname are fully qualified
    - Modifying headers
    - Logging errors
    - …
  - **RFC2476 introduces the idea of splitting MTA**
    - Let SA to share the load

Components of a mail (1)

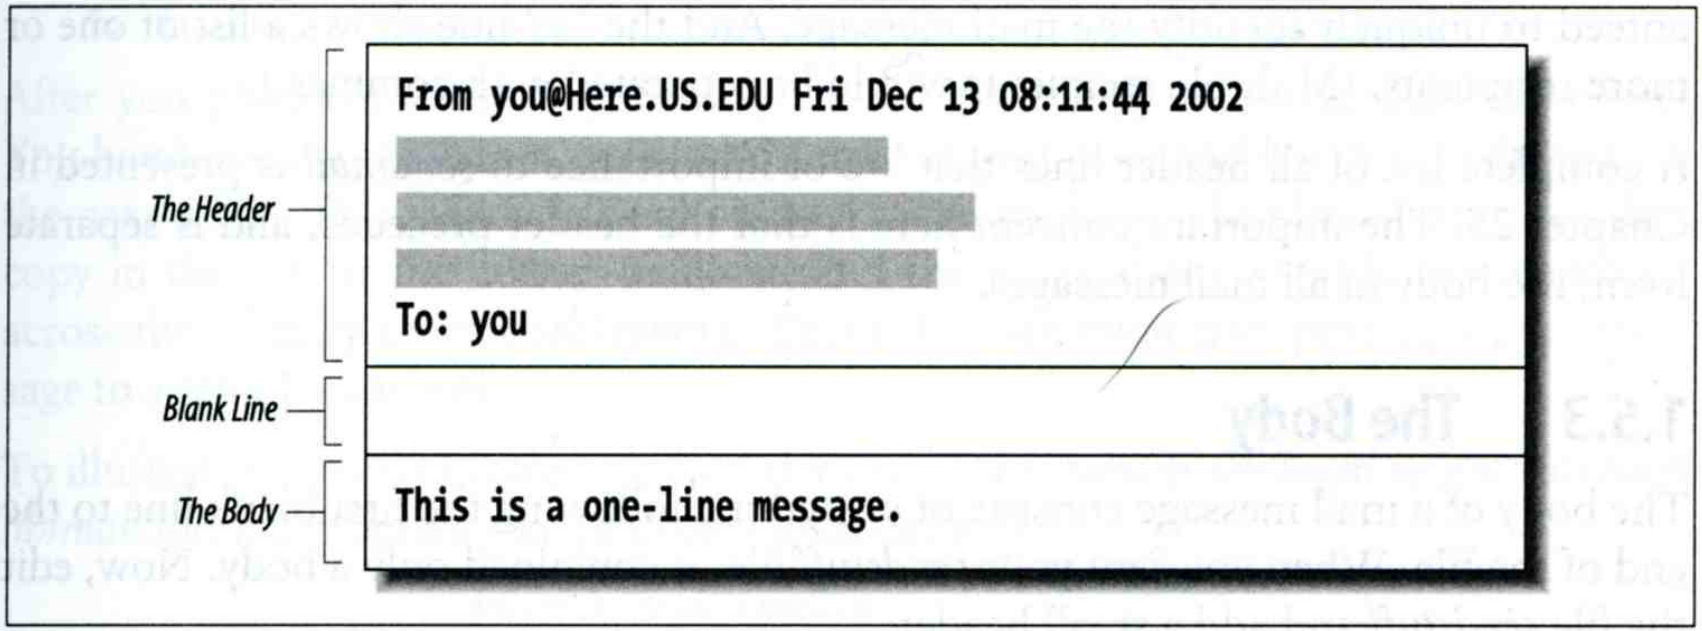

Figure 1-1. Every mail message is composed of a header and a body

Components of a mail (2)

- Three major components
  - **The envelope**
    - Invisible to users
    - Determine where the message should be delivered, or to whom it should be returned
  - **The headers**
    - Information about the messages, defined in RFC822
      - From, To, Date, Time, MTA, …
  - **The message body**
    - Plain text only
    - Various MIME contents are encoded as printable characters using radix-64 algorithm

Mail Addressing (1)

> Two kinds of email addresses:

- **Route based address**
  - Message will travel through several intermediate hosts to the destination
  - Format: host!path!user
    - > Ex: castle!sun!sierra!hplabs!ucbvax!winsor
    - > This mail is sent from “castle” host to the user “winsor” at “ucbvax” host
- **Location independent address**
  - Simply identify the final destination
  - Format: user@host.domain
    - > Ex: tytsai@tybsd.csie.nctu.edu.tw

Mail Addressing (2)

- Alias
  - **Map a username to something else, such as**
    - To a group of users
      - *Ex: cp86 → tytsai, chiaming, wujc, …*
    - To the same user at different machine
      - *Ex: tytsai@csie.nctu.edu.tw → tytsai@ee.nctu.edu.tw*
    - To another user
      - *Ex: admin@csie.nctu.edu.tw → tytsai@csie.nctu.edu.tw*

Mail Addressing (3)

- Where to send the mail?
  - **When you want to send the mail to tytsai@csie.nctu.edu.tw, the MTA will:**
    - First, lookup up the mail exchanger of "csie.nctu.edu.tw"
      - > % dig mx csie.nctu.edu.tw
      - > If there is any servers, choose the higher preference one
      - > If this preferred one can not be connected, choose another
      - > If all the mx servers can not be connected, mail it directly to the host
      - > Ex:

```
tytsai@ccduty:~/Mail/2004-12-18> dig mx csie.nctu.edu.tw

;; ANSWER SECTION:
csie.nctu.edu.tw.      8640    IN      MX      1 mx3.csie.nctu.edu.tw.
csie.nctu.edu.tw.      8640    IN      MX      5 mx1.csie.nctu.edu.tw.
csie.nctu.edu.tw.      8640    IN      MX      5 mx2.csie.nctu.edu.tw.
```

Mail Addressing (4)

- Why using “Mail eXchanger”?
  - **We can centralize all the mail tasks to group of servers**
  - **Multiple mail exchangers make it more robust**

Mail Headers (1)

- Defined by RFC822 which is obsoleted by RFC2822
  - **Mail reader will hide some uninteresting header information**

```
Date: Sat, 18 Dec 2004 11:50:13 +0000
From: Tsung-Yi Tsai <tytsai@tybsd.csie.nctu.edu.tw>
To: Tsung-Yi Tsai <tytsai@tybsd.csie.nctu.edu.tw>
Subject: 哈哈
User-Agent: Mutt/1.4.2.1i
```

Mail Headers (2)

```
From tytsai@tybsd.csie.nctu.edu.tw Sat Dec 18 11:50:13 2004
Return-Path: <tytsai@tybsd.csie.nctu.edu.tw>
Received: from tybsd.csie.nctu.edu.tw (localhost [127.0.0.1])
        by tybsd.csie.nctu.edu.tw (8.13.1/8.13.1) with ESMTP id iBIBoDmY029830
        for <tytsai@tybsd.csie.nctu.edu.tw>; Sat, 18 Dec 2004 11:50:13 GMT
        (envelope-from tytsai@tybsd.csie.nctu.edu.tw)
Received: (from tytsai@localhost)
        by tybsd.csie.nctu.edu.tw (8.13.1/8.13.1/Submit) id iBIBoDjv029829
        for tytsai@tybsd.csie.nctu.edu.tw; Sat, 18 Dec 2004 11:50:13 GMT
        (envelope-from tytsai)
Date: Sat, 18 Dec 2004 11:50:13 +0000
From: Tsung-Yi Tsai <tytsai@tybsd.csie.nctu.edu.tw>
To: Tsung-Yi Tsai <tytsai@tybsd.csie.nctu.edu.tw>
Subject: 哈哈
Message-ID: <20041218115013.GA29818@tybsd.csie.nctu.edu.tw>
Mime-Version: 1.0
Content-Type: text/plain; charset=utf-8
Content-Disposition: inline
Content-Transfer-Encoding: 8bit
User-Agent: Mutt/1.4.2.1i
Status: RO
Content-Length: 10
Lines: 1

現在呢
```

Mail Headers (3)

- Example
  - **User "eric" on "knecht.sendmail.org" sends a email to user "evi" on "anchor.cs.colorado.edu"**
    - % dig mx anchor.cs.colorado.edu
      > mroe.cs.colorado.edu

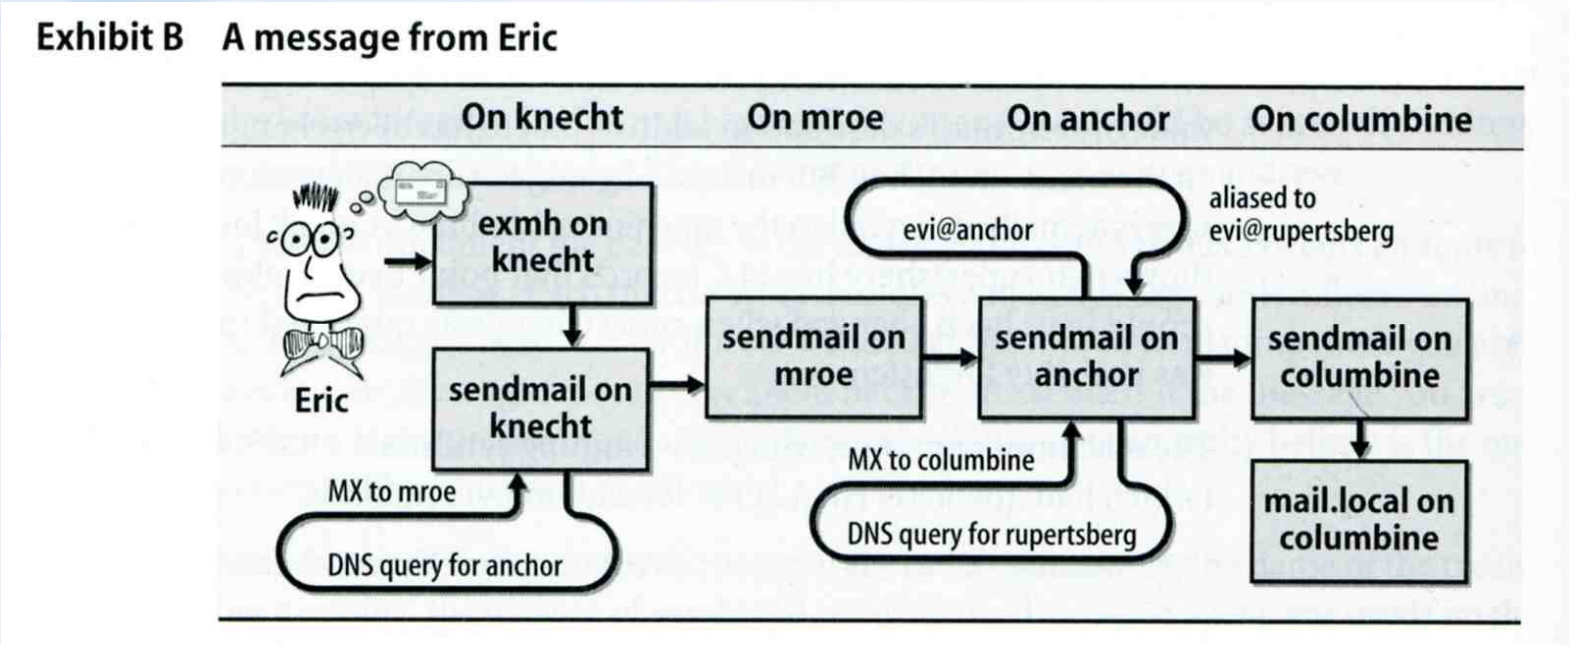

Mail Headers (4)

- Headers in this example
  - **From eric@knecht.sendmail.org**
    - Added by mail.local when the mail is put in user's mailbox
    - Used to separate message boundary
  - **Return-Path: eric@knecht.sendmail.org**
    - Used to send the error message to this address
    - May be different to the "From" address
  - **Received: from knecht.sendmail.org (localhost [127.0.0.1]) by knecht.sendmail.org (8.9.3/8.9.2) with ESMTP id GAA18984; Fri 1 Oct 1999 06:04:02 -800 (PST)**
    - Every machine that is ever processed this mail will add a "Received" record in top of headers
      - Sending machine
      - Receiving machine
      - Version of sendmail in receiving machine
      - Message unique identifier in receiving machine
      - Date and time

Mail Headers (5)

- Received: from anchor.cs.Colorado.EDU (root@anchor.cs.colorado.edu [128.138.242.1]) by columbine.cs.colorado.edu (8.9.3/8.9.2) with ESMTP id HAA21741 for <evi@rupertsberg.cs.colorado.edu>; Fri, 1 Oct 1999 07:04:25 -0700 (MST)
- Received: from more.cs.colorado.edu (more.cs.colorado.edu [128.138.243.1]) by anchor.cs.colorado.edu (8.9.3/8.9.2) with ESMTP id HAA26176 for <evi@anchor.cs.colorado.edu>; Fri, 1 Oct 1999 07:04:24 -0700 (MST)
- Received: from knecht.sendmail.org (knecht.sendmail.org [209.31.233.160]) by more.cs.colorado.edu (8.9.3/8.9.2) with ESMTP id HAA09899 fro <evi@anchor.cs.colorado.edu>; Fri, 1 Oct 1999 07:04:23 -700 (MST)
- Received: from knecht.sendmail.org (localhost [127.0.0.1]) by knecht.sendmail.org (8.9.3/8.9.2) with ESMTP id GAA18984; Fri 1 Oct 1999 06:04:02 -800 (PST)

Exhibit B A message from Eric

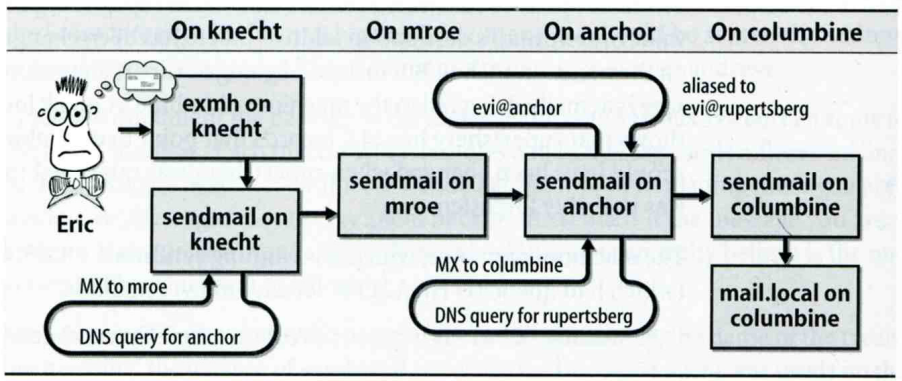

Mail Headers (6)

- **Message-Id: <199910011404.GAA18984@knecht.sendmail.org)**
  - Add by sender's MTA
- **X-Mailer: exmh version 2.0.2 2/24/98**
  - MUA
  - Non-standard header information
- **To: Evi Nemeth <evi@anchor.cs.colorado.edu>**
- **Subject: Re: hi**
- **Date: Fri, 1 Oct 1999 06:04:02 -800**

Mail System Architecture

- Components in a mail system architecture
  - **Mail servers for incoming and outgoing mails**
  - **Mail home**
  - **IMAP or POP to integrate PC and remote clients**
- Simplest architecture
  - **Only one machine**
    - This machine has sendmail to let you send and receive mail
    - This machine is also the mailbox home
    - This machine also provides IMAP or POP to let you download mail from PC

Mail System Architecture –
Scalable architecture for medium sites

- Centralize
  - **At least one machine for incoming message and**
    - Mail home can be the same host or another one
  - **At least one machine for outgoing message**
    - Each host run MSA and forward mail to the same mail-out server or send the mail directly

Exhibit C Mail system architecture

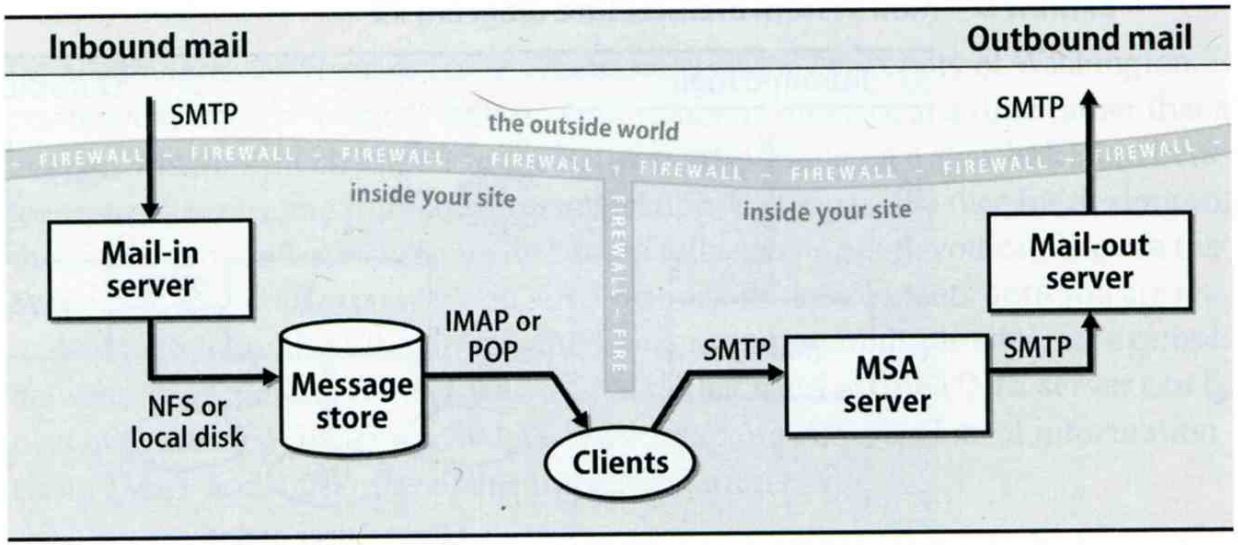

Mail Alias

- Several mechanisms to define aliases:
  - **Traditional method: in files**
  - **Traditional method with NIS**
  - **LDAP (Light-weight Directory Access Protocol)**
- When the sendmail wants to resolve name
  - **File-based method**
    - sendmail looks up files to resolve it by itself
  - **LDAP-based method**
    - sendmail call LDAP server to resolve the name and return the results

Mail Alias – Traditional aliasing mechanism (1)

- Aliases can be defined in three places
  - **In MUA's configuraiton file**
    - Read by MUA and expand the alias before injecting the message into the mail system
  - **In the system-wide /etc/mail/aliases file**
    - Read by MTA
    - The path to the system-wide alias file can be specified in sendmail's configuration file
  - **In user's forwarding file, ~/.forward**
    - Read by MTA after system-wide alias file

Mail Alias – Traditional aliasing mechanism (2)

- The format of an entry in aliases file
  1. **Local-name: recipient1,recipient2,…**
  - **Ex:**
    - admin: tytsai,jwwang
    - tytsai: tytsai@tybsd.csie.nctu.edu.tw

  2. **Local-name: :include:another-file**
  - **Ex:**
    - bsdTA: :include:/usr/local/mail/bsdTA

Contents of bsdTA

tytsai
kctao
chiahung
lwhsu
chlo
hwchu
liling

Mail Alias – Traditional aliasing mechanism (3)

3. **Local-name: absolute-path-file**
   - **Mails will be appended to this file**
   - **Ex:**
     - complaints: /dev/null
     - troubles: trouble_admin,trouble_log
     - trouble_admin: :include:/usr/local/mail/troadm
     - trouble_log: /usr/local/mail/logs/troublemail

4. **Local-name: "|program-path"**
   - **Route mail to stdin of program**
   - **Ex:**
     - autoftp: "|/usr/local/bin/ftpserver"

Mail Alias – Traditional aliasing mechanism (4)

- The hashed aliases DB
  - **/etc/mail/aliases is the plaintext aliases information**
  - **/etc/mail/aliases.db is the hashed version for efficiency**
  - **Use "newaliases" command to rebuild the hashed version when you change the aliases file**

Mail Alias – Traditional aliasing mechanism (5)

- User maintainable forwarding file
  - **In ~/.forward**
  - **Format: comma-separated**
  - **Ex:**
    - tytsai@gmail.com
    - \tytsai, tytsai@gmail.com, tytsai@yahoo.com.tw
  - **Must be owned by user and with permission of 600**
    - The path to .forward file should be writable only to user

Mail Alias – Traditional aliasing mechanism (6)

> Alias must

- **postmaster and MAILER-DAEMON**
  - Mail system maintainer
- **bin, sys, daemon, nobody, …**
  - System accounts (root)
- **root**
  - forward root mail to the administrator (.forward)

```
MAILER-DAEMON: postmaster
postmaster: root
bin:    root
bind:   root
daemon: root
games:  root
kmem:   root
mailnull: postmaster
nobody: root
operator: root
…
```