

# E-Mail System

lctseng (2020-2022, CC-BY)

? (?-2019)

國立陽明交通大學資工系資訊中心

Computer Center of Department of Computer Science, NYCU

# Components of an E-Mail (1)

- You can really see ...
  - Headers, which can be forged, altered, etc.
  - Body

The Header =>	Date: Mon, 22 Mar 2021 09:15:04 +0800 (CST) From: NCTU CSCC Help <help@cs.nctu.edu.tw> To: lctseng@cs.nctu.edu.tw Subject: [CSCC] Test Mail
Blank Line =>	
The Body =>	This is a test mail.

# Components of an E-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 RFC2822
      - Date, From, To, Content-Type, charset
      - Content-Length, MessageID, ...
      - **No checking consistent "To" in envelope and header**
  - The message body
    - Plain text only
    - Various MIME contents (attachments)
      - 7bit, quoted-printable, base64
      - 8bit, binary



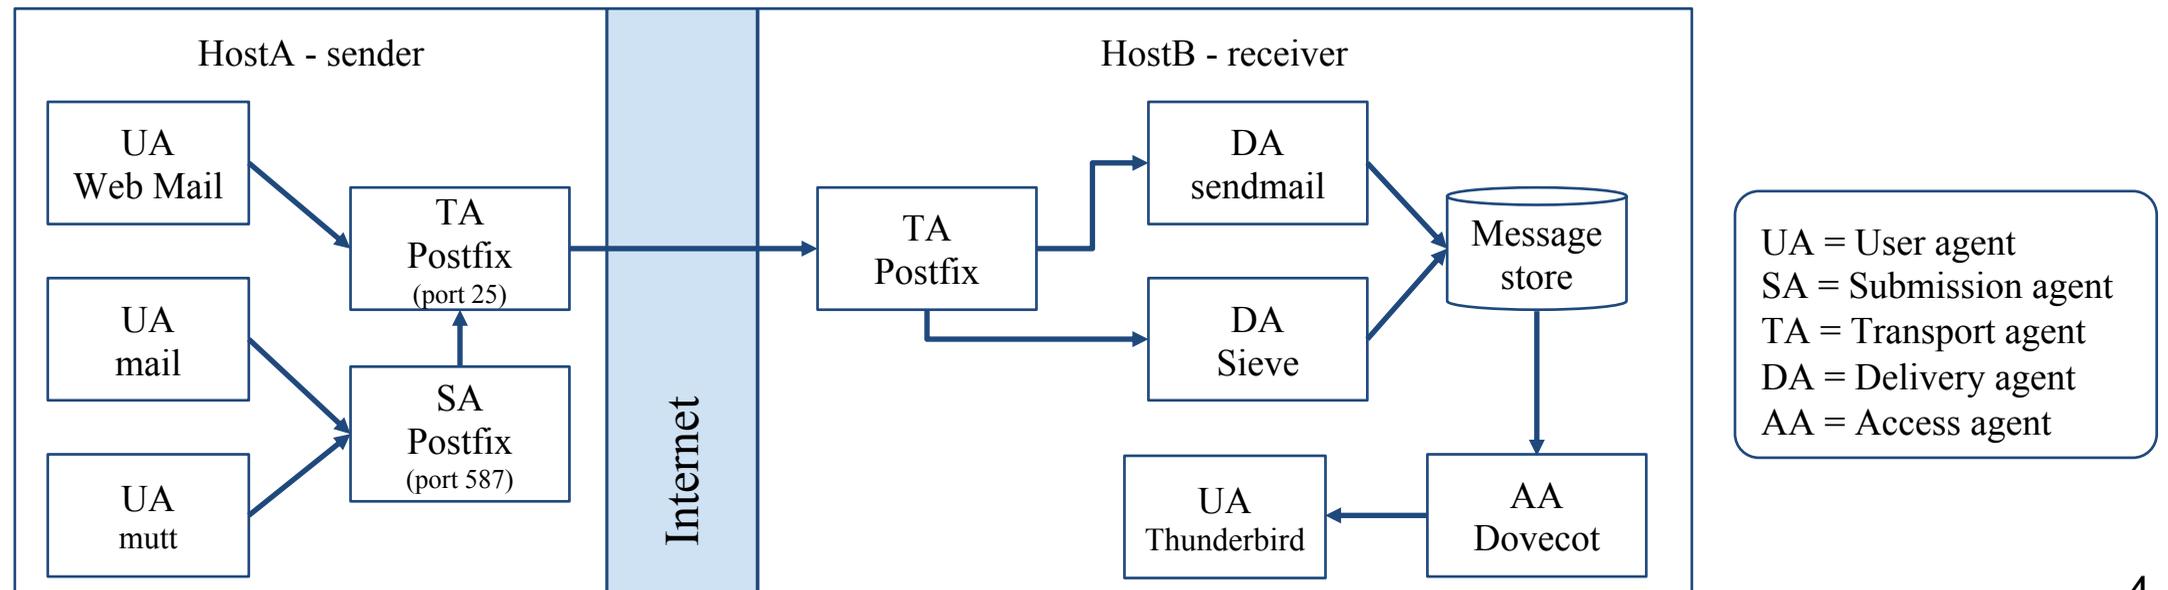
Mail system rely on this

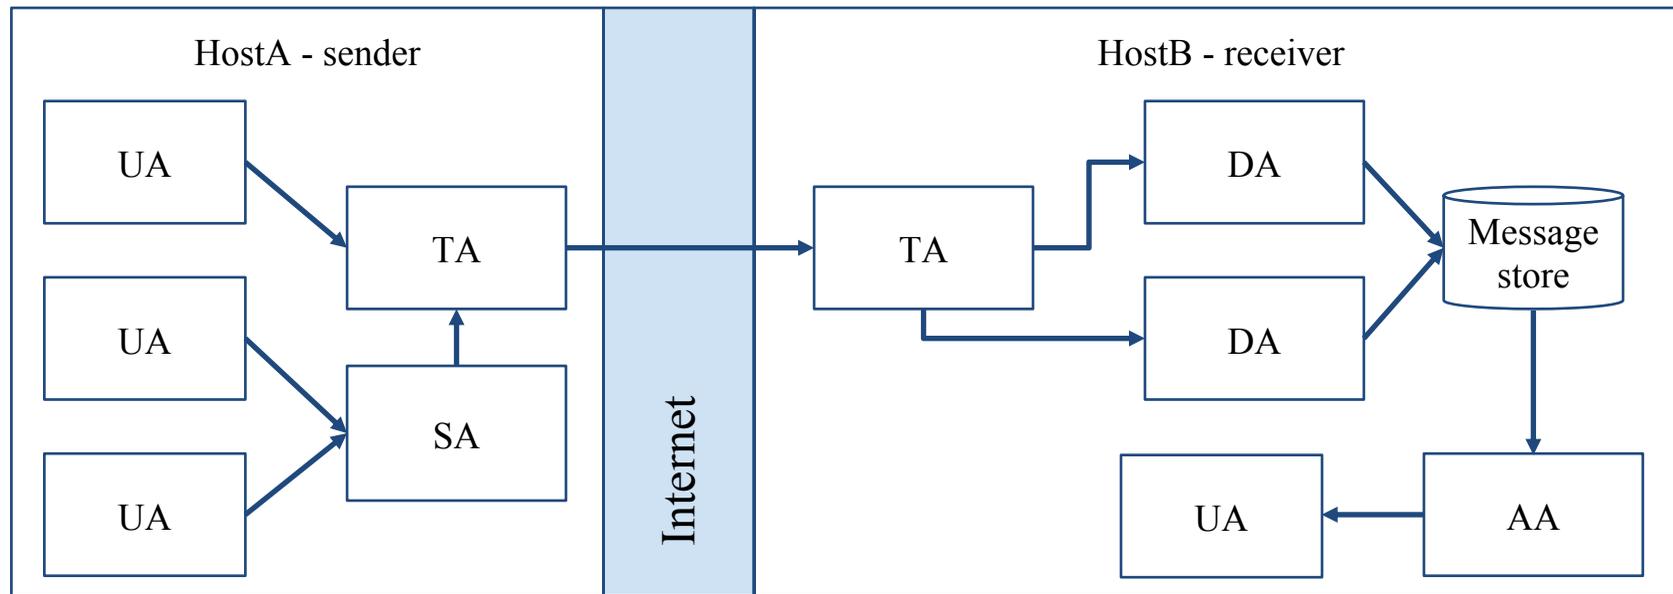


Can be anything!

# Mail System

- Major components
  - Mail User Agent (MUA)
    - Help user read and compose mails
  - Submission Agent (SA)
    - Route mails to local MTA
  - 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 POP3 or IMAP protocols



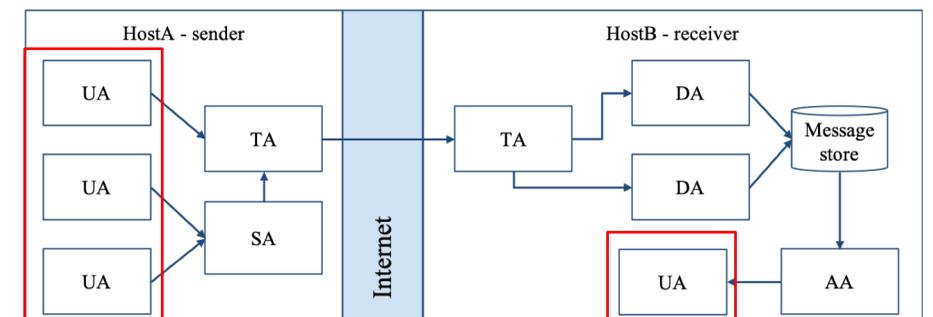


# Mail System – The User Agent (1)

- 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
  - image, video, **virus**, ...



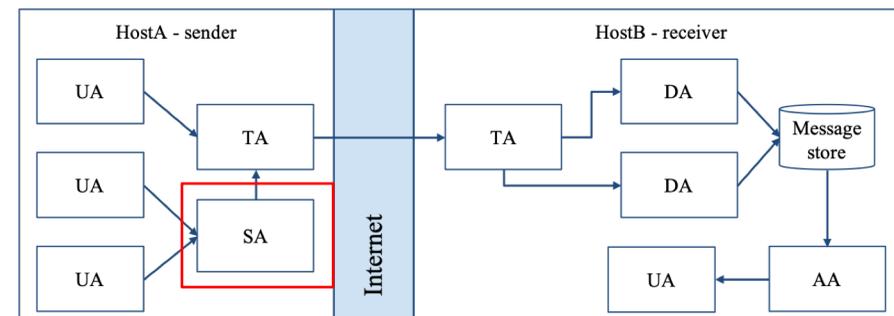
# Mail System – The User Agent (2)

- Popular Mail User Agents

User Agent	System Config.	User Config.	MIME	POP	IMAP	SMTP
mail	mail.rc	.mailrc				
mutt	/etc/muttrc	.muttrc	○	○	○	○
Netscape	-	-	○	○	○	○
Outlook Ep.	-	-	○	○	○	○
MS Outlook	-	-	○	○	○	○
Thunderbird	-	-	○	○	○	○
In Smartphones	-	-	○	○	○	○

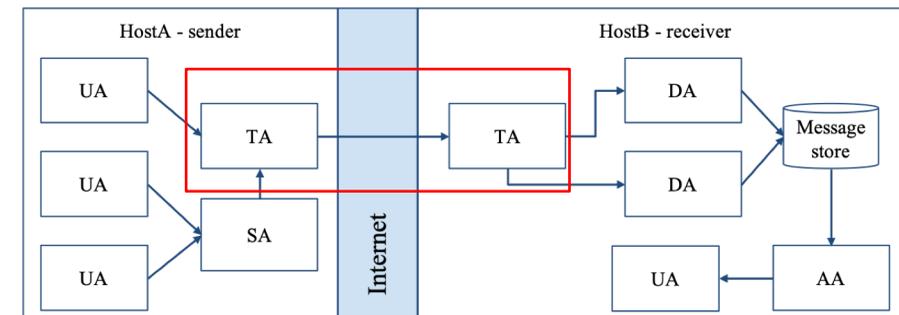
# Mail System – The Submission Agent

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



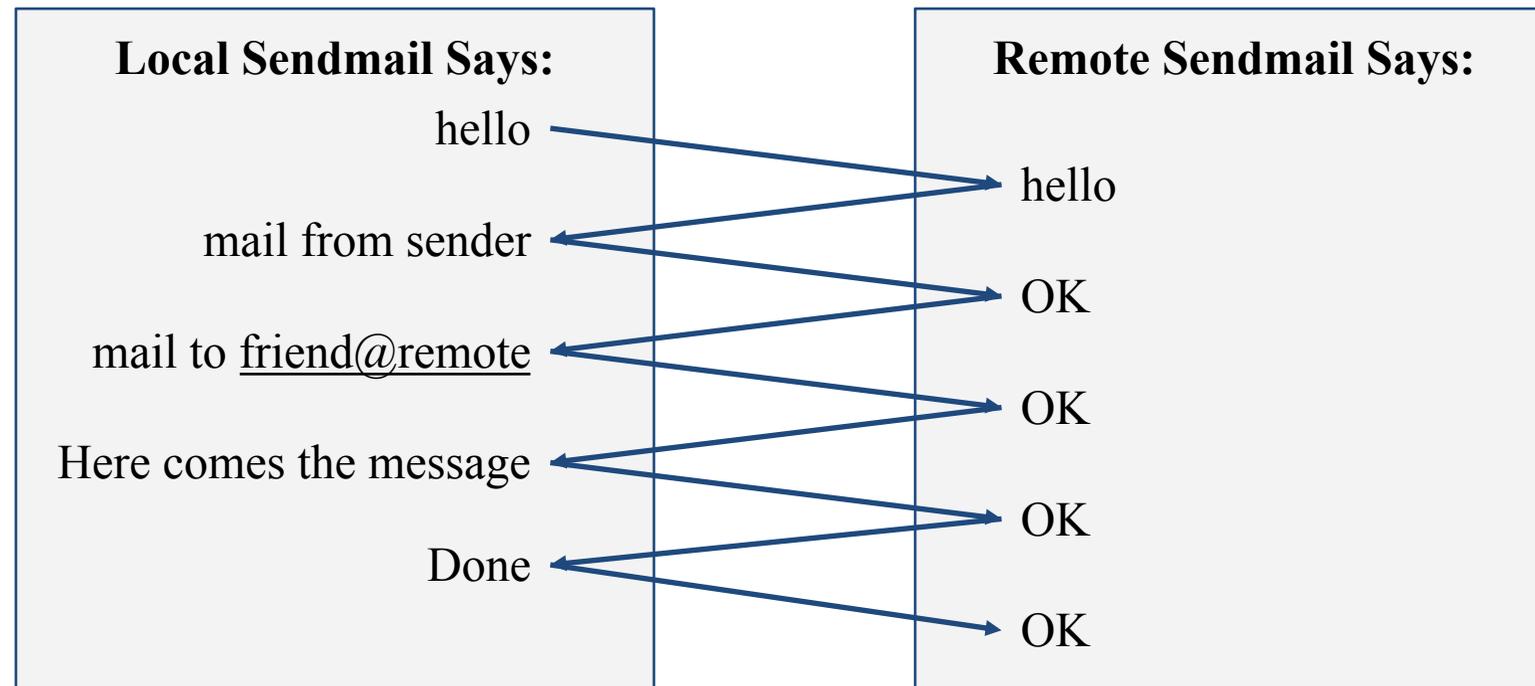
# Mail System – The Transport Agent (1)

- 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 2821 => ... => 5321 (2008)
  - Popular transport agents
    - sendmail <http://www.sendmail.org/>
    - Postfix <http://www.postfix.org/>
    - exim, qmail, ...



# Mail System – The Transport Agent (2)

- Conversation between MTAs
  - Threat of eavesdropping



# Mail System – The Transport Agent (3)

- Protocol: SMTP

```
$ telnet csmailgate 25
Trying 140.113.235.103...
Connected to csmailgate.
Escape character is '^]'.
220 csmailgate.cs.nctu.edu.tw ESMTP Postfix
ehlo bsd4.cs.nctu.edu.tw
250-csmailgate.cs.nctu.edu.tw
250-PIPELINING
250-SIZE 204800000
250-VRFY
250-ETRN
250-ENHANCEDSTATUSCODES
250-8BITMIME
250 DSN
```

```
mail from: <alice@cs.nctu.edu.tw>
250 2.1.0 Ok
rcpt to: <bob@cs.nctu.edu.tw>
250 2.1.5 Ok
data
354 End data with <CR><LF>.<CR><LF>
From: haha <devnull@cs.nctu.edu.tw>
To: admin@hinet.net

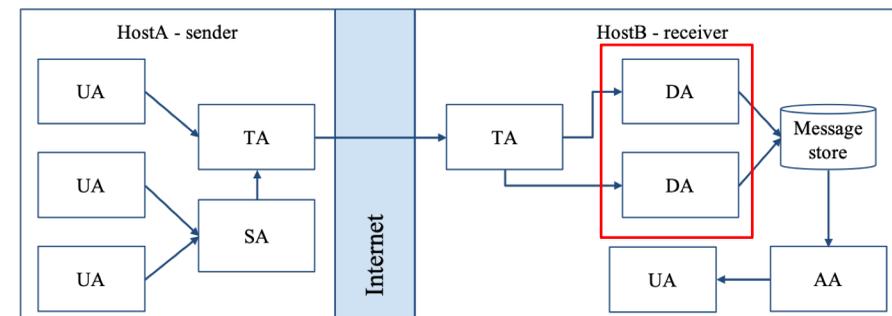
hehe... I spammed you!
.
250 2.0.0 Ok: queued as 81BD4FB4
quit
221 2.0.0 Bye
Connection closed by foreign host.
```

```
From: haha <devnull@cs.nctu.edu.tw>
To: admin@hinet.net
Message-Id: <20120501070002.81BD4FB4@csmailgate.cs.nctu.edu.tw>
Date: Tue, 1 May 2012 14:59:53 +0800 (CST)

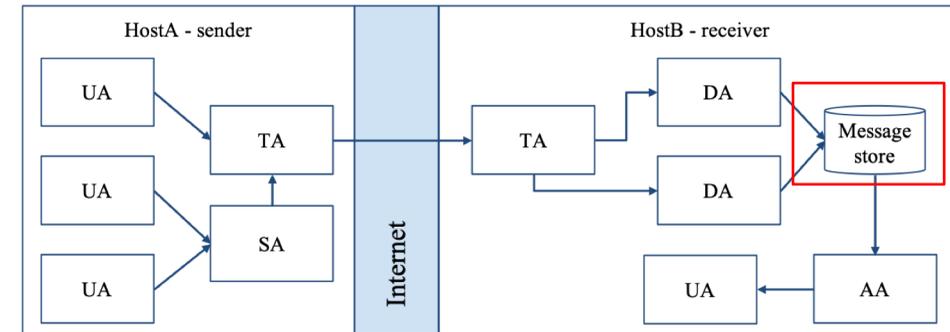
hehe... I spammed you!
```

# Mail System – The Delivery Agent

- Place mails in users' mailboxes
  - Accept mail from MTA and deliver the mail to the local recipients
  - Type of recipients
    - User
    - Program
      - Sieve filters, procmail (deprecated), ...
  - Sieve - mail filtering language (RFC 5228)
    - Many implementations
    - Pigeonhole - Sieve implementation
      - [Official documentation](#)



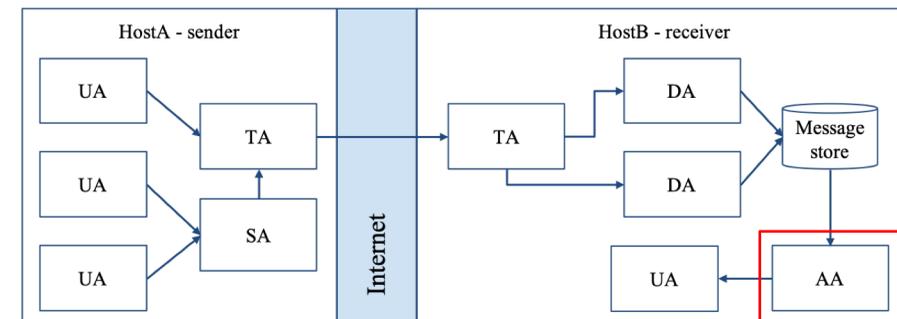
# Mail Storage



- 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
      - Eg. `/var/mail/lctseng`
    - 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
    - Easy to search, categorize

# Mail System – The Access Agent

- Help user download mail from server
  - Protocols
    - IMAP (Internet Message Access Protocol)
    - POP3 (Post Office Protocol – Version 3)



# Mail Addressing – Domain (1)

- Two kinds of email addresses:
  - Route based address (obsolete)
    - Message will travel through several intermediate hosts to the destination
    - Format: host!path!user
      - Ex: castle!sun!sierra!hplabs!alice
      - This mail is sent from "castle" host to the user "alice" at "hplabs" host
  - Location independent address
    - Simply identify the final destination
    - Format: user@host.domain
      - E.g. [ta@nasa.cs.nctu.edu.tw](mailto:ta@nasa.cs.nctu.edu.tw)

# Mail Addressing – Domain (2)

- Where to send the mail?
  - When you want to send a mail to lctseng@cs.nctu.edu.tw, the MTA will:
    - First, lookup up the mail exchanger of "cs.nctu.edu.tw"

```
$ dig mx cs.nctu.edu.tw
```

```
;; ANSWER SECTION:
```

```
cs.nctu.edu.tw.      3600    IN      MX      5 csmx2.cs.nctu.edu.tw.  
cs.nctu.edu.tw.      3600    IN      MX      10 csmx3.cs.nctu.edu.tw.  
cs.nctu.edu.tw.      3600    IN      MX      5 csmx1.cs.nctu.edu.tw.
```

- If there is any servers, try from servers with higher priority (smaller value)
- If no MX records, mail it directly to the host (A record)

# Mail Addressing – Domain (3)

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

# Mail Addressing – Alias

- Alias
  - Map a username to something else
    - Be careful of **mail looping**
- Several mechanisms to define aliases:
  - Traditional method: in files
  - Traditional method with NIS
  - LDAP (Light-weight Directory Access Protocol)
- When the mail server wants to resolve name
  - File-based method
    - Look up files to resolve by itself
  - LDAP-based method
    - 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 configuration 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 DA
    - The path to the system-wide alias file can be specified in mail server's configuration file
  - In user's forwarding file, `~/.forward`
    - Read by DA after system-wide alias file
    - [forward\(5\)](#)

# Mail Alias – Traditional aliasing mechanism (2)

- The format of an entry in aliases file
  1. Local-name: recipient1,recipient2,...
    - Ex:
      - admin: lwhsu,wangth,jnlin
      - lctseng: lctseng@cs.nctu.edu.tw
      - root: ta
  2. Local-name: **:include:**filename
    - Ex:
      - ta: :include:/usr/local/mail/TA

```
lwhsu
fyli
lctseng
jnlin
wangth
pml
```

Contents of TA

# Mail Alias – Traditional aliasing mechanism (3)

- The format of an entry in aliases file
  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"
      - nahw3: "|/home/nahw3/receive.py"

# 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
    - The file read from ":include:" is outside the aliases file

# Mail Alias – Traditional aliasing mechanism (5)

- User maintainable forwarding file
  - In `~/.forward`
  - Format: comma-separated
  - E.g.
    - `lctseng@gmail.com`
    - `\lctseng, lctseng@gmail.com, lctseng@nycu.edu.tw`
  - backslash + username
    - Bypassing further redirection (deliver to mail box directly)
  - 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)

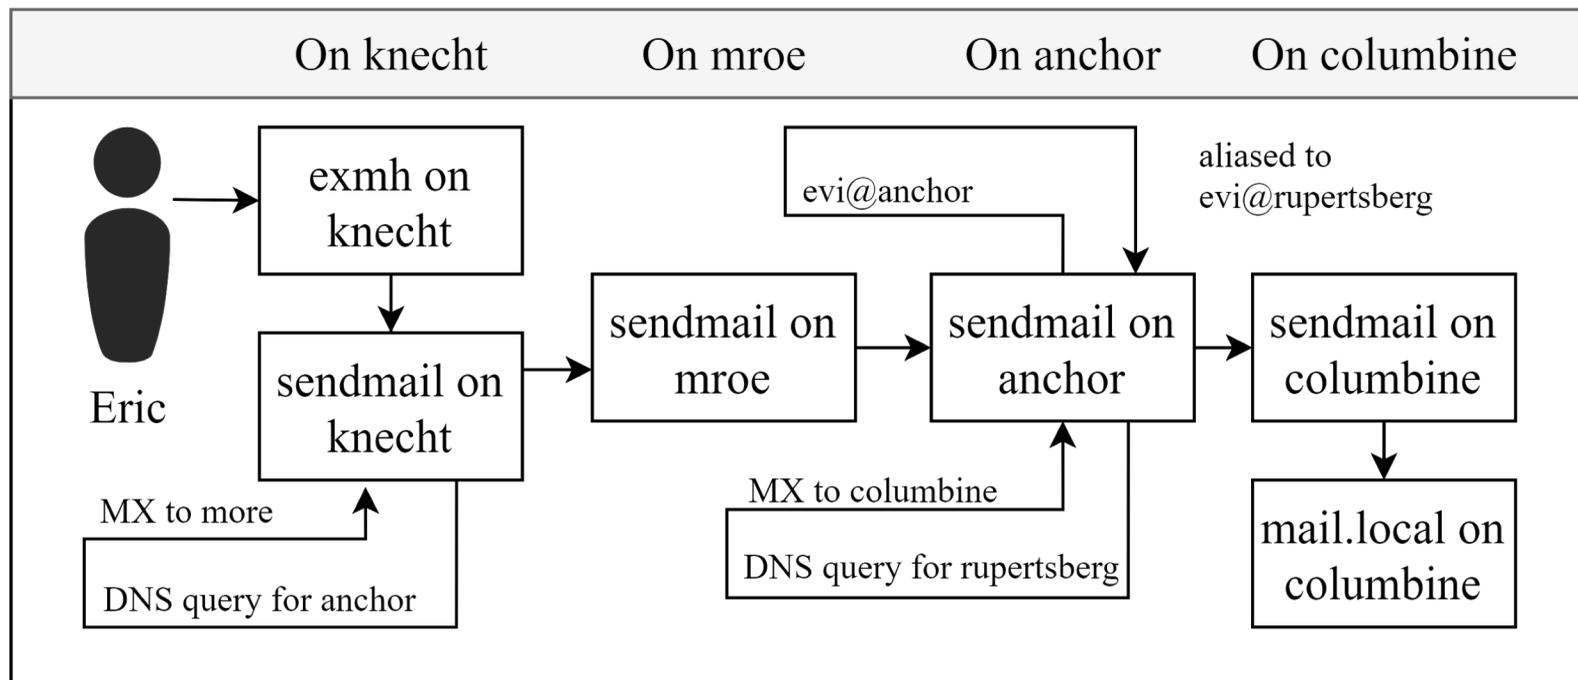
- Alias must
  - postmaster and MAILER-DAEMON
    - Mail system maintainer
  - bin, sys, daemon, nobody, ...
    - System accounts (root)
  - root
    - forward root mail to the administrator
      - /root/.forward
      - aliases

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

# Mail Transport Example

- User [eric@knecht.sendmail.org](mailto:eric@knecht.sendmail.org) sends a email to user [evi@anchor.cs.colorado.edu](mailto:evi@anchor.cs.colorado.edu)
  - \$ dig mx anchor.cs.colorado.edu
    - mroe.cs.colorado.edu

A message from Eric



# Mail Headers (1)

- Defined by RFC2822
  - Mail reader will hide some uninteresting header information

```
Date: Wed, 18 Apr 2007 14:05:04 +0800
From: 大小姐 <lkkg-girl@mail.richhome.net>
Subject: 笑狗好可怕
To: Yung-Hsiang Liu <liuyh@nabsd.cs.nctu.edu.tw>
User-Agent: Mutt/1.5.15 (2007-04-06)
```

你趕快把牠趕跑好不好？

# Mail Headers (2)

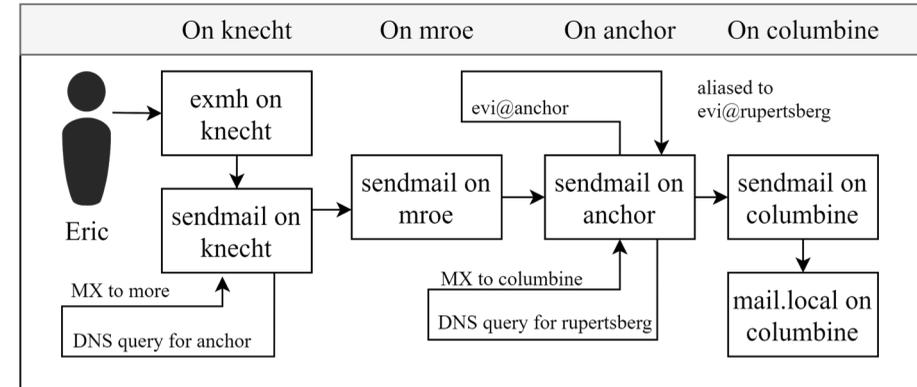
```
From chwong@chbsd.cs.nctu.edu.tw Wed Apr 18 14:07:21 2007
Return-Path: <chwong@chbsd.cs.nctu.edu.tw>
X-Original-To: liuyh@nabsd.cs.nctu.edu.tw
Delivered-To: liuyh@nabsd.cs.nctu.edu.tw
Received: from chbsd.cs.nctu.edu.tw (chbsd.csie.nctu.edu.tw [140.113.17.212])
    by nabsd.cs.nctu.edu.tw (Postfix) with ESMTTP id 22EC73B4D51
    for <chwong@nabsd.cs.nctu.edu.tw>; Wed, 18 Apr 2007 14:07:21 +0800 (CST)
Received: from chbsd.cs.nctu.edu.tw (localhost [127.0.0.1])
    by chbsd.cs.nctu.edu.tw (8.13.8/8.13.8) with ESMTTP id 13I654P3060925
    for <chwong@nabsd.cs.nctu.edu.tw>; Wed, 18 Apr 2007 14:05:04 +0800 (CST)
    (envelope-from chwong@chbsd.cs.nctu.edu.tw)
Received: (from chwong@localhost)
    by chbsd.cs.nctu.edu.tw (8.13.8/8.13.8/Submit) id 13I654AY060924
    for chwong@nabsd.cs.nctu.edu.tw; Wed, 18 Apr 2007 14:05:04 +0800 (CST)
    (envelope-from chwong)
Date: Wed, 18 Apr 2007 14:05:04 +0800
From: =?utf-8?B?5aSn5bCP5aeQ?= <lkk-girl@mail.richhome.net>
To: Yung-Hsiang Liu <liuyh@nabsd.cs.nctu.edu.tw>
Subject: =?utf-8?B?56yR54uX5aW95Y+v5oCV?=
Message-ID: <20070418060503.GA60903@chbsd.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/1.5.15 (2007-04-06)
Status: RO
Content-Length: 23
Lines: 1
```

你趕快把牠趕跑好不好？

# Mail Headers (3)

- Headers in the 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`
    - The envelope "mail from"
    - Used to send the error message to this address
    - May be different to the "From" address in usual header
  - Delivered-To: `evi@rupertsberg`
    - Final envelope "rcpt to"

A message from Eric



# Mail Headers (4)

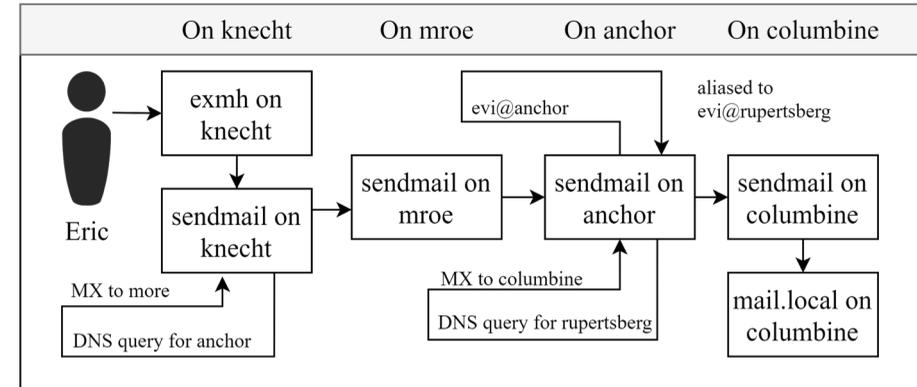
- Headers in the example

- 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

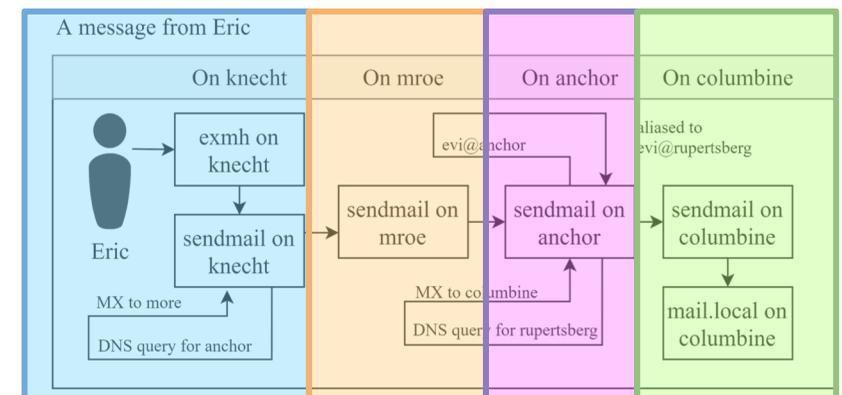
- Sending machine
- Receiving machine
- Mail server software in receiving machine
- Unique queue identifier of mail server in receiving machine
- Date and time

A message from Eric



# 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)



# 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

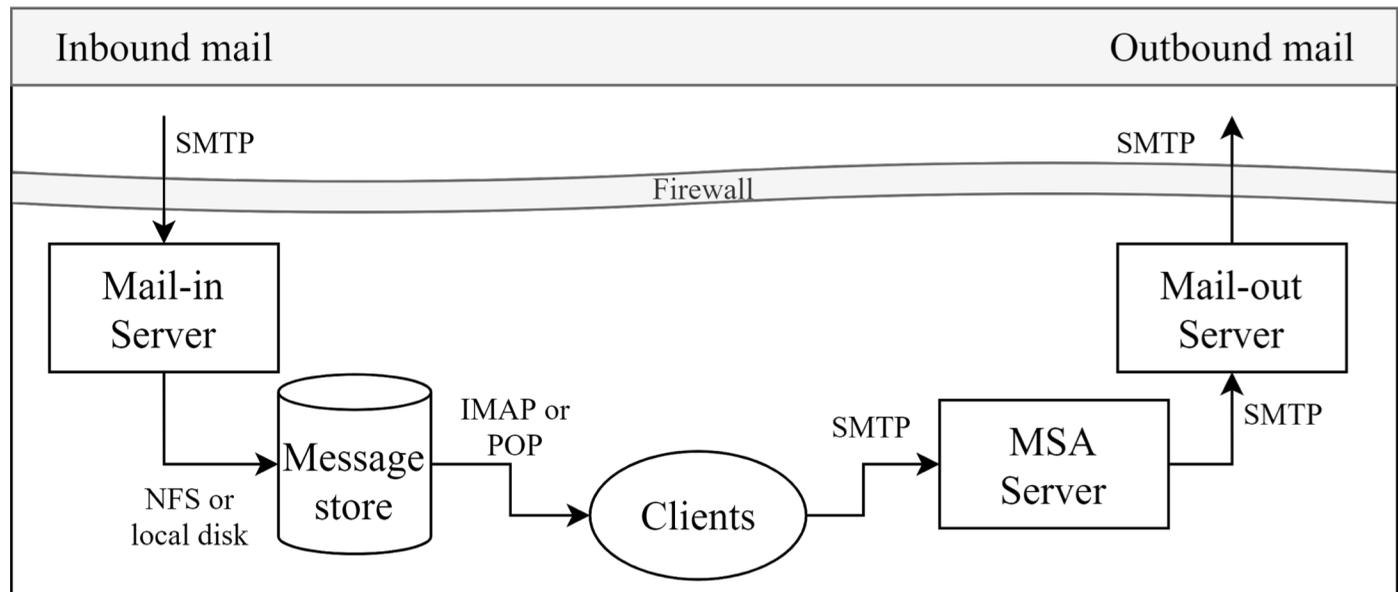
- Simplest architecture
  - Only one machine
    - Has MTA to let you send and receive mail
    - Provides storage for mailboxes
    - Provides IMAP or POP3 to let you download mail from PC
- Components in a mail system architecture
  - Mail servers for incoming and/or outgoing mails
  - Storage for mailboxes
  - IMAP or POP3 to integrate PC and remote clients
    - The issue of file locking

# 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

Mail System architecture



# To, CC, and BCC

- You should always make sure you mail the right people
  - The **To field** is for people that the message directly affects, and that you require actions from.
  - The **CC (or Carbon Copy) field** is for people you want to know about the message, but are not directly involved.
  - The **BCC field (Blind Carbon Copy)** is used when you want other people to receive the message, but you don't want the other recipients to know they got it.
- There are "To" and "CC," but not "BCC" in the email headers.
  - **Why** "No checking consistent 'To' in envelope and header"

# vacation

- [vacation\(1\)](#): E-mail auto-responder
    - returns a message, `~/.vacation.msg` by default
    - `~/.vacation.db`
      - default database file for `db(3)`
    - `~/.vacation.{dir,pag}`
      - default database file for `dbm(3)`
    - `~/.vacation.msg`
      - default message to send
  - Use with [forward\(5\)](#)
    - `\lctsend`, `|/usr/bin/vacation`
- 
- Stores messages people sent to you