



# E-Mail System

tsaimh (2024-2026, CC-BY)

lctseng (2020-2023, CC-BY)

? (?-2019)

# Components of an E-Mail (1)

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

The Header =>

```
Date: Thu, 30 Mar 2023 09:15:04 +0800 (CST)
From: NCTU CSCC Help <help@cs.nctu.edu.tw>
To: tsaimh@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
    - Text, attachments, ...



Mail systems rely on this



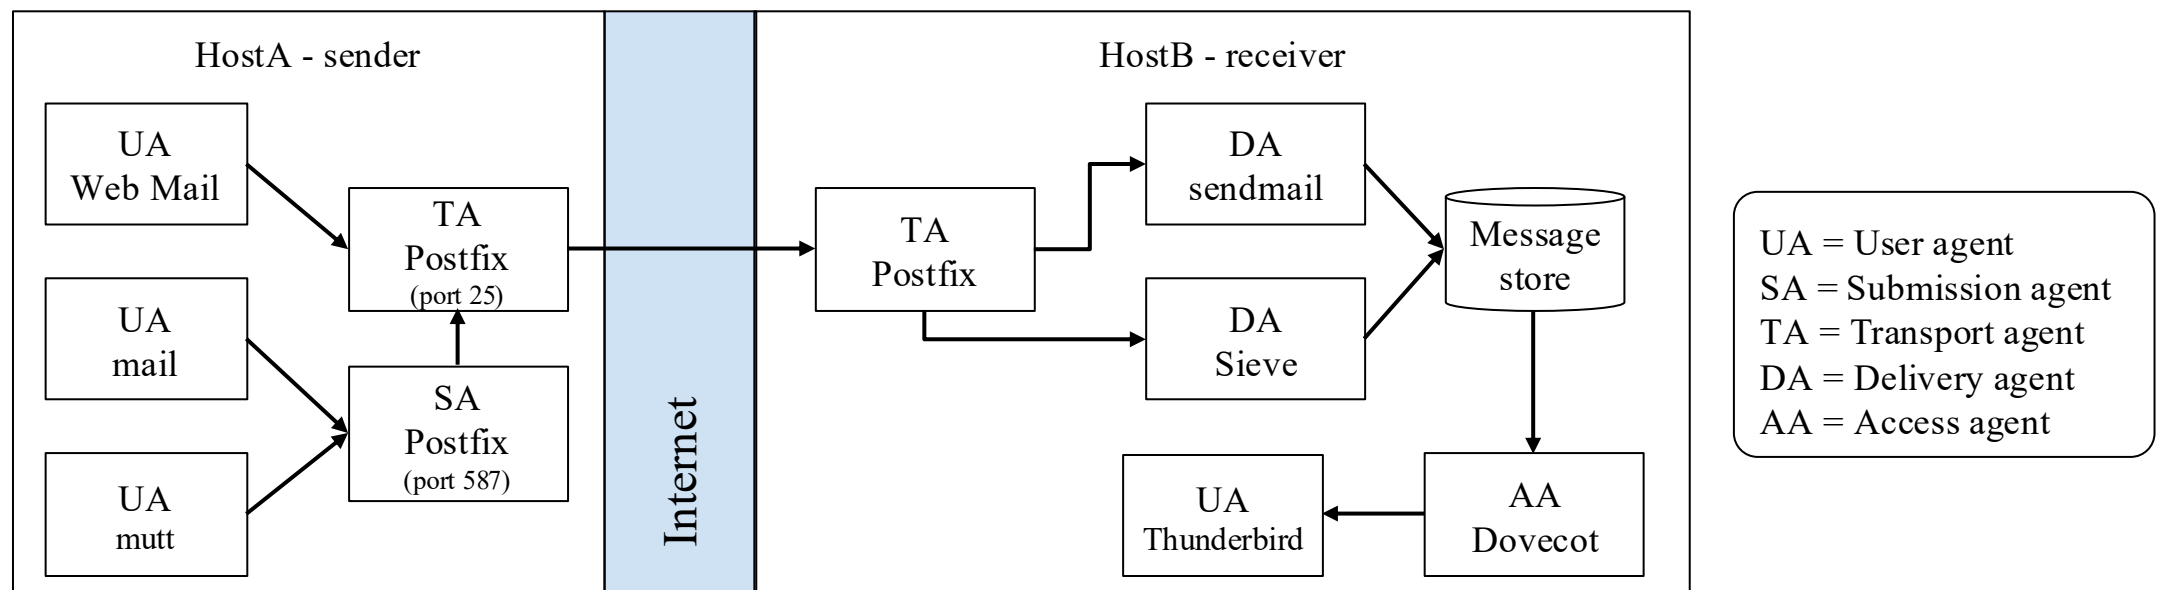
Can be anything!

# Mail System

smtp	25/tcp	mail	/etc/services
submission	587/tcp		# Submission [RFC4409]

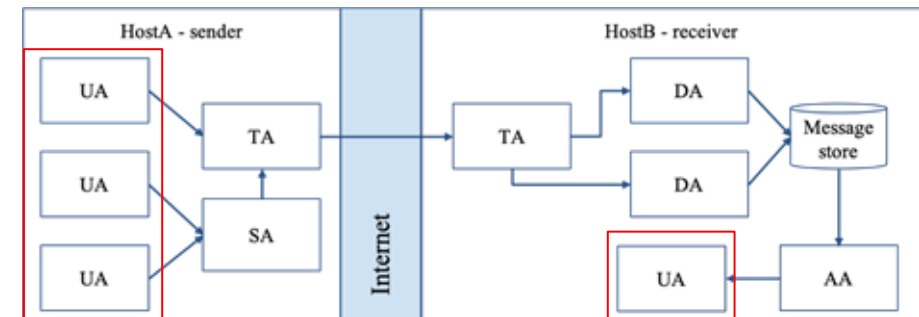
- Major components

- 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' mailboxes
- Access Agent (AA)
  - Connect the user agents to the mailboxes using POP3 or IMAP protocols



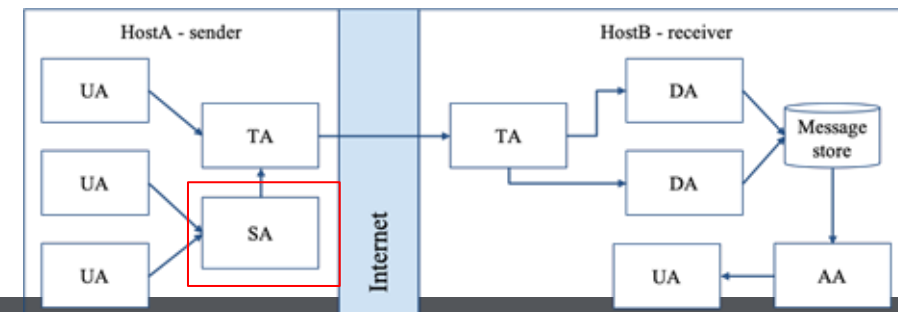
# Mail System – The User Agent

- Help user read and compose mails
  - UA must know mail format
    - Previously: Text only (command line agents)
    - Now: MIME
- ✂ MIME (Multipurpose Internet Mail Extensions)
  - Include several types of content that can be encoded in the mail
    - 7-bit base64, 8-bit binary, ...
    - image, video, **virus**, ...



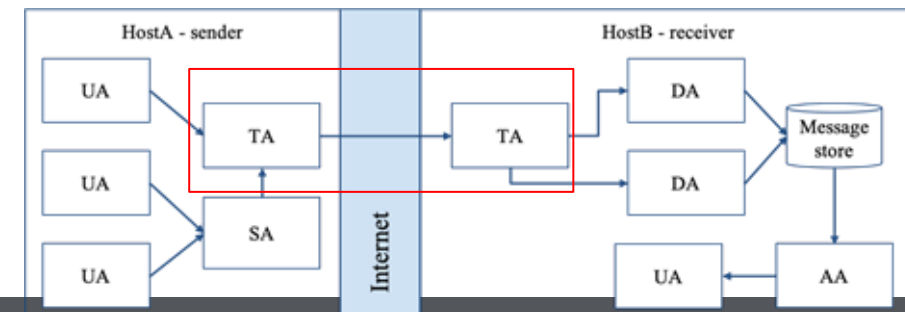
# Mail System – The Submission Agent

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



# Mail System – The Transport Agent


- 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](#) (1982)
    - ESMTP (Extended SMTP)
      - [RFC 2821](#) (2001) => [5321](#) (2008) => [7504](#) (2015)
  - Popular transport agents
    - sendmail <http://www.sendmail.org/>
    - Postfix <http://www.postfix.org/>
    - exim, qmail, ...



# Mail System – Mail (MX) Server Survey

## Basic Technology Breakdown

Description>	Number of Servers	Percent
Total Number of MX Servers Queried	1,100,168	100.0%
Number of Servers that didn't respond	223,341	20.30%
Server didn't open socket	220,688	20.06%
Server didn't provide banner	2,653	0.24%
Server provided banner	876,827	79.70%
Server banner identifies software in use	549,567	49.95%

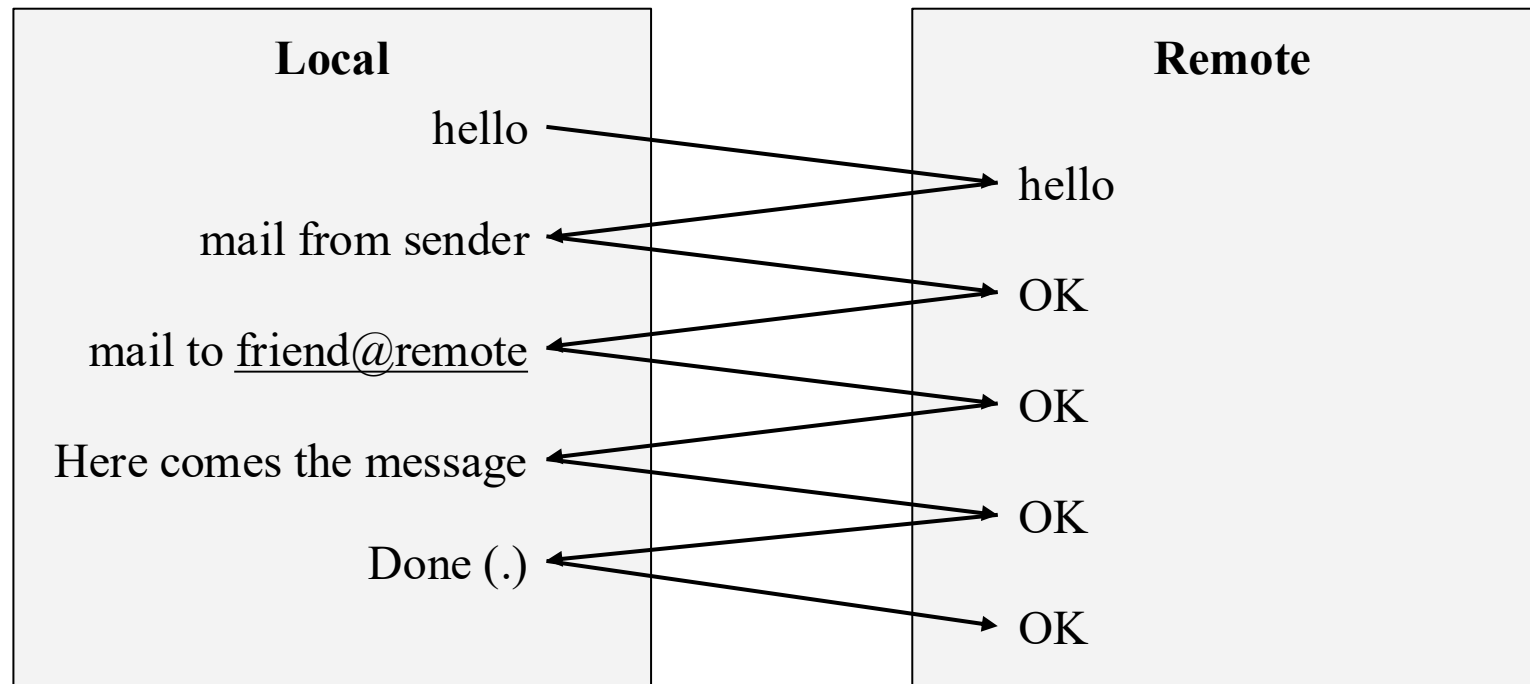


Server Type	Number of Servers	Percent
<a href="#">Exim</a>	302,529	55.05%
<a href="#">Postfix</a>	211,702	38.52%
<a href="#">Sendmail</a>	18,979	3.45%
<a href="#">MailEnable</a>	9,894	1.80%
<a href="#">MDaemon</a>	2,188	0.40%
<a href="#">Microsoft</a>	1,109	0.20%

Source: [https://secure1.securityspace.com/s\\_survey/data/man.202502/mxsurvey.html](https://secure1.securityspace.com/s_survey/data/man.202502/mxsurvey.html) , Statistics at March 1, 2025

# Mail System – Conversation between MTAs

- Conversation between MTAs
  - Threat of eavesdropping



# Mail System – SMTP Interaction

- Protocol: SMTP

```
$ telnet smtp.example.com 25
Trying 140.113.235.103...
Connected to smtp.example.com.
Escape character is '^]'.
220 smtp.example.com ESMTP Postfix
ehlo me.example.com
250-smtp.example.com
250-PIPELINING
250-SIZE 204800000
250-VRFY
250-ETRN
250-ENHANCEDSTATUSCODES
250-8BITMIME
250 DSN
```

```
mail from: <alice@example.com>
250 2.1.0 Ok
rcpt to: <bob@example.com>
250 2.1.5 Ok
data
354 End data with <CR><LF>.<CR><LF>
From: haha <devnull@example.com>
To: admin@foobar.net

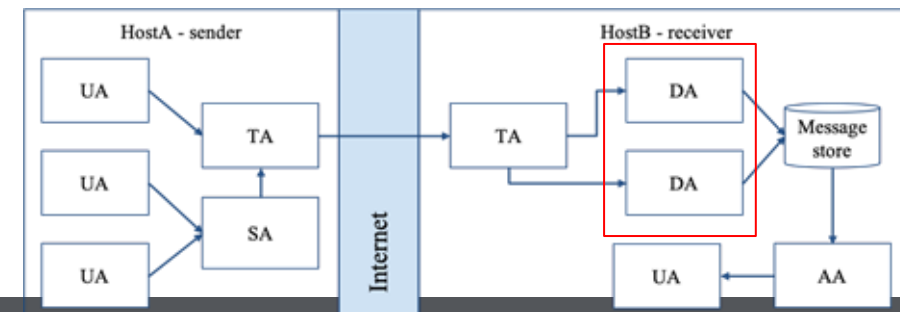
hehe... I spammed you!
.
250 2.0.0 Ok: queued as 81BD4FB4
quit
221 2.0.0 Bye
Connection closed by foreign host.
```

```
From: haha <devnull@example.com>
To: admin@foobar.net
Message-Id: <20230330070002.81BD4FB4@smtp.example.com>
Date: Thu, 30 Mar 2023 14:59:53 +0800 (CST)
```

hehe... I spammed you!

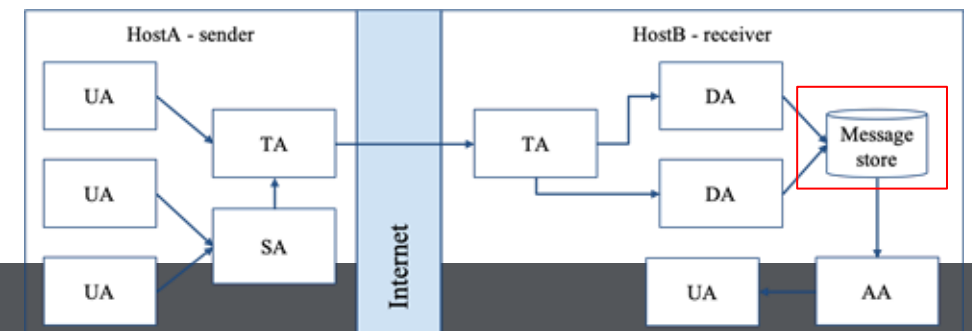
# Mail System – The Delivery Agent

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




# Mail Storage

- The place on the local machine where email is stored
  - Usually the directory: `/var/mail` or `/var/spool/mail`
    - Permission "775" and root:mail as the owner and group owner
      - `drwxrwxr-x 2 root mail 512 Dec 16 15:51 mail/`
    - Users' mails are stored in files named with each user's login name (Permission "600")
      - Eg. `/var/mail/tsaimh`
  - 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
      - E.g. sender!path!to!destination!recipient
      - This mail is sent from "sender" host to the user "recipient" at "destination" host
  - Location independent address (relies on DNS)
    - Simply identify the final destination
    - Format: user@host.domain
      - E.g. [ta@nasa.cs.nycu.edu.tw](mailto:ta@nasa.cs.nycu.edu.tw)

# Mail Addressing – Domain (2)

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

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

```
;; ANSWER SECTION:
```

```
cs.nycu.edu.tw.      3600  IN    MX    5 csmx2.cs.nycu.edu.tw.  
cs.nycu.edu.tw.      3600  IN    MX    10 csmx3.cs.nycu.edu.tw.  
cs.nycu.edu.tw.      3600  IN    MX    5 csmx1.cs.nycu.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"?
  - Centralize all the mail tasks to group of servers
    - Security enforcement, firewall control, ...
  - More robust
    - load balancing, fail over, ...

# Mail Addressing – Alias

- Alias
  - Map a username to something else
    - Mailing list
    - Be careful of **mail looping**
- Several mechanisms to define aliases:
  - Traditional method: in files
  - Traditional method + NIS
  - LDAP (Light-weight Directory Access Protocol)

# Mail Alias – Mechanisms (1)

- Places for defining alias
  - [sender] In configuration file of a MUA (e.g., `~/.mutt/muttrc`)
    - MUA expands the alias before injecting the message into the mail system
  - [receiver] In the system-wide `/etc/mail/aliases` file
    - Read by MDA
    - Deliver to the new destination when receiving the mail
  - [receiver] In user's forwarding file, `~/.forward`
    - Read by MDA after system-wide alias file
    - [forward\(5\)](#)

# Mail Alias – Mechanisms (2)

- The format of an entry in aliases file
  1. Local-name: recipient1,recipient2,...

- E.g.

- admin: tsaimh,lctseng,lwhsu,wangth,jnlin
- tsaimh: tsaimh@nycu.edu.tw
- root: ta

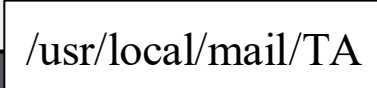
2. Local-name: **:include:**filename

- E.g.

- ta: :include:/usr/local/mail/TA



```
Flyotlin
analyst
efficacy38
enchen
bogay
ccy
```



```
/usr/local/mail/TA
```

# Mail Alias – Mechanisms (3)

- The format of an entry in aliases file

## 3. Local-name: **absolute**-path-file

- Appended to 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"

- Inject as STDIN

- Ex:

- autoftp: `"|/usr/local/bin/ftpserver"`
- nahw3: `"|/home/nahw3/receive.py"`

# Mail Alias – Mechanisms (4)

- The hashed aliases DB
  - `/etc/mail/aliases`: plain text aliases information
  - `/etc/mail/aliases.db`: hashed version for efficiency
  - `"newaliases"` command
    - Rebuild the hashed version when changing the aliases file
    - Files read from `":include:"` is **outside the aliases file**

# Mail Alias – Mechanisms (5)

- User maintainable forwarding file
  - `~/.forward`
  - Format: comma-separated
  - E.g.
    - `tsaimh@gmail.com`
    - `\tsaimh, tsaimh@gmail.com, tsaimh@nycu.edu.tw`
  - backslash + username
    - Bypassing further redirection (deliver to mailbox directly)
  - Must be owned by user and with **permission of 600**
    - The path to `.forward` file should be writable only to user

# Mail Alias – Mechanisms (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 Headers (1)

- Defined by [RFC2822](#)
  - Mail reader will hide some uninteresting header information

```
$ dig +noall +answer txt nycu.edu.tw
nycu.edu.tw. 604800 IN TXT      "v=spf1 ip4:140.113.2.64/26
ip4:211.76.241.6 ip4:140.113.98.162 ip4:140.113.98.175
ip4:140.113.9.141 ip4:120.126.43.22 ip4:140.113.7.200
ip4:140.113.40.99 include:_spf.google.com ~all"
```

【陽明交通大學 校園公告系統】 2026/3/21 新增訊息公告 收件匣 x

陽明交通大學校園公告系統 <noreply@nycu.edu.tw>

2026年3月21日 上午8:54 (23 小時前) ☆ 😊 ↶ ⋮

今日新增之公告 近 3 天到期之公告 當日活動公告

## 今日新增之公告

到期日	主旨
2026/04/03	外交部「臺灣醫衛貢獻網」及「新南向政策資訊平臺」網站將停止營運一案。
2026/04/03	【校外獎助學金公告】中華佛教善緣慈善會 清寒學生進步獎學金，申請詳見校外獎學金申請系統。
2026/04/03	【校外獎助學金公告】臺灣松樑教育公益促進協會助學金，申請詳見校外獎學金申請系統。
2026/04/03	中國鋼鐵股份有限公司徵「化學/化工工程師」。
2026/04/03	國立嘉義大學辦理「114年度高等教育深耕計畫教學創新暨社會責任成果展-原來抵嘉」資訊。
2026/04/03	馬偕醫護管理專科學校辦理「114年度高等教育深耕計畫成果發表會」活動資訊。
2026/04/03	中原大學設計學院徵求推薦院長候選人啟事，自即日起至115年4月2日收件截止一案。
2026/04/03	國立屏東大學「2026鏢旗魚海洋文化路徑與海岸聚落發展國際研討會」徵稿資訊。

- ↶ 回覆
- ↷ 轉寄
- 🗑️ 刪除
- 📧 標示為未讀取
- 🔒 封鎖「陽明交通大學校園公告系統」
- 🗑️ 回報為垃圾郵件
- 🐟 回報為網路釣魚郵件
- ≡ 篩選這類郵件
- 🗣️ 翻譯
- 🖨️ 列印
- ⬇️ 下載郵件
- ⏪ 顯示原始郵件

## 原始郵件

郵件 ID	<baddf64df245fe9a81009a8ade255cd@localhost.localdomain>
建立時間	2026年3月21日 上午8:15 (傳輸時間：2312 秒)
寄件者	陽明交通大學校園公告系統 <noreply@nycu.edu.tw>
收件者	蔡孟勳 <tsaimh@nycu.edu.tw>
主旨	【陽明交通大學 校園公告系統】 2026/3/21 新增訊息公告
SPF	PASS · IP 140.113.2.75 <a href="#">瞭解詳情</a>

下載原始郵件

複製到剪貼簿

```
Delivered-To: tsaimh@nycu.edu.tw
Received: by 2002:a05:7208:234c:b0:fa:4328:ed0c with SMTP id f12csp2034504rbe;
  Fri, 20 Mar 2026 17:54:01 -0700 (PDT)
X-Received: by 2002:a53:bbd0:0:b0:64d:6cf8:f9d with SMTP id 956f58d0204a3-64eaa7fdaa0mr3934476d50.58.1774054441020;
  Fri, 20 Mar 2026 17:54:01 -0700 (PDT)
ARC-Seal: i=1; a=rsa-sha256; t=1774054441; cv=none;
  d=google.com; s=arc-20240605;
  b=CQ5tr3uapBbP80y2ombwuwqbiwNFzaiC781P+QppRqBnwNidon8IeI0aX0AHRHoLIH
  jPLKKNof1WPlyDhgK0hWThu6rUg5730xIjYgTKvnrEq66AR/NBCNmL1BqMx9B+1/K
```

# Mail Headers (2)

Delivered-To: tsaimh@nycu.edu.tw

**Received:** by 2002:a05:7208:234c:b0:fa:4328:ed0c with SMTP id f12csp2034504rbe;

Fri, 20 Mar 2026 17:54:01 -0700 (PDT)

Return-Path: <noreply@nycu.edu.tw>

**Received:** from massmailing.nycu.edu.tw (massmailing.nctu.edu.tw. [140.113.2.75])

by mx.google.com with ESMTPS id 956f58d0204a3-64eabda477fsi3236715d50.91.2026.03.20.17.53.59

for <tsaimh@nycu.edu.tw>

(version=TLS1\_3 cipher=TLS\_AES\_256\_GCM\_SHA384 bits=256/256);

Fri, 20 Mar 2026 17:54:01 -0700 (PDT)

**Received-SPF:** pass (google.com: domain nycu.edu.tw configured 140.113.2.75 as internal address)

Authentication-Results: mx.google.com;

gateway.spf=pass (google.com: domain nycu.edu.tw configured 140.113.2.75 as internal address)

smtp.mailfrom=noreply@nycu.edu.tw smtp.remote-ip=140.113.2.75 policy.d=nycu.edu.tw

Date: Sat, 21 Mar 2026 08:15:29 +0800

To: "蔡孟勳" <tsaimh@nycu.edu.tw>

From: "陽明交通大學校園公告系統" <noreply@nycu.edu.tw>

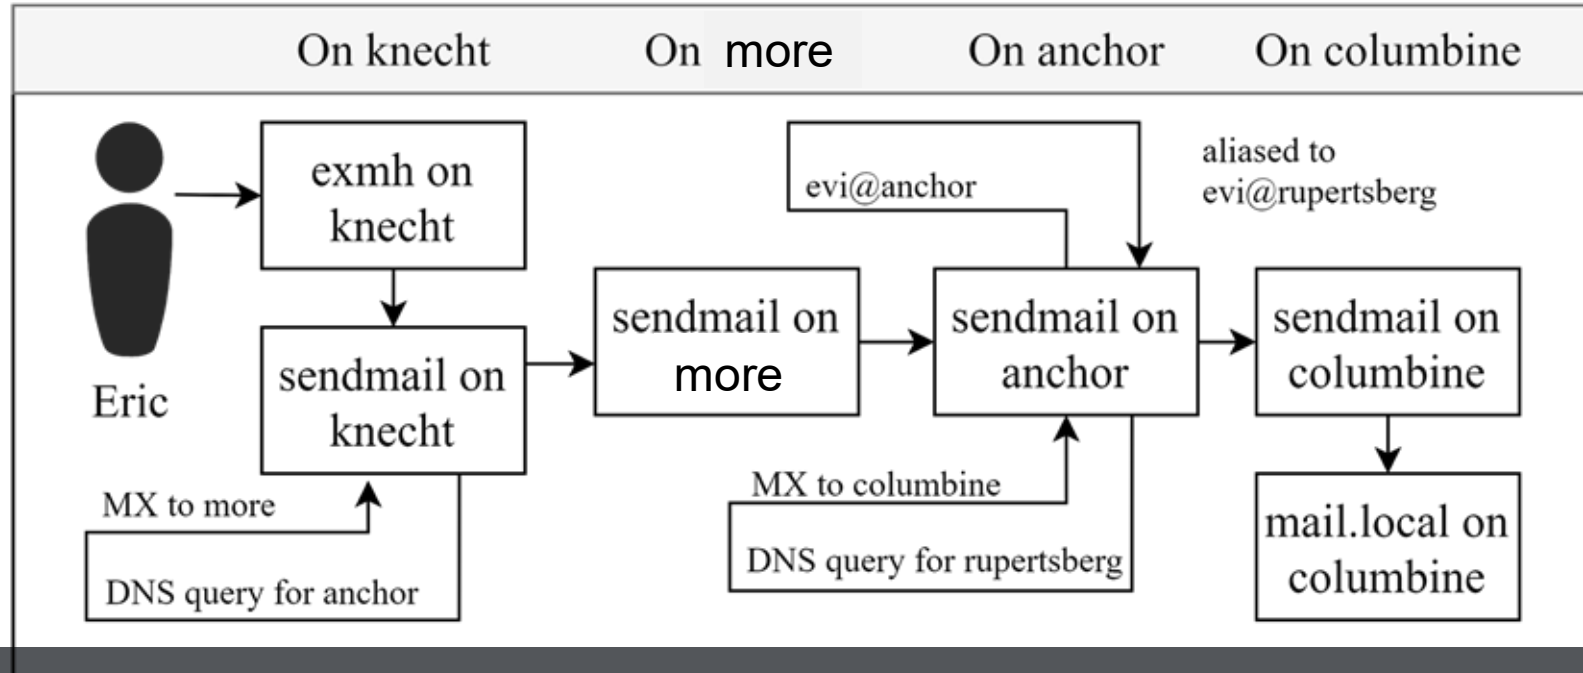
Subject: 【陽明交通大學 校園公告系統】2026/3/21 新增訊息公告

...

# Mail Transport Example (1)

- 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
    - more.cs.colorado.edu

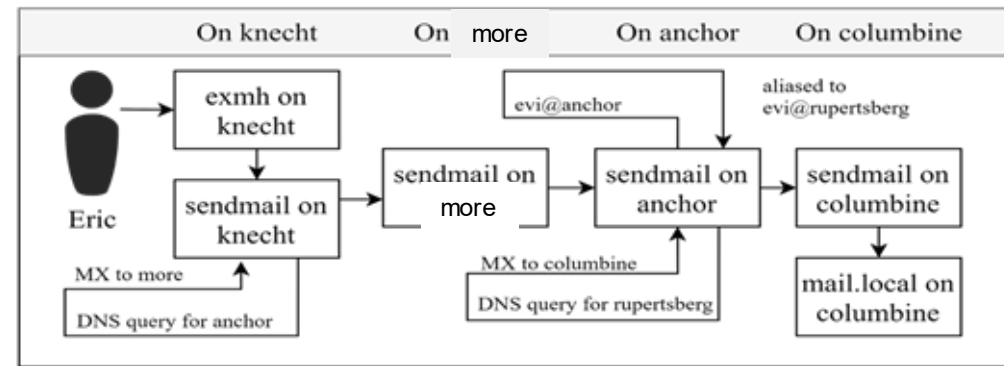
A message from Eric



# Mail Transport Example (2)

- Headers in the example
  - **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"
  - 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

A message from Eric



# Mail Transport Example (3)

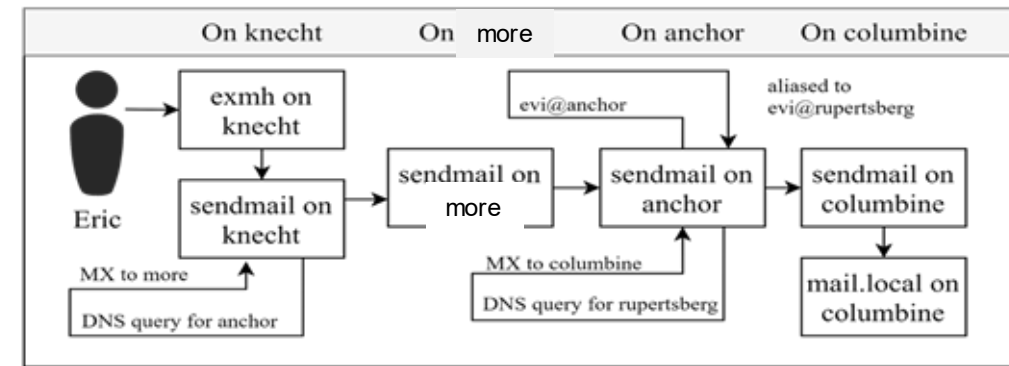
- Received Headers

- **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)

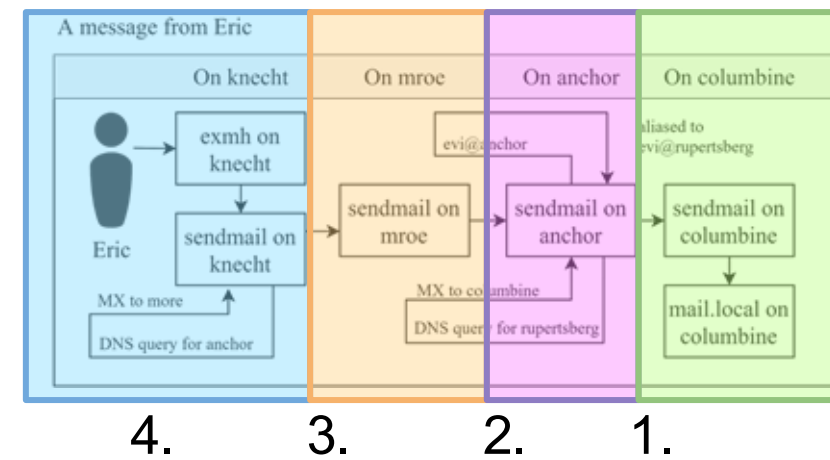
- **Every machine** that is ever processed this mail **will add a "Received"** record in **top** of headers

- 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 Transport Example (4)



1. 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)
2. 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)
3. 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 from <evi@anchor.cs.colorado.edu>; Fri, 1 Oct 1999 07:04:23 -700 (MST)
4. 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 Transport Example (5)

- From [eric@knecht.sendmail.org](mailto:eric@knecht.sendmail.org)
- To: Evi Nemeth <[evi@anchor.cs.colorado.edu](mailto:evi@anchor.cs.colorado.edu)>
- Subject: Re: hi
- Date: Fri, 1 Oct 1999 06:04:02 -800

# 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 see 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"