

Email Introduction

AfNOG 10
Kigali, Rwanda

Dorcas Muthoni


Scope

- How Email Appears to Work
- How Email Really Works
- Mail User Agent (MUA)
- Message Format
- Mail Delivery Agent (MDA)/ Mail Transfer Agent (MTA)
- Network Cloud
- Email Queue
- MTA to MTA Transfer
- Firewalls, Spam and Virus Filters
- Delivery
- RFCs

How Email Appears to Work



How Email Really Works




Mail User Agent (MUA)

- Application the originating sender uses to compose and read email

Pine, MH, Elm, mutt, mail, Eudora, Marcel, Mailstrom,

Thunderbird, Pegasus, Express, Netscape, Outlook, ...

- You can have multiple MUAs on one system - end user choice



Message Format

- **Envelope**

 - Routing information for the "postman"

- **Message Header**

 - Sender

 - Recipients (simple, lists, copies, blind copies)

 - Other fields of control (date, subject)

- **Message Body**

 - Free text

 - Structured document (i.e.: MIME)

Message Format

From: Philip Hazel <ph10@cus.cam.ac.uk>

To: Julius Caesar <julius@ancient-rome.net>

Cc: Mark Anthony <MarkA@cleo.co.uk>

Subject: How Internet mail works

Julius,

I'm going to be running a course on ...

- Format was originally defined by RFC 822 in 1982
- Now superseded by RFC 2822
- Message consists of
 - Header lines
 - A blank line
 - Body lines

Message Format

- Embedded MUA uses inter-process call to send to MTA
- Freestanding MUA uses SMTP to send mail
- Headers added by the MUA before sending


From: Philip Hazel <ph10@cus.cam.ac.uk>
To: Julius Caesar <julius@ancient-rome.net>
cc: Mark Anthony <MarkA@cleo.co.uk>
Subject: How Internet mail works

Date: Fri, 10 May 2002 11:29:24 +0100 (BST)
Message-ID: <Pine.SOL.3.96.990117111343.
19032A-100000@taurus.cus.cam.ac.uk>
MIME-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII

Julius,

I'm going to be running a course on ...

Mail Delivery Agent (MDA)/ Mail Transfer Agent (MTA)



- MDA/MTA **accepts the email, then routes it to local mailboxes or forwards** it if it isn't locally addressed
- An email **can encounter a network cloud within a large company** or ISP, or the largest network cloud in existence: the Internet.

Mail Delivery Agent (MDA)/ Mail Transfer Agent (MTA)

- Headers added by MTAs

```
Received: from taurus.cus.cam.ac.uk  
([192.168.34.54] ident=exim)  
by mauve.csi.cam.ac.uk with esmtp  
(Exim 4.00) id 101qxX-00011X-00;  
Fri, 10 May 2002 11:50:39 +0100
```

```
Received: from ph10 (helo=localhost)  
by taurus.cus.cam.ac.uk with local-smtp  
(Exim 4.10) id 101qin-0005PB-00;  
Fri, 10 May 2002 11:50:25 +0100
```

```
From: Philip Hazel <ph10@cus.cam.ac.uk>  
To: Julius Caesar <julius@ancient-rome.net>  
cc: Mark Anthony <MarkA@cleo.co.uk>
```

...

A message in transit

- A message is transmitted with an *envelope*:

MAIL FROM:<ph10@cus.cam.ac.uk>

RCPT TO:<julius@ancient-rome.net>

- The envelope is separate from the RFC 2822 message
- Envelope** (RFC 2821) fields need **not be the same** as the **header** (RFC 2822) fields
- MTAs are (mainly) concerned with envelopes

Just like the Post Office...

- Error (“bounce”) messages have null senders

MAIL FROM:<>

An SMTP session

```
telnet relay.ancient-rome.net 25
220 relay.ancient-rome.net ESMTP Exim ...
EHLO taurus.cus.cam.ac.uk
250-relay.ancient-rome.net ...
250-SIZE 10485760
250-PIPELINING
250 HELP
MAIL FROM:<ph10@cus.cam.ac.uk>
250 OK
RCPT TO:<julius@ancient-rome.net>
250 Accepted
DATA
354 Enter message, ending with "."
Received: from ...

(continued >>>>)
```


```
From: ...
To: ...
etc...
250 OK id=10sPdr-00034H-00
quit
221 relay.ancient-rome.net closing
conn...
```

SMTP return codes

```
2xx OK
3xx send more data
4xx temporary failure
5xx permanent failure
```

Network Cloud


- large company network or ISP, or the largest network cloud in existence: the Internet.
- may encompass a **multitude of mail servers, DNS servers, routers, lions, tigers, bears (wolves!)** and other devices and services
- devices may be **protected by firewalls, spam filters and malware detection software** that may **bounce or even delete an email**



Email Queue

- The email enters an email queue with other outgoing email messages.
- If there is a high volume of mail in the queue—**either because there are many messages or the messages are unusually large, or both—**

the **message will be delayed in the queue until the MTA processes the messages ahead of it.**



MTA to MTA Transfer

- Email **clears the queue, enters the Internet network cloud, where it is routed along a host-to-host** chain of servers
- The sending MTA **handles all aspects of mail delivery until the message has been either accepted or rejected** by the receiving MTA
- Each MTA needs to **"stop and ask directions" from the DNS in order to identify the next MTA in the delivery chain**
- Exact route **depends partly on server availability and mostly on which MTA can be found to accept email** for the domain specified in the address
- **ABUSE:** *Some spammers specify any part of the path, deliberately routing their message through a series of relay servers in an attempt to obscure the true origin of the message.*

DNS resolution and transfer process

- **To find the recipient's IP address and mailbox**, the MTA must **drill down through the DNS system**, which consists of a set of servers distributed across the Internet beginning with the root nameservers

root servers refer requests for a given domain to the root nameservers that handle requests for that tld

- *MTA can bypass this step because it has already knows which domain nameservers handle requests for these .tlds e.g. telecom.ma*

asks the appropriate DNS server which Mail Exchange (MX) servers have knowledge of the subdomain or local host in the email address

DNS server responds with an MX record: a prioritized list of MX servers for this domain


To the DNS server, the server that accepts messages is an MX server. When is transferring messages, it is called an MTA.

MTA contacts the MX servers on the MX record in order of priority until it finds the designated host for that address domain


sending MTA asks if the host accepts messages for the recipient's username at that domain (i.e., username@domain.tld) and transfers the message

Firewalls, Spam and Virus Filters

- An email encountering a firewall may be **tested by spam and virus filters** before it is allowed to pass inside the firewall
- filters test to see **if the message qualifies as spam or malware**
- If the message contains **malware, the file is usually quarantined and the sender is notified**
- If the message is identified as spam, **it will probably be deleted without notifying the sender.**



Delivery



- If the message makes it past the filters:

The MTA calls a local MDA to deliver the mail to the correct mailbox, where it will sit until it is retrieved by the recipient's MUA

RFCs

- Documents that define email standards are called "Request For Comments (RFCs)", and are **available on the Internet through the Internet Engineering Task Force (IETF) website**
www.ietf.org
- There are **many RFCs and they form a somewhat complex, interlocking set of standards**, great information for anyone interested in gaining a deeper understanding of email.
- Most pertinent RFCs:
 - RFC 822, 2822: Standard for the Format of ARPA Internet Text Messages
 - RFC 2821: Simple Mail Transfer Protocol
 - RFC 1122, 1123, 1651, 1653, 1830, MIME RFCs...

Troubleshooting Email Issues

- **transient failures**

If a transient error occurs, the MTA will hang onto the message, periodically retrying the delivery until it either succeeds or fails, or until the MTA decides that the transient issue is really a permanent condition.

- **permanent failures**

If the MTA cannot deliver the message (it has received a fatal error message or failed to complete the transfer after repeated attempts), it bounces the message back to the sender. If the sender is a mailing list, the bounce may be handled by automated bounce-handling software.

Q&A