

# Advanced Unix System Administration

Lecture 23  
April 25, 2007

Steven Luo  
<sluo+decal@OCF.Berkeley.EDU>

# Email

- A brief history of Internet email
  - Mail between users on multi-user machines had existed since the 1960s
  - With the advent of ARPANET (1970s), mail began to be transferred over networks
    - Initially, mail transfer was performed over FTP
  - Sites connected intermittently (usually via dialup serial link) transferred mail using UUCP
    - UUCP allows for transfer of files over serial links
    - ARPANET-UUCP gateways allowed much faster transfers

# Email

- History con't
  - Without direct connections between all sites, and no way of determining how best to get a message there, senders specified the routing of their messages
    - The infamous bang-paths:  
...foovax!barbox!quuxnet!me, where presumably everyone could find a path to foovax
    - People would often specify paths from multiple hosts: ...{ucbvax,menlo70}!barbox!me
    - Classic example:  
...{decvax,philabs}!mcvax!moskvax!kremvax!cher  
nenko

# Email

- The modern era: SMTP
  - The protocol is defined in RFC 2821, the message format in RFC 2822
  - Messages are transferred in a simple text format: headers, newline, body
    - Headers: From, To, Subject, Date, ...
    - Content of the body may not be text in modern usage – see the MIME RFCs
    - Note that the information in the headers is NOT used in routing the message

# Email

- SMTP con't
  - Routing an email message
    - The destination of an email is determined by its envelope recipient, specified to the server on sending
    - The sending mail server looks for an MX record for the destination site to determine who to talk to
      - MX records are tried in order of the priority given
      - If no MX record is present, an MX record pointing to the name itself with priority 0 is assumed
    - If the message doesn't get there, the envelope sender is returned a bounce

# Email

- SMTP con't
  - A typical SMTP conversation
    - HELO hostname or EHL0 hostname
    - MAIL FROM: <envelope-sender@site>
    - RCPT TO: <envelope-recipient@othersite>
    - DATA
    - QUIT
  - The server replies with the traditional 1xx-5xx codes

# Email

- Problems with SMTP
  - No identity verification – anyone can send email claiming to be anyone else
    - In fact, most mailing list software depends on this behavior
    - One of the enabling factors behind spam, blowback bounces, etc.
  - Legacy compatibility causes problems
    - Long timeouts, retries, etc. for unreliable networks lead to messages spending in mail queues before bouncing
    - Bang paths still (theoretically) supported!

# Email

- Mail Transfer Agents

- sendmail

- The original Internet mailer, an extension by Eric Allman of the original UUCP-era mailers
    - Complex configuration, but can do (literally) anything
    - Inefficient, slow, terrible security record
    - Includes a nice mail filtering architecture in newer versions (milter)
    - Not recommended for new installs unless you have special needs



# Email

- MTAs con't
  - qmail
    - Dan Bernstein's attempt at a secure mailer
    - Multi-binary, privilege-separated architecture provides better security and faster mail delivery
    - Introduced maildir mail storage (one file per email); often faster than mbox (one file per mail spool)
    - Unmaintained, unmaintainable by others due to license issues, has some quirky behaviors
    - Not recommended for new installs

# Email

- MTAs con't
  - Exim
    - Developed at Cambridge as a fast Internet mailer
    - Single-binary architecture sometimes criticized, but has reasonable security record
    - Extremely fast in environments where most mail is deliverable immediately, but has poor queue management
    - Can embed Perl to process email

# Email

- MTAs con't
  - Postfix
    - Wietse Venema's attempt at a secure mailer
    - Multi-binary, privilege-separated architecture
    - Very fast, simple configuration
    - Some processing isn't possible inside Postfix, but external hooks are possible
      - Such filtering setups can be a bit complex
    - Sendmail milters are supported to some degree in newer versions