ECE Mail System Overview

Pablo J. Rebollo

ECE Network Operations Center

Agenda

- Overview of ECE mail system
- How mail system works
- SPAM!!!
- ECE mail system statistics and examples
- Problems
- References

Mail system

- Previous server
  - Sun UltraEnterprise 450
    - 4 X UltraSparc 300 MHz
    - 2 Gigabytes of RAM
    - 10 x 9 Gigabytes hard drives (SCSI)
    - Solaris
  - Postfix (SMTP)
  - Inboxes in MBOX format
  - UW IMAP, and QPopper (POP3)
  - Text file for user information (/etc/passwd)

Mail System

- Current server
  - Dell PowerEdge 1750
    - 2 X Intel Xeon 3.2 GHz with HT
    - 4 gigabytes of RAM
    - 2 X 36 GB (SCSI), RAID 1 for OS
    - 14 x 73 GB (SCSI), RAID 5 for users, web pages, etc
    - Linux
  - Postfix (SMTP, SMTPS, SASL, TLS)
  - Cyrus (IMAP, POP3, TLS, maildir inboxes)
  - LDAP for user information

Mail System (cont.)

- Current system
  - Over 1,400 inboxes
  - Over 40,000 messages received per week
  - Over 10,000 messages received are SPAM
  - Over 10,000 messages sent per week
- Additional services
  - Mail gateway (Spamassassin, ClamAV)
  - Greylisting (OpenBSD spamd)

Mail System (cont.)

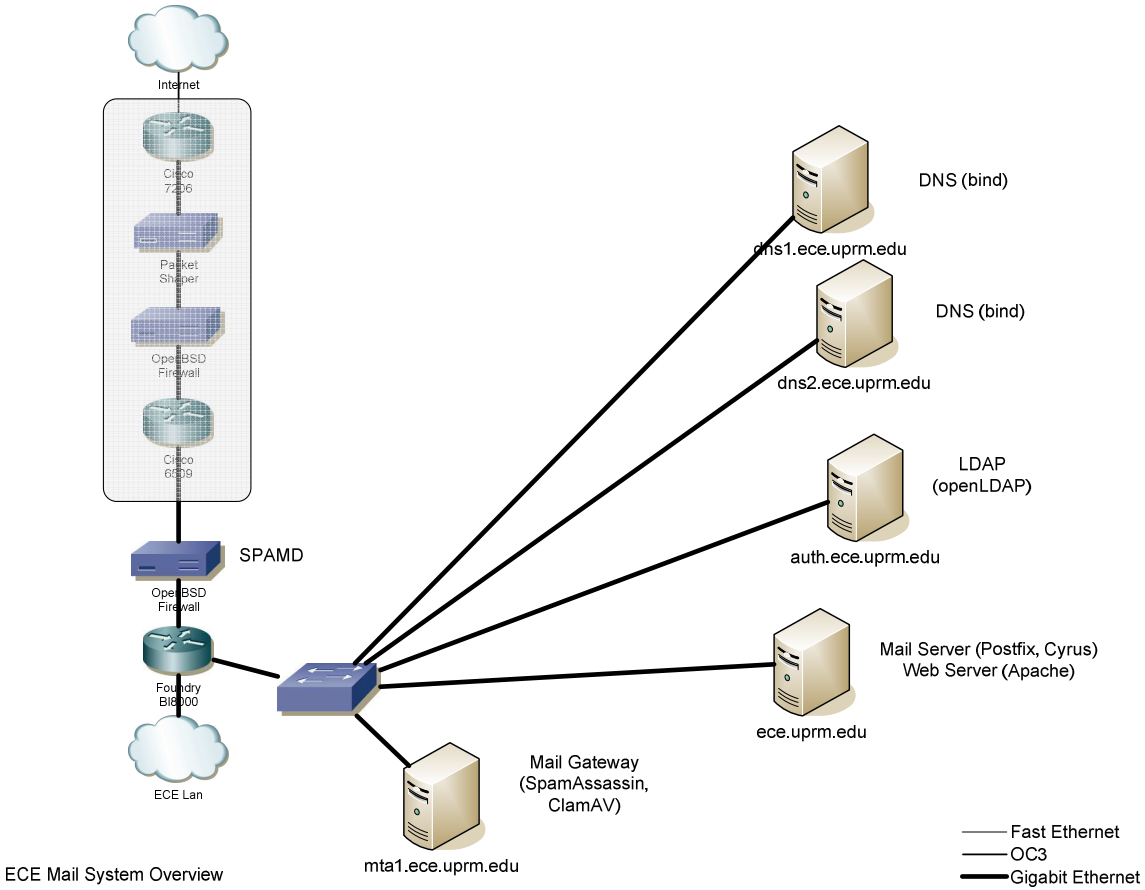

ECE Mail System Overview

How mail system works

- User sends an email with a client
- The client sends the email to the designated SMTP server.
- The SMTP server look for the MX record for the recipient domain.
- The SMTP server sends the email to the MX.
- The recipient domain mail server receives the message and store it into the user INBOX.
- Finally, the user reads the new message with an email client using IMAP or POP3.

How mail system works (cont.)

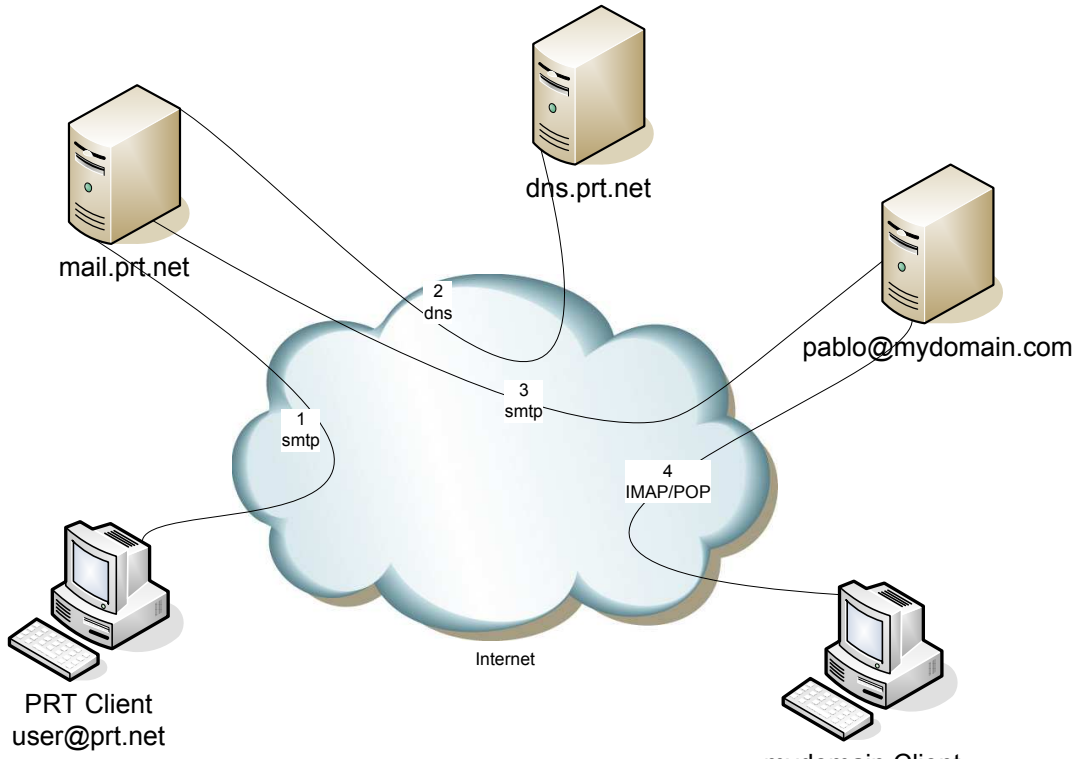

1) Client sends the messages to mail.prt.net (SMTP)
2) mail.prt.net query the MX record for mydomain.com (DNS)
3) mail.prt.net send the message to mydomain.com (SMTP)
4) Recipient reads the message (IMAP/POP)

SPAM!!!

- The biggest problem is SPAM. Users don’t want to receive SPAM. SPAM consumes bandwidth and other resources.
- To reduce the amount of spam, several techniques has been implemented.
  - Mailgateway (Spamassassin, ClamAV, FuzzyOcr)
  - OpenBSD spamd for greylisting and tarpitting.

Techniques to deal with SPAM

- Spamassassin
  - OSS used to identify SPAM by assigning scores based on several tests. If the score exceeds a threshold, then the message is tagged as SPAM (***SPAM***).
  - The software accepts custom made tests.
- ClamAV
  - OSS used to identify viruses. The system downloads new definitions every hour. Messages with viruses aren't delivered to users.
- FuzzyOCR
  - OSS who perform OCR (optical character recognition) to images contained in mail messages. This technique can hit system CPU.

Techniques to deal with SPAM

- Greylisting
  - “In name, as well as operation, greylisting is related to whitelisting and blacklisting. What happen is that each time a given mailbox receives an email from an unknown contact (ip), that mail is rejected with a "try again later"-message.This, in the short run, means that all mail gets delayed at least until the sender tries again - but this is where spam loses out! Most spam is not sent out using RFC compliant MTAs; the spamming software will not try again later.” (from: greylisting.org)
- SPF (Sender Policy Framework)
  - The idea is to advertise the authorized mail server for a specific domain. This is achieved by publishing a TXT record for a domain.
- Postfix SASL
  - This option force users to be authenticated first when sending email to external accounts (relaying) when they aren’t connected to ECE facilities.

Stats & Examples

- Mailgateway Statistics

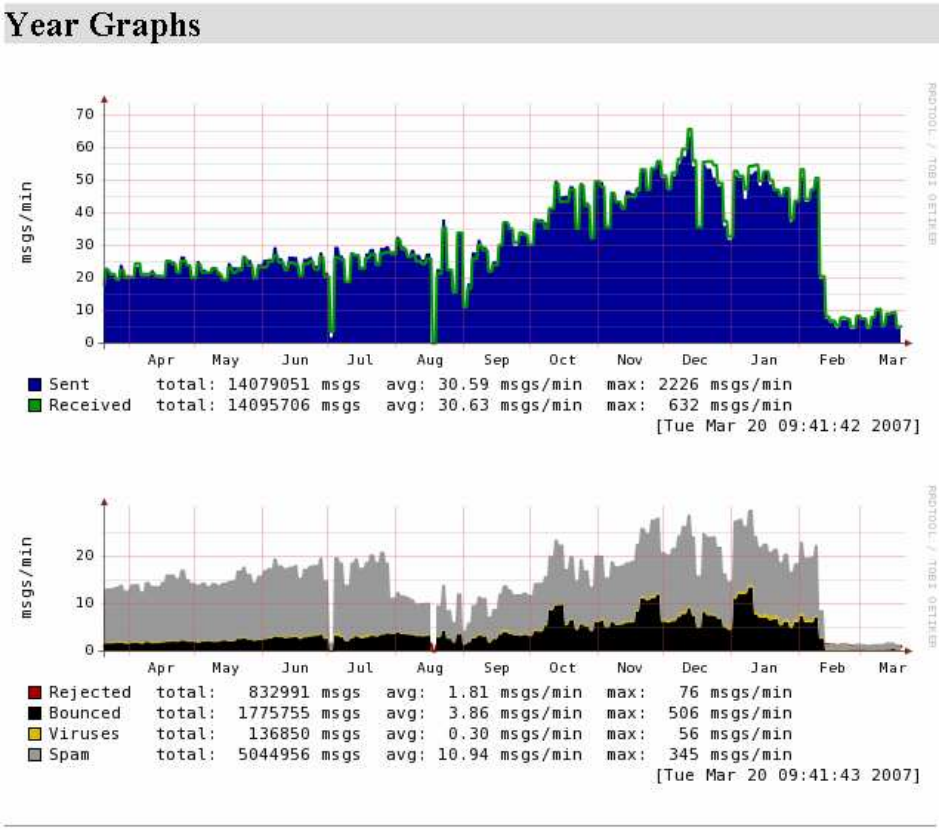

Stats & Examples (cont.)

- Spamd Statistics

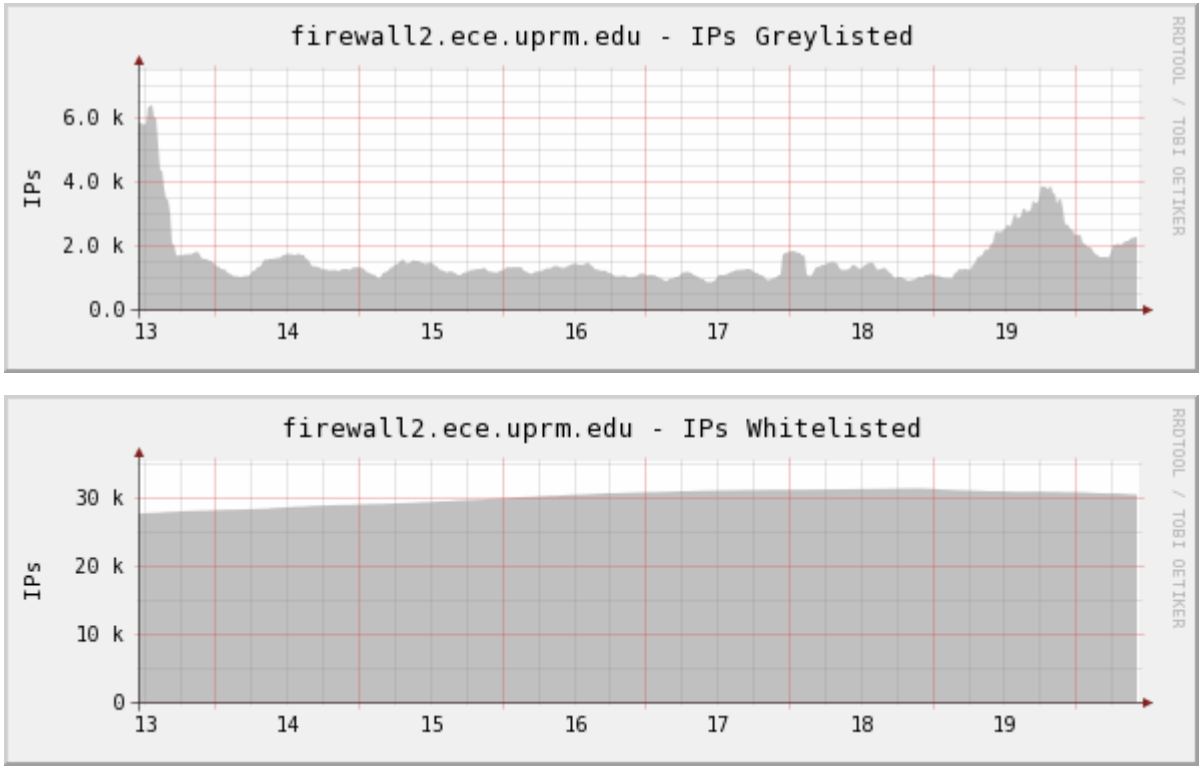

Stats & Examples (cont.)

- DNS Query

```
pablor@noc:~> dig ece.uprm.edu ANY

; <<>> DiG 9.2.4 <<>> ece.uprm.edu ANY
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 23336
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 7, AUTHORITY: 0, ADDITIONAL: 4

;; QUESTION SECTION:
;ece.uprm.edu.                  IN      ANY

;; ANSWER SECTION:
ece.uprm.edu.           3600    IN      SOA     dns1.ece.uprm.edu. pablor.ece.uprm.edu. 2007031900 3600 3600 604800 86400
ece.uprm.edu.           43200   IN      NS      dns1.ece.uprm.edu.
ece.uprm.edu.           43200   IN      NS      dns1.uprm.edu.
ece.uprm.edu.           43200   IN      NS      dns2.ece.uprm.edu.
ece.uprm.edu.           43200   IN      A       136.145.57.24
ece.uprm.edu.           43200   IN      MX      0 mta1.ece.uprm.edu.
ece.uprm.edu.           43200   IN      TXT     "v=spf1 a mx ~all"

;; ADDITIONAL SECTION:
dns1.ece.uprm.edu.      43200   IN      A       136.145.57.3
dns1.uprm.edu.          3600    IN      A       136.145.30.2
dns2.ece.uprm.edu.      43200   IN      A       136.145.58.3
mta1.ece.uprm.edu.      43200   IN      A       136.145.57.11

;; Query time: 1 msec
;; SERVER: 136.145.57.3#53(136.145.57.3)
;; WHEN: Tue Mar 20 10:59:26 2007
;; MSG SIZE  rcvd: 260
```

Stats & Examples (cont.)

Spamassassing report

Content analysis details: (14.5 points, 5.0 required)

pts	rule name	description
1.1	EXTRA_MPART_TYPE	Header has extraneous Content-type:...type= entry
2.0	DATE_IN_FUTURE_03_06	Date: is 3 to 6 hours after Received: date
0.5	HTML_40_50	BODY: Message is 40% to 50% HTML
0.0	HTML_MESSAGE	BODY: HTML included in message
4.3	BAYES_99	BODY: Bayesian spam probability is 99 to 100%
3.8	LONGWORDS	Long string of long words
3.0	DC_PNG_UNO_LARGO	Message contains a single large inline gif
-0.1	AWL	AWL: From: address is in the auto white-list

Problems

- The most common problem is with false positives. To deal with this kind of problem is important to have users feedback.
- Another problem can be delivery delays due to greylisting process. This could be solved by having a static whitelist.

References

- Postfix
  - http://www.postfix.org/
- Cyrus
  - http://cyrusimap.web.cmu.edu/
- Spamassassin
  - http://spamassassin.apache.org/
- ClamAV
  - http://www.clamav.net/
- FuzzyOCR
  - http://wiki.apache.org/spamassassin/FuzzyOcrPlugin
- Greylisting
  - http://www.greylisting.org/
- OpenBSD spamd
  - http://www.openbsd.org/cgi-bin/man.cgi?query=spamd&sektion=8
- SPF
  - http://www.openspf.org/
- OpenBSD spamd - greylisting and beyond
  - http://www.ualberta.ca/~beck/nycbug06/spamd/index.html