

CSE 265:

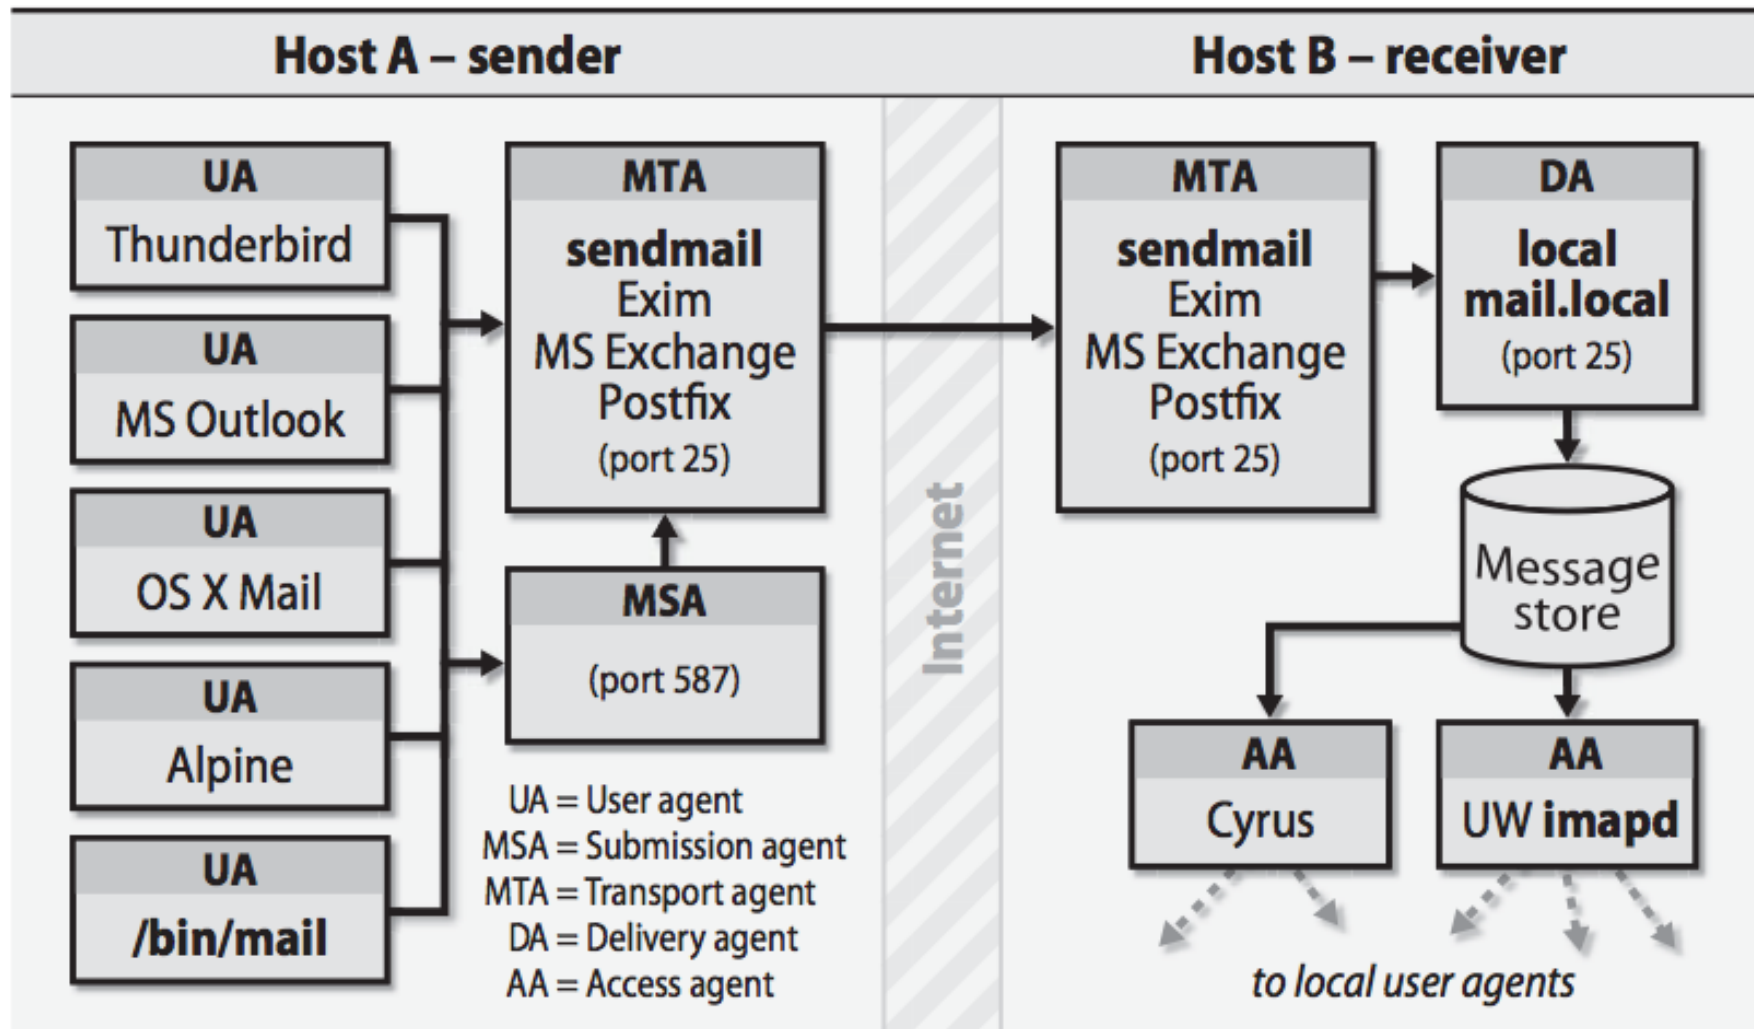
System and Network Administration

- Electronic Mail
 - Mail systems
 - Addressing, mail headers
 - Client/server philosophy, mail homes
 - Aliases, mail routing, mailing list software
 - sendmail
 - Security
 - Performance

Mail systems

- Four components
 - **Mail user agent** (MUA) to read and compose mail
 - **Mail transport agent** (MTA) route messages
 - **Delivery agent** that stores messages for later retrieval by users
 - Optional **access agent** to connect user agent to message store

The big picture



User agents

- Provide means to read and compose email



- Outlook, Thunderbird, Eudora, pine, elm, IMP, /bin/mail, emacs, web-based gmail, and more
- Often have system-wide and personal configuration files
- Multipurpose Internet Mail Extensions (MIME) encoding for different text formats and attachments

Transport agents

- Transport agents accept mail from a user agent, and deliver mail to the correct hosts
 - PMDF, **postfix**, smail, Exim, **sendmail**
- Speak the Simple Mail Transport Protocol (SMTP) or Extended SMTP (ESMTP)
- Run on port 25 (unencrypted)



Delivery agents

- Accepts mail from a transport agent, and delivers to the local recipient
- Delivery can be to
 - a person's mailbox
 - a mailing list
 - a file
 - a program
- Agents include
 - `/bin/mail` for local users
 - `/bin/sh` for programs
 - `/usr/bin/procmail` for user-configurable delivery



Access agents

- Agents include
 - imapd – IMAP server
 - insecure, port 143
 - secure, port 993
 - spop – POP server
 - insecure, port 109 (pop2), 110 (pop3)
 - secure, port 995



Mail submission agents (MSA)

- High volume sites may need a separate mail submission agent
- Preprocess messages
 - Ensure hostnames are fully qualified
 - Modify broken headers
 - Log errors
 - Re-write headers
- Usually runs on port 587 or 465 (smtps)
- sendmail can act as an MSA (as well as MTA)

Mail messages

- Three components
 - The envelope
 - Where the message is to be delivered, plus where to return if undeliverable
 - Different from header lines From: and To:
 - Supplied separately to the MSA
 - The headers
 - Collection of property-value pairs
 - Includes date and times and agents through which the message has passed
 - The body
 - Actual contents (in plain text)

Sample mail headers #1

From rjd0@lehigh.edu Wed Sep 26 16:50:49 2001
Received: from rain.CC.Lehigh.EDU (rain.CC.Lehigh.EDU [128.180.39.20])
by genie.eecs.lehigh.edu (8.9.3/8.9.3) with ESMTTP id QAA03440
for <brian@cse.lehigh.edu>; Wed, 26 Sep 2001 16:50:34 -0400 (EDT)
Received: from lehigh.edu (iceBook.CC.Lehigh.EDU [128.180.3.8])
by rain.CC.Lehigh.EDU (8.11.5/8.11.5) with ESMTTP id f8QKoIT24177
for <brian@cse.lehigh.edu>; Wed, 26 Sep 2001 16:50:24 -0400
Message-ID: <3BB23F7A.A1005AC8@lehigh.edu>
Date: Wed, 26 Sep 2001 16:50:01 -0400
From: Robin Deily <rjd0@lehigh.edu>
Organization: Lehigh University
X-Mailer: Mozilla 4.75C-CCK-MCD {C-UDP; EBM-APPLE} (Macintosh; U; PPC)
X-Accept-Language: en
MIME-Version: 1.0
To: "Brian D. Davison" <brian@cse.lehigh.edu>
Subject: Re: commercial internet outage
References: <Pine.SOL.3.91.1010926112807.18638A@pan>
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
Status: RO
X-Status:
X-Keywords:
X-UID: 2

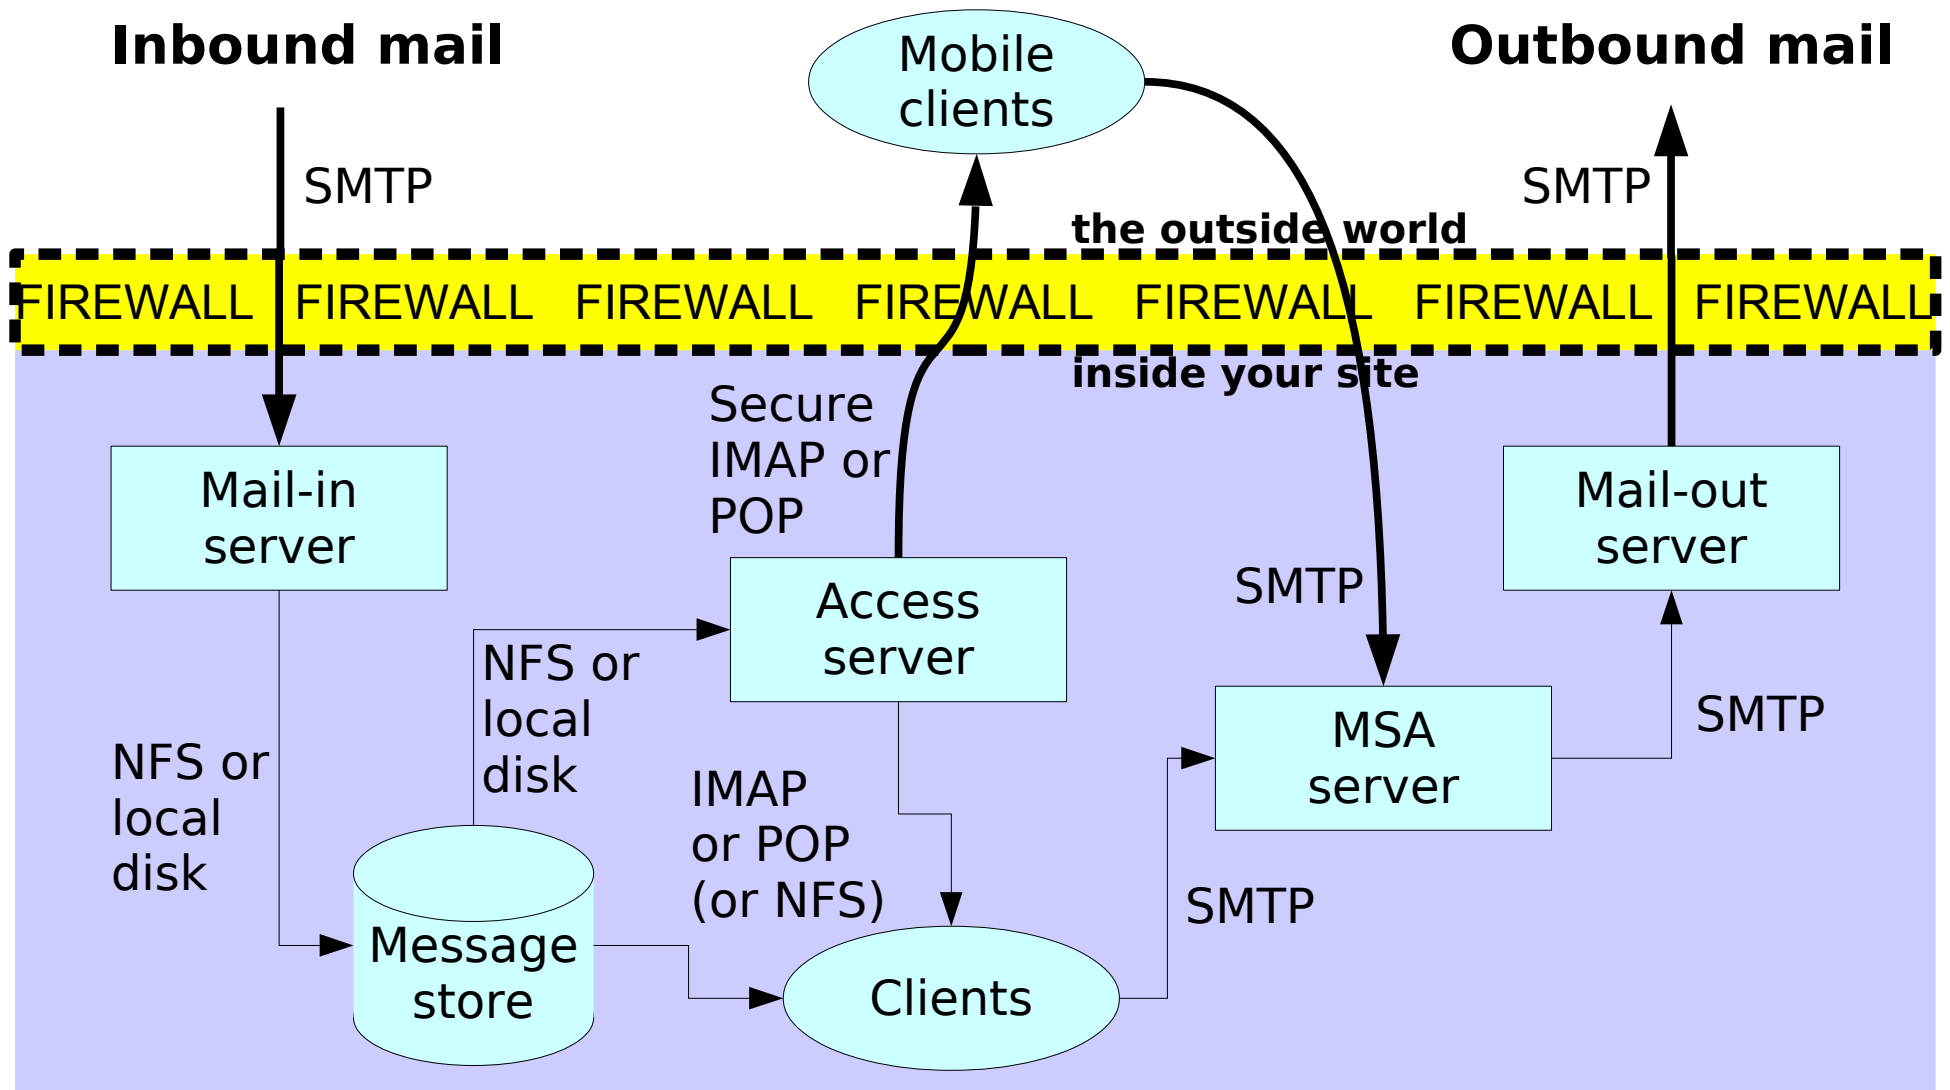
Sample mail headers #2

From BBUOVA@yahoo.com Fri Mar 19 12:37:49 2004
Received: from rain.CC.Lehigh.EDU (rain.CC.Lehigh.EDU [128.180.39.20])
by genie.eecs.lehigh.edu (8.12.10/8.12.10) with ESMTTP id i2JHbmN9014501
for <brian@cse.lehigh.edu>; Fri, 19 Mar 2004 12:37:48 -0500 (EST)
Received: from alias.acm.org (alias.acm.org [199.222.69.90])
by rain.CC.Lehigh.EDU (8.12.11/8.12.11) with ESMTTP id i2JHZ2Sa006893
for <davison@lehigh.edu>; Fri, 19 Mar 2004 12:35:03 -0500
Received: from 12-219-103-195.client.mchsi.com ([12.219.103.195])
by alias.acm.org (ACM Email Forwarding Service) with SMTP id COB73880;
Fri, 19 Mar 2004 12:35:00 -0500
X-Message-Info: EUKNoBG22bAWz/vLgLAarLmRbForUh0F
Received: from deface-l13.besiege.aol.com ([239.93.237.144]) by tp9-h40.hotmail
.com with Microsoft SMTPSVC(5.0.2195.6824);
Sat, 20 Mar 2004 12:23:54 +0300
From: Olin Pack <BBUOVA@yahoo.com>
To: davidlow@acm.org
Subject: wknd-wonder is here! homestead
Date: Sat, 20 Mar 2004 08:19:54 -0100 EST
Message-ID: <75395305408904.00820.60856274@yucatan-t14.aol.com>
Mime-Version: 1.0
Content-Type: multipart/alternative;
boundary="--7357593428207540603"
Content-Length: 873

Mail architecture

- Typical architecture
 - Servers for incoming and outgoing mail
 - A mail home for each user in an organization
 - IMAP or POP for access by users (PCs, Macs, remote clients)
- A mail server needs
 - to accept outgoing mail from user agents and inject into mail system
 - to receive incoming mail from outside world
 - to deliver mail to end-user's mailboxes
 - to allow users to access mail via IMAP (or perhaps POP)

Sample architecture



Aliases and mail forwarding

- Mail can be re-routed by admins or users
 - when sending user's agent config file has a replacement
 - when there is an entry in `/etc/aliases`
 - when the receiving user has a `~/.forward` file
- Sample `/etc/aliases` entries:

```
webmaster:    steinberg,hodgson
support:      :include:/usr/local/mail/lists/support.ml
help:         support
```
- **newaliases** rebuilds alias database
- Sample `.forward` files:
 - `"| IFS=' ' && exec /usr/bin/procmail -t || exit 75 # brian"`
 - `user@newaddress.com`

Mailing lists

- sendmail treats entries in /etc/aliases that :include: files as mailing lists
- If an alias for owner-mylist exists, sendmail uses the value of that alias as the envelope sender
 - This makes list bounces go to the list owner, rather than to the poster of the message
 - If the bounced message also bounces, then the value of the alias owner-owner gets the message (or postmaster, otherwise)
- Many packages help to maintain mailing lists
 - Majordomo, mailman, ListProc, SmartList, etc.

sendmail

- One common MTA for Linux
- sendmail does most of the work
 - understands recipients' addresses
 - chooses an appropriate delivery or transport agent
 - rewrites addresses to be understood by delivery agent
 - reformats headers as required
 - generates error messages and returns messages to senders if undeliverable
- System daemon explicitly started at boot

sendmail modes

- -b flag determines modes
 - -bd daemon mode, listen on port 25
 - -bD, but in foreground rather than background
 - -bp print mail queue (same as mailq)
 - -bt address test mode
 - -bv verify mail addresses only (don't send mail)
- -q30m attempts to process the mail queue every 30 minutes

mail queue

- Mail messages are stored in the queue directory `/var/spool/mqueue`
 - when the system is too busy to deliver them immediately
 - when a destination machine is unavailable
- `/usr/bin/mailq` to view
 - separate files for headers, body, error messages

```
          /var/spool/mqueue (24 requests)
-----Q-ID----- --Size--  -----Q-Time-----  -----Sender/Recipient-----
i2JKcuR26576      4230 Fri Mar 19 15:38 MAILER-DAEMON
      8BITMIME      (Deferred: Connection timed out with sbcglobal.com.)
                   <mchohl@sbcglobal.com>
i2K2G7R12880*    3479 Fri Mar 19 21:16 MAILER-DAEMON
                   (Deferred: Connection timed out with 168.com.)
                   <enxwesbkqen@168.com>
```

sendmail configuration

- /etc/sendmail.cf – only read at startup
- Specifies
 - choice of delivery agents
 - address rewriting rules
 - mail header formats
 - options
 - security precautions
 - spam resistance
- Raw config file is almost **unreadable**
- Use a preprocessor (m4) instead

sendmail and m4

- **m4** is a generic macro preprocessor
 - macros have form
 - name(arg1, arg2, ..., argn)
 - **dnl** is built-in macro to ignore until newline
 - used to convert sendmail.mc to sendmail.cf
 - strings use open and close quote `example'
- Typical process
 - 1) edit .mc file with changes
 - 2) rebuild config file
 - 3) install config file in right directory
 - 4) restart sendmail

sendmail m4 primitives

- OSTYPE(`linux')
 - OS-specific flags, file locations, etc.
- **define(`ALIAS_FILE',`/etc/aliases,nis:mail.aliases')**
 - Define which sources and ordering of aliases
- MAILER(smtp) and/or MAILER(procmail)
 - Specify which local mailers are enabled
- FEATURE(`use_cw_file')
 - /etc/mail/local-host-names contains all names for system
- FEATURE(`always_add_domain')
 - adds the local hostname to local addresses when needed

Virtual Users

- sendmail supports domain aliasing for incoming mail

- FEATURE(`virtusertable')

- Examples

```
info@foo.com    foo-info        # route to local user
info@bar.com    bar-info        # another local user
@baz.org        jane@elsewhere.com # all mail to jane
@zokni.org      %1@elsewhere.com # same user, dif. domain
```

- Still need

- MX records for each domain (to receive such mail)
- cw entries for each domain (to enable relay)

Sample sendmail.mc

```
divert(-1)
dnl This is the sendmail macro config file. If you make changes to this,
dnl generate a new /etc/sendmail.cf by running the following command:
dnl      m4 /etc/mail/sendmail.mc > /etc/sendmail.cf
dnl
include(`/usr/lib/sendmail-cf/m4/cf.m4')
VERSIONID(`linux setup for Red Hat Linux')dnl
OSTYPE(`linux')
define(`confDEF_USER_ID',`8:12')dnl
define(`confAUTO_REBUILD')dnl
define(`confTO_CONNECT',`1m')dnl
define(`confDONT_PROBE_INTERFACES',true)dnl
define(`ALIAS_FILE',`/etc/aliases')dnl
define(`confUSERDB_SPEC',`/etc/mail/userdb.db')dnl
define(`confPRIVACY_FLAGS',`goaway,authwarnings,restrictqrun')dnl
FEATURE(`no_default_msa',`dnl')dnl
FEATURE(`smrsh',`/usr/sbin/smrsh')dnl
FEATURE(`mailertable',`hash -o /etc/mail/mailertable.db')dnl
FEATURE(`virtusertable',`hash -o /etc/mail/virtusertable.db')dnl
FEATURE(redirect)dnl
FEATURE(always_add_domain)dnl
FEATURE(use_cw_file)dnl
FEATURE(use_ct_file)dnl
```

sample sendmail.cf portions (1)

```
Cwlocalhost
# file containing names of hosts for which we receive email
Fw/etc/mail/local-host-names

#####
#   Format of headers   #
#####

H?P?Return-Path: <$g>
HReceived: $?sfrom $s $.?$_($?s$|from $.$_)
           $.?{auth_type}(authenticated${auth_ssf} (${auth_ssf} bits)$.)
           $.by $j ($v/$Z)$?r with $r$. id $i$?{tls_version}
           (using ${tls_version} with cipher ${cipher} (${cipher_bits} bits) verifi
ed ${verify})$.?u
           for $u; $|;
           $. $b
H?D?Resent-Date: $a
H?D?Date: $a
H?F?Resent-From: $?x$x <$g>$|$g$.
H?F?From: $?x$x <$g>$|$g$.
H?x?Full-Name: $x
# HPosted-Date: $a
# H?l?Received-Date: $b
H?M?Resent-Message-Id: <$t.$i@$j>
H?M?Message-Id: <$t.$i@$j>
```


sample sendmail.cf portions (2)

```
#####  
### Ruleset 3 -- Name Canonicalization ###  
#####  
Scanonify=3  
  
# handle null input (translate to <@> special case)  
R$@          $@ <@>  
  
# strip group: syntax (not inside angle brackets!) and trailing semicolon  
R$*          $: $1 <@>          mark addresses  
R$* < $* > $* <@>          $: $1 < $2 > $3          unmark <addr>  
R@ $* <@>          $: @ $1          unmark @host:..  
R$* :: $* <@>          $: $1 :: $2          unmark node::addr  
R:include: $* <@>          $: :include: $1          unmark :include:..  
R$* [ IPv6 : $+ ] <@>          $: $1 [ IPv6 : $2 ]          unmark IPv6 addr  
R$* : $* [ $* ]          $: $1 : $2 [ $3 ] <@>          remark if leading colon  
R$* : $* <@>          $: $2          strip colon if marked  
R$* <@>          $: $1          unmark  
R$* ;          $1          strip trailing semi  
R$* < $+ ;; > $*          $@ $2 ;; <@>          catch <list:;>  
R$* < $* ; >          $1 < $2 >          bogus bracketed semi
```

sendmail.mc continued

```
define(`PROCMAIL_MAILER_PATH',`/usr/bin/procmail')dnl
FEATURE(local_procmail,`,`,`procmail -t -Y -a $h -d $u')dnl
FEATURE(`access_db',`hash -o /etc/mail/access.db')dnl
FEATURE(`blacklist_recipients')dnl
FEATURE(dnsbl,`dnsbl.njabl.org',`Message from ${client_addr} rejected -
  see http://njabl.org/lookup?\${client\_addr}')
FEATURE(`dnsbl',`relays.ordb.org',`"550 Email rejected due to sending
  server misconfiguration - see
  http://www.ordb.org/faq/#why\_rejected"')dnl
FEATURE(`dnsbl',`psbl.surriel.com',`*** SPAM Blocked --
  See http://psbl.surriel.com/')dnl
FEATURE(`dnsbl',`dnsbl.sorbs.net',`"554 Rejected " ${client_addr} "
  found in dnsbl.sorbs.net"')dnl
FEATURE(`dnsbl',`dnsbl-1.uceprotect.net',`"554 Rejected " ${client_addr}
  "is BLACKLISTED at LEVEL 1 by UCEPROTECT-NETWORK. To be removed see
  http://www.uceprotect.net"')dnl
EXPOSED_USER(`root')dnl
MAILER(smtp)dnl
MAILER(procmail)dnl
Cwlocalhost.localdomain
```

Debugging

- SMTP is a simple protocol with only 14 commands
 - Can use telnet to connect to an SMTP server and issue commands manually
- sendmail uses syslog – messages get placed into `/var/log/maillog` (on RHEL)

```
Mar 22 10:55:10 localhost sendmail[26115]: i2MFt9D26115: ruleset=check_relay,  
  arg1=mx-01.suga-n-spice.com, arg2=64.201.119.12, relay=mx-01.suga-n-spice.com  
  [64.201.119.12], reject=553 5.3.0 *** SPAM Blocked from 64.201.119.12 - See  
  http://bl.csma.biz/.  
Mar 22 10:55:10 localhost sendmail[26115]: NOQUEUE: mx-01.suga-n-spice.com  
  [64.201.119.12] did not issue MAIL/EXPN/VRFY/ETRN during connection to MTA
```

Final comments

- My server/domains were online 1995~2010
 - Well-publicized domains and email addresses
 - Posted to mailing lists, newsgroups, and in Web pages
- Few accounts; each got hundreds of SPAM/day
- Using the **dnsbl** feature with multiple sites has blocked (not filtering) ~2000 messages per day
 - some still get through (perhaps 5%)
- Find list of dnsbl sites at
 - <http://www.dnsbl.info/>
- Check potential spammer/relay IPs in multiple lists
 - <http://multirbl.valli.org/lookup/> or <http://www.mxtoolbox.com>