

Bell Laboratories Cover Sheet for Technical Memorandum

The information contained herein is for the use of employees of Bell Laboratories and is not for publication. (See GEI 13.9-3)

Title- **Electronic Mail Without Aliases**

Date- **September 30, 1980**

TM- **80-1273-7**
 80-1274-11

Other Keywords- **UNIX**
 Electronic mail

Author	Location	Extension	Charging Case-
R. Jane Elliott	MH 2C-511	2879	39199
Michael Lesk	MH 2C-572	6377	Filing Case- 39199-11

ABSTRACT

Electronic mail at Bell Labs is usually routed to a system and a user-id. There are now over 100 BTL UNIX systems connected by the UUCP network, and it is more and more of a nuisance to remember where your correspondents log in and what user-id they use. The new utility *post* looks up the real names of people in a database of BTL UNIX users, and sends mail appropriately. The most specific kind of address is of the form "ruby.jane.elliott-mh-1273" but everything except "elliott" is optional. When someone has no computer address, the mail is routed to a nearby clerk, with appropriate labels so that the mail need merely be printed and then put in the company mail. If the eventual address is not only non-electronic but outside the company (e.g. at AT&T) the mail is routed to a clerk near the sender and addressed for the U.S. Mail.

The names, room numbers, and department numbers of all employees are updated automatically from the standard company files; userids can be updated by the users themselves, simply by sending the appropriate electronic mail. A line of the form

```
UPDATE: research!mel IS michael.e.lesk-mh-1274
```

mailed to *`postoffice`!mailroom* will update the appropriate listing.

As with other data bases, there is the usual choice between distributing the data base or routing all queries to a single site. The *post* command may be used with local directories of users, or all mail may be sent to a central site which has a phone book and then the central site can route and deliver the message. In addition, it is possible to route some mail locally and then forward other mail to the central site. The individual user may also have a directory of abbreviations; thus the directory is really a hierarchy of directories running through personal directories, local site directories, and finally large central directories.

Pages Text	5	Other	0	Total	5
No. Figures	1	No. Tables	0	No. Refs.	2

Bell Laboratories

Subject: **Electronic Mail Without Aliases**
Case- 39199 -- File- 39199-11

date: **September 30, 1980**

from: **R. Jane Elliott
Michael Lesk**

TM: **80-1273-7
80-1274-11**

MEMORANDUM FOR FILE

Introduction

There are over 100 Bell Labs UNIX systems now in the UUCP network. It's hard to remember where everyone you know works so that you can send them electronic mail. The trend from somewhat meaningful system names (e.g. *mhwa* or *ihuxa*) to meaningless ones (e.g. *alice*, *eagle*, or *groucho*) makes it even more difficult to remember who works where. As a result, electronic mail is largely used within single systems, with only 13% of all mail being to another system. By contrast, 19% of regular BTL paper mail goes not just to other groups but to other locations. To make it easier to send electronic mail, the command *post* maintains data bases of userids and real names. It can be used to address electronic mail to people by their real names.

We maintain a data base which gives, for all Bell Labs employees, their real name, department number, location and room number, UNIX system name, and user-id. Non-employees can also be entered in the data base if we wish to send them mail (e.g. AT&T users of BTL computer systems). The real name, department, and room address of each employee are taken from the Bell Labs phone book and regularly updated. We keep a list of UNIX names and user-ids. We store the data in a B-tree, using Peter Weinberger's programs,¹ and can retrieve either the address of a particular individual or the real name of a computer user. A smaller data base of clerks is also maintained. This is a list, similar to the last, of people with computer addresses who can be asked to take paper mail and put it in mailboxes.

Implementation

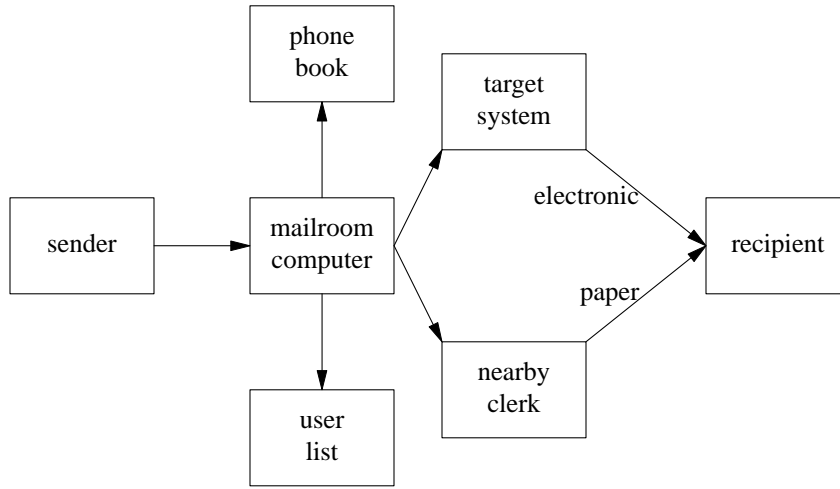
Given a letter and a destination name, the program first identifies the person in its name list. If the name is ambiguous, and the user is at the terminal, a menu of possible individuals is presented and the user is asked to choose one. When the program has identified the addressee, it dispatches the mail if possible. Several routes are available:

1. If the addressee has an electronic mail address, the mail is sent electronically.
2. If the addressee is at Bell Labs but has no login user-id, the mail is formatted for paper and sent electronically to the closest clerk to the addressee. "Closest" is defined by the organization number; i.e. same department if possible, otherwise same laboratory if possible, otherwise same division. If no clerk is registered anywhere in that organization, then the program selects a clerk at that main location.

1. P. J. Weinberger, *UNIX B-trees*, TM 80-1272-xx.

3. If the addressee is not at Bell Labs, but the sender is, the mail is formatted for paper and sent electronically to the closest clerk to the sender. This clerk then puts the letter in the external mail.
4. If we have no electronic address for the addressee, and neither the addressee nor the sender is at Bell Labs, then at present we abandon the mail. There is no reason in principle not to have a list of clerks outside BTL, at least at Western and AT&T, but at the moment we have decided not to keep such a list.

In any case, the mail is identified as to the sender: we look up the user invoking the command in our data base, and supply his real name and department in the mail. The pattern of operation is shown in the figure below. The mail moves electronically until it reaches either the final destination or the clerk. If the destination is at BTL, the clerk is near the recipient; otherwise the clerk is near the sender.



Clearly, *post* can only be used to communicate with sites in *uucp* communication with the mailroom computer. We hope that some administrative solution, perhaps similar to the one proposed by Grampp, Koenig, and Ritchie,² will provide for general updating of *uucp* lists. In any case keeping these mail lists up to date on one computer is less work than keeping them up to date everywhere.

Usage

The *post* command is used just as the regular *mail* command is used. There is no difference at all if you intend to read mail, since *post* invoked without arguments just executes ordinary *mail*. If the user has a shell environment parameter MAILP set, it provides the name of the mail program; otherwise “mail” is used. This permits a user who prefers Berkeley mail or other mail variants to have it invoked instead of the standard mail command. To send mail, you type *post addressee1 addressee2 ...* and supply the message on the standard input. The usual *-R* option is also recognized (this specifies “return receipt requested” to give positive confirmation that the mail has been received and read). The format for addressee is *firstname.middle.lastname-location-department* in which everything but the last name is optional. Upper and lower case letters are equivalent throughout, and the location and department may be in either order. Location should be specified as the ordinary Bell Labs 2-letter code; AT&T building numbers are not recognized. The first name, middle name, and department may be searched as prefixes, i.e. initials will match names and as many digits of the department number as are known may be used.

To update your listing, send mail to “postoffice`!mailroom” of the form

```
UPDATE: system!userid IS john.a.smith-ho-3456
```

where 3456 is the department number. This produces a listing in the directory as indicated, changing the previous one if one existed. The Bell Labs location and room number are obtained from the standard Bell

2. F. T. Grampp, A. R. Koenig, and D. M. Ritchie, *Directory assistance for UUCP*, MF 3644-800318.01.

Labs telephone book data and are not supplied by the user. To register a U.S. mail address only, and not a computer address, send

```
UPDATE: john.smith USADDR WECO.; 220 Broadway; New York 10007
```

Sometimes, it may be desirable to define a particular userid, but not to direct mail there. It may, for example, be a secondary userid for the user in question; or, although primary, it may be a userid for someone who rarely uses the computer. It is nevertheless useful for *post* to know who owns the userid, so that it can identify mail sent from that userid. To indicate ownership of a userid without getting mail sent to it, send

```
UPDATE: system!userid ANOTHER john.smith
```

to the same place (``postoffice`!mailroom`). To simply delete a listing, without replacing it, send mail of the form

```
UPDATE: system!userid WAS john.a.smith-ho-3456
```

to the usual place. Changes in room number or location are picked up automatically from the Bell Labs telephone directory and need not be sent in. We hope that cooperating systems will, perhaps once a week, send in a current list of their users so that people will not need to send their own change of address forms. To aid in system administration, there is another update command `NEWPASSWD` which is used in the normal way, i.e.

```
UPDATE: system!userid NEWPASSWD john.a.smith-ho-3456
```

and is taken to mean "IS" if the userid does not already appear in the lists and "ANOTHER" if it does. This is thus an appropriate command when a new password file is received from a remote site and all the names in it, old and new, are to be fed to the updating program.

There are two ways this command can be made available on other systems. First, the data bases could be distributed. Although perfectly feasible, this involves distributing about 1,200,000 bytes of information. Secondly, a site can define *post* as

```
uux - "`postoffice`!post -u (`uname -n`!`logname`) $*"
```

where *uname* and *logname* are the standard USG commands which print, respectively, the system name and the user name. Systems deriving from Research versions will need to replace ``uname`` by the explicit name of the system and *logname* by the equivalent command, which is named *getuid*. This will invoke *post* remotely via *uux*; the mailroom system will then process the mail and forward it. In this case, diagnostics are returned when a direct invocation of the program would have resulted in a choice of several names.

We have specified the name of the mailroom computer through a *postoffice* command which prints its name. At present the appropriate name is *alice* and the *postoffice* command should just be *echo alice* or any equivalent. In the future the "mailroom" computer may be moved to a different site, depending on administrative decisions. Several sites might be used, particularly if it is desired to take advantage of local area networks in some locations. Another reason for using several sites might be to permit security-sensitive sites to avoid having to join even the UUCP public network and yet exchange mail among themselves. In either case, the data bases would have to be distributed, possibly with deletions for security reasons.

Local and Personal Lists

To avoid typing of long addresses, and to facilitate the maintenance of mailing lists, *post* permits users to maintain their own lists of abbreviations. All names are first looked for in a file `$HOME/lib/mail.aliases` where *\$HOME* is the invoker's login directory. This file, if it exists, should contain lines like

```
jane   elliott-127
sandy  a.g.fraser-mh
steve  s.c.johnson-1273
...
1273   jane sandy steve ...
```

and can be used either for giving mailing lists or expanding abbreviations. If an address matches the first word on any line, mail is sent to the names that follow on the line. The following names must either be in normal *post* addressee format, or must be ordinary *system!userid* names (discouraged, since it requires you to do your own list maintenance), or must be defined elsewhere in the file. Infinite loops of definitions are rejected (infinite is defined as five or greater). The entire local directory system requires, of course, that you actually have a copy of *post* on your system, and not merely the *uux* command to execute on a remote postoffice.

When *post* is unable to find a name in its telephone book, it can forward the request to another system. We imagine that systems are likely to maintain a database of their own users. This prevents local messages from requiring a remote call (the time for *uucp* between systems is 2 minutes or so). The non-local mail can still be forwarded to a central site. If this is desired, the local site must have its own copy of the *post* command, installed as such, and must have the *uux* line given above as a command *genpost*. The *genpost* command will be invoked by *post* for the names it can not handle locally.

At present, we envisage one central Bell Laboratories directory. However, if some group of systems is connected by a hardwired network, and wishes to speed up service, they may want to have a copy of the full directory. For example, it may be reasonable to provide another postoffice system at Indian Hill. It is also possible, if this system spreads to Bell System computers outside BTL, that there should be a master list of all participating UNIX systems and their users, to permit company to company electronic mail. The disadvantage of this hierarchical arrangement, of course, is that it delays the error message on ordinary misspellings of names. Some experience will be necessary to decide how best to arrange matters.

Future Plans

Conceivably, traffic may get so large that we will be forced to dedicate a computer system to electronic mail. An average mail message to a remote site takes 6.8 seconds of 11/70 CPU time (all but about 1 second is in ordinary mail). Although this could be reduced somewhat by avoiding as many process invocations as now occur, even so it seems there is a capacity for thousands of remote messages per day, which seems adequate for a long time. The mail traffic from people on our system, for example, is around 140 messages sent per day, of which 30 go to remote sites.

A dedicated machine could be rather small; it would need only perhaps a 5Mbyte disk, and all of the mail programs fit in about 30K address space. However, note that a large number of ACUs would be desirable; consider, for example, the 11/70, which takes 7 CPU seconds per call of 2 minutes real time each. This implies about 18 calls at a time, and if half are outgoing a need for 9 ACUs. An 11/23, which is much slower, might need perhaps 3 ACUs.

In the event that special hardware is dedicated to electronic mail, it may be necessary to charge separately in order to recover the costs. We thus envisage the eventual design and sale of UNIX stamps, to be used for electronic postage. However, following the example of the Trucial States and San Marino, we see no reason why these could not be sold to collectors, thus turning a net profit.

Acknowledgment

The help and encouragement of A. G. Fraser is gratefully acknowledged; also, the suggestions of R. F. Rosin have been most valuable.

R. Jane Elliott

Michael Lesk

