

Changes to The Rand MH Message Handling System: MH #6.5 for 4.3BSD UNIX

Marshall T. Rose
Northrop Research and Technology Center
One Research Park
Palos Verdes Peninsula, CA 90274

April 12, 1990

Abstract

This document describes the user-visible change to the UCI version of the Rand MH system that were made from `mh.5` to MH #6.5. It is based on the `mh.6` changes document, but has been updated to accurately reflect the MH distributed with 4.3BSD UNIX¹. This document does not describe bug-fixes, per se, or internal changes, unless these activities resulted in a visible change for the MH user.

This document is meant to supplement, not supersede, the standard MH User's manual[MRose85].

Comments concerning this documentation should be addressed to the Internet mailbox `Bug-MH@ICS.UCI.EDU`.

⁰This document (version #2.10) was L^AT_EXset April 12, 1990 with lplain v2.09-10/29/85.

¹UNIX is a trademark of AT&T Bell Laboratories.

Acknowledgements

The MH system described herein is based on the original Rand MH system. It has been extensively developed (perhaps too much so) by Marshall T. Rose and John L. Romine at the University of California, Irvine. Einar A. Stefferud, Jerry N. Sweet, and Terry P. Domae provided numerous suggestions to improve the UCI version of MH. Of course, a large number of people have helped MH along. The list of “MH immortals” is too long to list here. However, Van Jacobson deserves a special acknowledgement for his tireless work in improving the performance of MH. Some programs have been speeded-up by a factor of 10 or 20. All of users of MH, everywhere, owe a special thanks to Van.

Disclaimer

The Regents of the University of California wish to make it known that:

Although each program has been tested by its contributor, no warranty, express or implied, is made by the contributor or the University of California, as to the accuracy and functioning of the program and related program material, nor shall the fact of distribution constitute any such warranty, and no responsibility is assumed by the contributor or the University of California in connection herewith.

Conventions

In this document, certain L^AT_EX-formatting conventions are adhered to:

1. The names of UNIX commands, such as *comp*, are presented in *text italics*.
2. Arguments to programs, such as 'msgs', are presented in **typewriter style** and delimited by single-quotes.
3. UNIX pathnames and envariables, such as

/usr/uci/ and *\$SIGNATURE*,

are presented in *slanted roman*.

4. Text presenting an example, such as

`comp -editor zz`

is presented in **typewriter style**.

General Changes

Unlike the changes between **mh.4** and **mh.5**, a large number of user-visible changes have been made in **mh.6**. These changes have been in the form of bug fixes and several generalizations. The majority of these will not affect novice users. In addition, **mh.6** is a great deal faster than **mh.5**: all programs have been speeded-up significantly, thanks to work done by Van Jacobson as part of the process of including **mh.6** in the 4.3BSD UNIX distribution.

This document describes all user-visible changes to **mh.5** from its initial release to the intermediate release of **MH #6.5**.

System-5 Support

In addition to support for BSD UNIX, V7 UNIX and Xenix² variants of UNIX, **MH** finally has support for the AT&T variant of UNIX, System 5. Hopefully this will greatly expand the number of system which can run **MH**. Ironically, it appears that five ports of earlier versions of **MH** (including **mh.5**) were done, but news of the work was not widespread.³

Documentation

Several new documents have been included in the **mh.6** distribution: The paper *MH.5: How to process 200 messages a day and still get some real work done* was presented at the 1985 Summer Usenix Conference and Exhibition in Portland, Oregon. Another paper, *MH: A Multifarious User Agent*, has been accepted for publication by Computer Networks. Both describe **MH**, the former from a more technical and somewhat humorous perspective, the latter from a more serious and research-oriented perspective. In addition, a third paper has been included, *Design of the TTI Prototype Trusted Mail Agent*, which describes a so-called “trusted” mail agent built on top of **MH**. This paper was presented at the Second International Symposium on Computer Message Systems in Washington, D.C. A fourth paper, *MZnet: Mail Service for Personal Micro-Computer Systems*, is also included. This paper, which

²Xenix is a trademark of Microsoft Corporation.

³In fact, three groups in one large organization ported **MH** independently, each without knowledge of the others' work.

was presented at the First International Symposium on Computer Message Systems in Nottingham, U.K., describes a CP/M⁴-based version of MH.

In addition, the MH tutorial, *The Rand MH Message Handling System: Tutorial*, and, *The Rand MH Message Handling System: The UCI BBoards Facility*, have both been updated by Jerry N. Sweet.

For MH administrators (PostMasters and the like), there's an entirely new document, *The Rand MH Message Handling System: Administrator's Guide*. It explains most of the "ins and outs" of maintaining an MH system.

Finally, all of the manual entries and the MH manual have had a thorough working over. The documentation is expanded, more accurate, and more detailed.

Help Listings

When any MH command is invoked with the `-help` switch, in addition to listing the syntax of the command and version information, the MH configuration options will be listed. MH has so many configuration options, that when debugging problems, this information is invaluable.

The MH Profile

There are two new profile entries worth noting: **MH-Sequences** tells MH the name of the file to record public sequences in. Users of *vm*, a proprietary, visual front-end to MH, make use of this to disable the public sequences feature of MH.

The profile entry **Unseen-Sequence** names those sequences which should be defined as those messages recently incorporated by *inc*. The *show* program knows to remove messages from this sequence once it thinks they have been seen. If this profile entry is not present, or is empty, then no sequences are defined. Otherwise, for each name given, the sequence is first zero'd and then each message incorporated is added to the sequence. As such, this profile entry is rather analogous to the **Previous-Sequence** entry in the user's MH profile.

In addition, the **Alternate-Mailboxes** entry in the profile has been expanded to support simple wild-carding. Also, the default for this profile entry

⁴CP/M is a trademark of Digital Research Corporation.

is now the user's mail-id at any host. This change was made since MH can no longer reliably figure out what the user's real outgoing address looks like.

Finally, when the *install-mh* program is automatically invoked by MH, it won't prompt the user for information. Instead, it notes that it's setting up the default environment. In addition, the MH administrator may set-up a file called *mh.profile* in the MH library area which is consulted by *install-mh* when initializing the user's *.mh.profile*.

The MH Context

The *folder*, *scan*, and *show* programs have been modified to update the user's MH context prior to writing to the user's terminal. This allows the MH user interrupt output to the terminal and still have the expected context. This is especially useful to interrupt long *scan* listings. This change also introduces a subtle bug between *show* and messages denoted by the **Unseen-Sequence**. See *show*(1) for the details.

Addresses and 822 support

MH now fully supports the RFC-822 routing syntax for addresses (it used to recognize the syntax, but ignore the information present). In addition, there are three major modes for support of non-822 addressing in MH:

- **BERK**
This is useful on sites running **SendMail**. It doesn't support full 822-style addressing, in favor of recognizing such formats as ACSnet, and so on. For sites that can't run in an 822-compliant environment, this is the option to use (at the price of sacrificing some of the power of 822-style addressing). This also drastically reduces the address formatting facilities described below.
- **DUMB**
Although not as liberal as **BERK**, the **DUMB** option is useful on sites in which the message transport system conforms to the 822 standard, but wants to do all the defaulting itself.
- **BANG**
From out in left field, the **BANG** option favors **UUCP**-style addressing

over 822-style addressing. Hopefully when all the UUCP sites around get around to adopting domain-style addresses, this option won't be needed.

The *ap* program (mentioned momentarily) and the *ali* program both support a '-normalize' switch indicate if addresses should be resolved to their "official" hostnames.

New Programs

There are five new programs available: The *ap* program is the MH stand-alone address parser. It's useful for printing address in various formats (and for debugging address strings). The *dp* program is similar, but works on dates instead of addresses.

The *msgchk* program checks for new mail, possibly using the Post Office Protocol, POP, described below.

A new receive mail hook, the *rcvstore* program, which was written by Julian L. Onions is available.

Finally, a visual front-end to *msh* called *vmh* has been included. (This program is discussed in greater detail later on.)

Message Numbering

MH now no longer restricts the number of messages which may reside in a folder (beyond that of system memory constraints). This means that message numbers larger than 2000 are permissible. Hopefully this will make life easier for people reading the network news using MH.

The WhatNow Shell

In *mh.6*, there is now the concept of a unified **What now?** processor that the four composition programs, *comp*, *dist*, *forw*, and *repl* all invoke. This permits a greater flexibility in building mail applications with MH. As a result, there's a new program, *whatnow*, which acts as the default **What now?** program. Consult the *whatnow(1)* manual entry for all the details. The only important user-visible change is the **headers** option went away, which wasn't used that much anyway.

The only other thing worth noting is that unless MH has been compiled with the UCI option, the user's *\$HOME/.signature* file is not consulted for the user's personal name.

Format Strings

A general format string facility has been added to allow MH users to tailor the output of certain commands.

The *inc*, *scan*, *ap*, and *dp* programs all consult a file containing format strings. Format strings, which look a lot like *printf(3)* strings, give these MH commands precise instructions on how to format their output.

As a result, the *inc* and *scan* programs no longer have the `'-size'`, `'-nosize'`, `'-time'`, `'-notime'`, `'-numdate'`, and `'-nonumdate'` switches. These switches have been replaced with the `'-form_formatfile'` switch and the `'-format_string'` switch. The former directs the program to consult the named file for the format strings. The latter directs the program to use the named string as the format. To get the behavior of the old `'-time'` option, use the `'-form_scan.time'` option. Similarly, to get the effect of `'-size'`, use `'-form_scan.size'`.

A fun form to use is `'-form_scan.timely'` with *scan*. Try it sometime.

The *repl* command uses a file containing format files to indicate how the reply draft should be constructed. Note that reply templates prior to **mh.6** are incompatible with **mh.5**.⁵ Don't worry though, it's quite easy to convert the templates by hand. (Those clever enough to have written a reply template to begin with won't have *any* problem.)

Similarly, when the *forw* program is constructing a digest, it uses a file containing format strings to indicate how to build the encapsulating draft.

Finally, you can use these facilities in *mhl* as well.

News

The depreciated MH news system (from **mh.1**) is now de-supported. Use the "hoopy" BBoards facility instead.

⁵In fact, reply templates between **mh.6** and **MH #6.5** are incompatible.

BBoards

MH maintainers take note: the default home directory for the bboards login has changed from `/usr/bboards/` to `/usr/spool/bboards/`. Use the `bbhome` directive in your MH configuration file to set it back to the old value if you wish.

In addition, the `aliases` field for a BBoard in the BBoards file is now deemed useful only for addressing, not for user input to `bbc`. This means when giving the name of a BBoard to `bbc`, only the official name should be used.

A final note for mailsystem maintainers: the MMDF-II BBoards channel and the SendMail BBoards mailer have been modified to use the standard message encapsulation format when returning failed messages to the list maintainer. This means that the failure notices that the maintainer receives can simply be *burst*.

New Switches in `bbc`

The `bbc` program permits you to specify the `mshproc` to use on the command line by using the `'-mshproc_program'` option. There's also a `'-rcfile_file'` option which does "the obvious thing". In addition, options which aren't understood by `bbc` are passed along to the `mshproc`.

In addition, the following commands pass any unrecognized switches on to the program that they invoke: `bbc`, `next`, `show`, `prev`, and `vmh`.

Distributed BBoards

If both BBoards and POP (see the next section) are enabled, then distributed BBoards can be supported on top of the POP service. This allows the MH user to read BBoards on a server machine instead of the local host (which saves a lot of wasted disk space when the same BBoards are replicated several times at a site with several hosts). See the *Administrator's Guide* for information on how this can be made completely transparent to the MH user.

If you have several machines at your site running 4.2BSD UNIX and connected by an Ethernet⁶ (or other high-speed LAN), you *want* this software.

⁶Ethernet is a trademark of the Xerox Corporation.

Visual Front-End to msh

A simple window management protocol has been implemented for MH programs that might wish to act as a back-end to a sophisticated visual front-end.

The first implementation of a server side (front-end) program is *vmh*, which uses *curses*(3) to maintain a split-screen interface. Perhaps look for a *mhtool* program for the SUN next!

The *msh* program has been modified to speak the client side (back-end) of this protocol, if so directed. At present, *msh* is the only program in the MH distribution which implements the client side of the window management protocol.

Updates in msh

Prior to quitting, the *msh* command now asks if the *packf*'d file you've been perusing should be updated if you've modified it and the file is writable by you. The file can be modified by using *burst*, *rmm*, *rmm*, or *sortm* commands. The file can also be modified by using the *refile* command without the '-link' option. (Or course, the '-link' option doesn't actually link anything to the file.)

Distributed Mail

MH now contains a powerful facility for doing distributed mail (having MH reside on a host different than the message transport agent). For general information, consult either the *MH.5: How to process 200 messages a day and still get some real work done* paper, or the *MH: A Multifarious User Agent* paper. For specific information, consult the *Administrator's Guide*. Here's a brief synopsis:

This POP facility in MH is based on a modification of the ARPA Post Office Protocol (POP). A POP *subscriber* is a remote user, on a POP *client host*, that wishes to pick-up mail on a POP *service host*.

There are two ways to administer POP:

- Naive Mode

Each user-id in the *passwd*(5) file is considered a POP subscriber. No

changes are required for the mailsystem on the POP service host. However, this method requires that each POP subscriber have an entry in the password file. The POP server will fetch the user's mail from wherever maildrops are kept on the POP service host. This means that if maildrops are kept in the user's home directory, then each POP subscriber must have a home directory.

- Smart Mode

This is based on the notion that the list of POP subscribers and the list of login users are completely separate name spaces. A separate database (similar to the *BBoards*(5) file) is used to record information about each POP subscriber. Unfortunately, the local mailsystem must be changed to reflect this. This requires two changes (both of which are simple):

1. Aliasing

The aliasing mechanism is augmented so that POP subscriber addresses are diverted to a special delivery mechanism. **MH** comes with a program, *popaka*(8), which generates the additional information to be put in the mailsystem's alias file.

2. Delivery

A special POP delivery channel (for **MMDF-II**) or POP mailer (for **SendMail**) performs the actual delivery (**mh.6** supplies both). All it really does is just place the mail in the POP spool area.

Clever mailsystem people will note that the POP mechanism is really a special case of the more general **BBoards** mechanism.

These two different philosophies are not compatible on the same POP service host: one or the other, but not both, may be run.

In addition, there is one user-visible difference, which the administrator controls the availability of. The difference is whether the POP subscriber must supply a password to the POP server:

- ARPA standard method

This uses the standard ARPA technique of sending a username and a password. The appropriate programs (*inc*, *msgchk*, and possibly *bbc*) will prompt the user for this information.

- **UNIX remote method**
This uses the Berkeley UNIX reserved port method for authentication. This requires that the two or three mentioned above programs be *setuid* to root. (There are no known holes in any of these programs.)

These two different philosophies are compatible on the same POP service host: to selectively disable RPOP for hosts which aren't trusted, either modify the *.rhosts* file in the case of POP subscribers being UNIX logins, or zero the contents of network address field of the *pop(5)* file for the desired POP subscribers.

The *inc* command also has two other switches when MH is enabled for POP: `'-pack_file'` and `'-nopack'`. Normally, *inc* will use the POP to incorporate mail from a POP service host into an MH folder (`+inbox`). However, there are some misguided individuals who prefer to *msh* to read their maildrop. By using the `'-pack_file'` option, these individuals can direct *inc* to fetch their maildrop from the POP service host and store it locally in the named file. As expected, *inc* will treat the local file as a maildrop, performing the appropriate locking protocols. And, if the file doesn't exist, the user is now asked for confirmation.

Rcvmail hooks

In order to offer users of MH increased rcvmail hook functionality, the *slocal* program has been upgraded to support the semantics of the MMDF-II mail-delivery mechanism. This means that users of **mh.6** can maintain identical *.maildelivery* files regardless of the underlying transport system. See *mhook(1)* for all the details.

Change in rcvdist

The *rcvdist* rcvmail hook now uses the MH formatting facility when redistributing a message.

Field change in rcvpack

The *rcvpack* rcvmail hook now adds the field name `Delivery-Date:` instead of `Cron-Date:` to messages it *packs*.

GNU Emacs Support

James Larus' *mh-e* macro package for GNU Emacs (version 17) is included in the distribution. When loaded in Emacs, this provides a handy front-end.

Other Changes

Here's the miscellany:

Continuation Lines

Alias files used by MH, display templates used by *mhl*, and format files used by *forw*, *repl*, and *scan* all support a standard continuation line syntax. To continue a line in one of these files, simply end the line with the backslash character ('\'). All the other files used by MH are in 822-format, so the 822-continuation mechanism is used.⁷

Default Date Format

MH now uses numeric timezones instead of locally-meaningful alpha timezones when generating mail. This change was made to encourage the use of unambiguous, globally-meaningful timezone strings. A local configuration option can disable this correct behavior. All of the *mhl* templates have been modified to use locally-meaningful alpha timezones when displaying messages.

New switch in ali

The *ali* command now has a '-noalias' switch to prevent system-wide aliases from being interpreted.

Modifications to show

The '-format', '-noformat', '-pr', and '-nopr' options to *show* have gone away in favor of a more general mechanism. The '-showproc_program' option tells

⁷Looking back, it would have been best had all files in MH used the 822-format.

show (or *next* or *prev*) to use the named program as the `showproc`. The `-noshowproc` option tells *show*, et. al., to use the `cat(1)` program instead of a `showproc`. As a result, the profile entry `prproc` is no longer used.

Switch change in `inc`

The `-ms_ms-file` switch in *inc* has been changed to `-file_name` to be more consistent.

Front-End to `mhl`

When outputting to a terminal, the *mhl* program now runs the program denoted by the profile entry `moreproc`. If this entry is not present, the default is the UCB *more* program. If the entry is non-empty, then that program is spliced between *mhl* and the user's terminal. The author uses the *less* program as his `moreproc`.

Of course, if *mhl* isn't outputting to a terminal, then `moreproc` is not invoked.

Finally, to aid in the construction of replies, a prefix string may be specified for the `body` component of the message being replied-to. Simply use the `component=` construct in *mhl* for `body:`.

Confirmation in `packf`

If the file specified by the `-file_name` switch doesn't exist, the user is now asked for confirmation.

Complex Expressions in `pick`

The *pick* command now handles complex boolean expressions.

Defaults change in `prompter` and `burst`

The `-prepend` option is now the default in *prompter*. The `-noinplace` option is now the default in *burst*.

Fcc:s and post

If multiple Fcc:s for a message are specified during posting, *post* will try much harder to preserve links.

Interactive option in rmf

The *rmf* program has been changed to support an ‘-interactive’ switch. If given, then the user is prompted regarding whether the folder should be deleted. If the folder to be removed is not given by the user, this switch is defaulted to on.

Trusted Mail Interface

MH now has an interface for so-called “trusted mail” applications. Although the modifications to MH to support this are in the public domain, the actual library that MH uses is not. Contact Professor David J. Farber (Farber@UCI) for more information.

References

- [MRose85] Marshall T. Rose and John L. Romine. *The Rand MH Message Handling System: User’s Manual*. Department of Information and Computer Science, University of California, Irvine, mh.6 edition, November, 1985. UCI Version.