

Using maildir format

Why should I use maildir?

Two words: **no locks**. An MUA can read and delete messages while new mail is being delivered: each message is stored in a separate file with a unique name, so it isn't affected by operations on other messages. An MUA doesn't have to worry about partially delivered mail: each message is safely written to disk in the **tmp** subdirectory before it is moved to **new**. The maildir format is reliable even over NFS.

How are unique names created?

Unless you're writing messages to a maildir, the format of a unique name is none of your business. A unique name can be anything that doesn't contain a colon (or slash) and doesn't start with a dot. Do not try to extract information from unique names.

Okay, so you're writing messages. A unique name has three pieces, separated by dots. On the left is the result of `time()` or the second counter from `gettimeofday()`. On the right is the result of `gethostname()`. (To deal with invalid host names, replace `/` with `\057` and `:` with `\072`.) In the middle is a delivery identifier, discussed below.

As the terminology suggests, **every delivery to this maildir must have its own unique name**. When a maildir is shared through NFS, every machine that delivers to the maildir **must** have its own hostname. Within one machine, every delivery within the same second **must** have a different delivery identifier.

Modern delivery identifiers are created by concatenating enough of the following strings to guarantee uniqueness:

- `#n`, where `n` is (in hexadecimal) the output of the operating system's `unix_sequencenumber()` system call, which returns a number that increases by 1 every time it is called, starting from 0 after reboot.
- `Xn`, where `n` is (in hexadecimal) the output of the operating system's `unix_bootnumber()` system call, which reports the number of times that the system has been booted. Together with `#`, this guarantees uniqueness; unfortunately, most operating systems don't support `unix_sequencenumber()` and `unix_bootnumber`.
- `Rn`, where `n` is (in hexadecimal) the output of the operating system's `unix_cryptorandomnumber()` system call, or an equivalent source such as `/dev/urandom`. Unfortunately, some operating systems don't include cryptographic random number generators.
- `In`, where `n` is (in hexadecimal) the UNIX inode number of this file. Unfortunately, inode numbers aren't always available through NFS.
- `Vn`, where `n` is (in hexadecimal) the UNIX device number of this file. Unfortunately, device numbers aren't always available through NFS. (Device numbers are also not helpful with the standard UNIX filesystem: a maildir has to be within a single UNIX device for `link()` and `rename()` to work.)
- `Mn`, where `n` is (in decimal) the microsecond counter from the same `gettimeofday()` used for the left part of the unique name.
- `Pn`, where `n` is (in decimal) the process ID.

- Qn , where n is (in decimal) the number of deliveries made by this process.

Old-fashioned delivery identifiers use the following formats:

- n , where n is the process ID, and where this process has been forked to make one delivery. Unfortunately, some foolish operating systems repeat process IDs quickly, breaking the standard time+pid combination.
- n_m , where n is the process ID and m is the number of deliveries made by this process.

What can I put in info?

When you move a file from **new** to **cur**, you have to change its name from *uniq* to *uniq:info*. Make sure to preserve the *uniq* string, so that separate messages can't bump into each other.

info is morally equivalent to the Status field used by mbox readers. It'd be useful to have MUAs agree on the meaning of *info*, so I'm keeping a list of *info* semantics. Here it is.

info starting with "1,": Experimental semantics.

info starting with "2,": Each character after the comma is an independent flag.

- Flag "P" (passed): the user has resent/forwarded/bounced this message to someone else.
- Flag "R" (replied): the user has replied to this message.
- Flag "S" (seen): the user has viewed this message, though perhaps he didn't read all the way through it.
- Flag "T" (trashed): the user has moved this message to the trash; the trash will be emptied by a later user action.
- Flag "D" (draft): the user considers this message a draft; toggled at user discretion.
- Flag "F" (flagged): user-defined flag; toggled at user discretion.

New flags may be defined later. Flags must be stored in ASCII order: e.g., "2,FRS".

Can a maildir contain more than tmp, new, cur?

Yes:

- .qmail: used to do direct deliveries with qmail-local.
- bulletintime: empty file, used by system-wide bulletin programs.
- bulletinlock: empty file, used by system-wide bulletin programs.
- seriallock: empty file, used to serialize AutoTURN.