Gnus Manual

by Lars Magne Ingebrigtsen

The Gnus Newsreader

Gnus is the advanced, self-documenting, customizable, extensible unreal-time newsreader for GNU Emacs.

Oops. That sounds oddly familiar, so let's start over again to avoid being accused of plagiarism:

Gnus is a message-reading laboratory. It will let you look at just about anything as if it were a newsgroup. You can read mail with it, you can browse directories with it, you can `ftp` with it—you can even read news with it!

Gnus tries to empower people who read news the same way Emacs empowers people who edit text. Gnus sets no limits to what the user should be allowed to do. Users are encouraged to extend Gnus to make it behave like they want it to behave. A program should not control people; people should be empowered to do what they want by using (or abusing) the program.

1 Starting Gnus

If your system administrator has set things up properly, starting Gnus and reading news is extremely easy—you just type `M-x gnus` in your Emacs.

If you want to start Gnus in a different frame, you can use the command `M-x gnus-other-frame` instead.

If things do not go smoothly at startup, you have to twiddle some variables in your '`~/.gnus`' file. This file is similar to '`~/.emacs`', but is read when gnus starts.

If you puzzle at any terms used in this manual, please refer to the terminology section (see Section 10.4 [Terminology], page 229).

1.1 Finding the News

The `gnus-select-method` variable says where Gnus should look for news. This variable should be a list where the first element says *how* and the second element says *where*. This method is your native method. All groups not fetched with this method are foreign groups.

For instance, if the '`news.somewhere.edu`' NNTP server is where you want to get your daily dosage of news from, you'd say:

```
(setq gnus-select-method '(nntp "news.somewhere.edu"))
```

If you want to read directly from the local spool, say:

```
(setq gnus-select-method '(nnspool ""))
```

If you can use a local spool, you probably should, as it will almost certainly be much faster.

If this variable is not set, Gnus will take a look at the `NNTPSERVER` environment variable. If that variable isn't set, Gnus will see whether `gnus-nntpserver-file` ('`/etc/nntpserver`' by default) has any opinions on the matter. If that fails as well, Gnus will try to use the machine running Emacs as an NNTP server. That's a long shot, though.

If `gnus-nntp-server` is set, this variable will override `gnus-select-method`. You should therefore set `gnus-nntp-server` to `nil`, which is what it is by default.

You can also make Gnus prompt you interactively for the name of an NNTP server. If you give a non-numerical prefix to `gnus` (i.e., `C-u M-x gnus`), Gnus will let you choose between the servers in the `gnus-secondary-servers` list (if any). You can also just type in the name of any server you feel like visiting. (Note that this will set `gnus-nntp-server`, which means that if you then `M-x gnus` later in the same Emacs session, Gnus will contact the same server.)

However, if you use one NNTP server regularly and are just interested in a couple of groups from a different server, you would be better served by using the `B` command in the group buffer. It will let you have a look at what groups are available, and you can subscribe to any of the groups you want to. This also makes '`.newsrc`' maintenance much tidier. See Section 2.9 [Foreign Groups], page 21.

A slightly different approach to foreign groups is to set the `gnus-secondary-select-methods` variable. The select methods listed in this variable are in many ways just as native as the `gnus-select-method` server. They will also be queried for active files during startup

(if that's required), and new newsgroups that appear on these servers will be subscribed (or not) just as native groups are.

For instance, if you use the `nnmbox` back end to read your mail, you would typically set this variable to

```
(setq gnus-secondary-select-methods '((nnmbox "")))
```

1.2 The First Time

If no startup files exist, Gnus will try to determine what groups should be subscribed by default.

If the variable `gnus-default-subscribed-newsgroups` is set, Gnus will subscribe you to just those groups in that list, leaving the rest killed. Your system administrator should have set this variable to something useful.

Since she hasn't, Gnus will just subscribe you to a few arbitrarily picked groups (i.e., '`*.newusers`'). (*Arbitrary* is defined here as *whatever Lars thinks you should read*.)

You'll also be subscribed to the Gnus documentation group, which should help you with most common problems.

If `gnus-default-subscribed-newsgroups` is `t`, Gnus will just use the normal functions for handling new groups, and not do anything special.

1.3 The Server is Down

If the default server is down, Gnus will understandably have some problems starting. However, if you have some mail groups in addition to the news groups, you may want to start Gnus anyway.

Gnus, being the trusting sort of program, will ask whether to proceed without a native select method if that server can't be contacted. This will happen whether the server doesn't actually exist (i.e., you have given the wrong address) or the server has just momentarily taken ill for some reason or other. If you decide to continue and have no foreign groups, you'll find it difficult to actually do anything in the group buffer. But, hey, that's your problem. Blllrph!

If you know that the server is definitely down, or you just want to read your mail without bothering with the server at all, you can use the `gnus-no-server` command to start Gnus. That might come in handy if you're in a hurry as well. This command will not attempt to contact your primary server—instead, it will just activate all groups on level 1 and 2. (You should preferably keep no native groups on those two levels.)

1.4 Slave Gnusae

You might want to run more than one Emacs with more than one Gnus at the same time. If you are using different '`.newsrc`' files (e.g., if you are using the two different Gnusae to read from two different servers), that is no problem whatsoever. You just do it.

The problem appears when you want to run two Gnusae that use the same `.newsrc` file.

To work around that problem some, we here at the Think-Tank at the Gnus Towers have come up with a new concept: *Masters* and *slaves*. (We have applied for a patent on

this concept, and have taken out a copyright on those words. If you wish to use those words in conjunction with each other, you have to send \$1 per usage instance to me. Usage of the patent (*Master/Slave Relationships In Computer Applications*) will be much more expensive, of course.)

Anyway, you start one Gnus up the normal way with `M-x gnus` (or however you do it). Each subsequent slave Gnusae should be started with `M-x gnus-slave`. These slaves won't save normal '`.newsrc`' files, but instead save *slave files* that contain information only on what groups have been read in the slave session. When a master Gnus starts, it will read (and delete) these slave files, incorporating all information from them. (The slave files will be read in the sequence they were created, so the latest changes will have precedence.)

Information from the slave files has, of course, precedence over the information in the normal (i.e., master) `.newsrc` file.

1.5 Fetching a Group

It is sometimes convenient to be able to just say "I want to read this group and I don't care whether Gnus has been started or not". This is perhaps more useful for people who write code than for users, but the command `gnus-fetch-group` provides this functionality in any case. It takes the group name as a parameter.

1.6 New Groups

If you are satisfied that you really never want to see any new groups, you can set `gnus-check-new-newsgroups` to `nil`. This will also save you some time at startup. Even if this variable is `nil`, you can always subscribe to the new groups just by pressing `U` in the group buffer (see Section 2.13 [Group Maintenance], page 28). This variable is `ask-server` by default. If you set this variable to `always`, then Gnus will query the back ends for new groups even when you do the `g` command (see Section 2.17.1 [Scanning New Messages], page 34).

1.6.1 Checking New Groups

Gnus normally determines whether a group is new or not by comparing the list of groups from the active file(s) with the lists of subscribed and dead groups. This isn't a particularly fast method. If `gnus-check-new-newsgroups` is `ask-server`, Gnus will ask the server for new groups since the last time. This is both faster and cheaper. This also means that you can get rid of the list of killed groups altogether, so you may set `gnus-save-killed-list` to `nil`, which will save time both at startup, at exit, and all over. Saves disk space, too. Why isn't this the default, then? Unfortunately, not all servers support this command.

I bet I know what you're thinking now: How do I find out whether my server supports `ask-server`? No? Good, because I don't have a fail-safe answer. I would suggest just setting this variable to `ask-server` and see whether any new groups appear within the next few days. If any do, then it works. If none do, then it doesn't work. I could write a function to make Gnus guess whether the server supports `ask-server`, but it would just be a guess. So I won't. You could `telnet` to the server and say `HELP` and see whether it lists '`NEWGROUPS`' among the commands it understands. If it does, then it might work. (But there are servers that lists '`NEWGROUPS`' without supporting the function properly.)

This variable can also be a list of select methods. If so, Gnus will issue an `ask-server` command to each of the select methods, and subscribe them (or not) using the normal methods. This might be handy if you are monitoring a few servers for new groups. A side effect is that startup will take much longer, so you can meditate while waiting. Use the mantra "dingnusdingnusdingnus" to achieve permanent bliss.

1.6.2 Subscription Methods

What Gnus does when it encounters a new group is determined by the `gnus-subscribe-newsgroup-method` variable.

This variable should contain a function. This function will be called with the name of the new group as the only parameter.

Some handy pre-fab functions are:

`gnus-subscribe-zombies`

> Make all new groups zombies. This is the default. You can browse the zombies later (with `A z`) and either kill them all off properly (with `S z`), or subscribe to them (with `u`).

`gnus-subscribe-randomly`

> Subscribe all new groups in arbitrary order. This really means that all new groups will be added at "the top" of the group buffer.

`gnus-subscribe-alphabetically`

> Subscribe all new groups in alphabetical order.

`gnus-subscribe-hierarchically`

> Subscribe all new groups hierarchically. The difference between this function and `gnus-subscribe-alphabetically` is slight. `gnus-subscribe-alphabetically` will subscribe new groups in a strictly alphabetical fashion, while this function will enter groups into its hierarchy. So if you want to have the 'rec' hierarchy before the 'comp' hierarchy, this function will not mess that configuration up. Or something like that.

`gnus-subscribe-interactively`

> Subscribe new groups interactively. This means that Gnus will ask you about **all** new groups. The groups you choose to subscribe to will be subscribed hierarchically.

`gnus-subscribe-killed`

> Kill all new groups.

`gnus-subscribe-topics`

> Put the groups into the topic that has a matching `subscribe` topic parameter (see Section 2.16.5 [Topic Parameters], page 32). For instance, a `subscribe` topic parameter that looks like
>
>   "nnslashdot"
>
> will mean that all groups that match that regex will be subscribed under that topic.
>
> If no topics match the groups, the groups will be subscribed in the top-level topic.

A closely related variable is `gnus-subscribe-hierarchical-interactive`. (That's quite a mouthful.) If this variable is non-`nil`, Gnus will ask you in a hierarchical fashion whether to subscribe to new groups or not. Gnus will ask you for each sub-hierarchy whether you want to descend the hierarchy or not.

One common mistake is to set the variable a few paragraphs above (`gnus-subscribe-newsgroup-method`) to `gnus-subscribe-hierarchical-interactive`. This is an error. This will not work. This is ga-ga. So don't do it.

1.6.3 Filtering New Groups

A nice and portable way to control which new newsgroups should be subscribed (or ignored) is to put an *options* line at the start of the '`.newsrc`' file. Here's an example:

```
options -n !alt.all !rec.all sci.all
```

This line obviously belongs to a serious-minded intellectual scientific person (or she may just be plain old boring), because it says that all groups that have names beginning with '`alt`' and '`rec`' should be ignored, and all groups with names beginning with '`sci`' should be subscribed. Gnus will not use the normal subscription method for subscribing these groups. `gnus-subscribe-options-newsgroup-method` is used instead. This variable defaults to `gnus-subscribe-alphabetically`.

If you don't want to mess with your '`.newsrc`' file, you can just set the two variables `gnus-options-subscribe` and `gnus-options-not-subscribe`. These two variables do exactly the same as the '`.newsrc`' '`options -n`' trick. Both are regexps, and if the new group matches the former, it will be unconditionally subscribed, and if it matches the latter, it will be ignored.

Yet another variable that meddles here is `gnus-auto-subscribed-groups`. It works exactly like `gnus-options-subscribe`, and is therefore really superfluous, but I thought it would be nice to have two of these. This variable is more meant for setting some ground rules, while the other variable is used more for user fiddling. By default this variable makes all new groups that come from mail back ends (`nnml`, `nnbabyl`, `nnfolder`, `nnmbox`, and `nnmh`) subscribed. If you don't like that, just set this variable to `nil`.

New groups that match this regexp are subscribed using `gnus-subscribe-options-newsgroup-method`.

1.7 Changing Servers

Sometimes it is necessary to move from one NNTP server to another. This happens very rarely, but perhaps you change jobs, or one server is very flaky and you want to use another.

Changing the server is pretty easy, right? You just change `gnus-select-method` to point to the new server?

Wrong!

Article numbers are not (in any way) kept synchronized between different NNTP servers, and the only way Gnus keeps track of what articles you have read is by keeping track of article numbers. So when you change `gnus-select-method`, your '`.newsrc`' file becomes worthless.

Gnus provides a few functions to attempt to translate a '.newsrc' file from one server to another. They all have one thing in common—they take a looong time to run. You don't want to use these functions more than absolutely necessary.

If you have access to both servers, Gnus can request the headers for all the articles you have read and compare `Message-IDs` and map the article numbers of the read articles and article marks. The *M-x gnus-change-server* command will do this for all your native groups. It will prompt for the method you want to move to.

You can also move individual groups with the *M-x gnus-group-move-group-to-server* command. This is useful if you want to move a (foreign) group from one server to another.

If you don't have access to both the old and new server, all your marks and read ranges have become worthless. You can use the *M-x gnus-group-clear-data-on-native-groups* command to clear out all data that you have on your native groups. Use with caution.

After changing servers, you **must** move the cache hierarchy away, since the cached articles will have wrong article numbers, which will affect which articles Gnus thinks are read.

1.8 Startup Files

Now, you all know about the '.newsrc' file. All subscription information is traditionally stored in this file.

Things got a bit more complicated with GNUS. In addition to keeping the '.newsrc' file updated, it also used a file called '.newsrc.el' for storing all the information that didn't fit into the '.newsrc' file. (Actually, it also duplicated everything in the '.newsrc' file.) GNUS would read whichever one of these files was the most recently saved, which enabled people to swap between GNUS and other newsreaders.

That was kinda silly, so Gnus went one better: In addition to the '.newsrc' and '.newsrc.el' files, Gnus also has a file called '.newsrc.eld'. It will read whichever of these files that are most recent, but it will never write a '.newsrc.el' file. You should never delete the '.newsrc.eld' file—it contains much information not stored in the '.newsrc' file.

You can turn off writing the '.newsrc' file by setting `gnus-save-newsrc-file` to `nil`, which means you can delete the file and save some space, as well as exiting from Gnus faster. However, this will make it impossible to use other newsreaders than Gnus. But hey, who would want to, right? Similarly, setting `gnus-read-newsrc-file` to `nil` makes Gnus ignore the '.newsrc' file and any '.newsrc-SERVER' files, which is convenient if you have a tendency to use Netscape once in a while.

If `gnus-save-killed-list` (default `t`) is `nil`, Gnus will not save the list of killed groups to the startup file. This will save both time (when starting and quitting) and space (on disk). It will also mean that Gnus has no record of what groups are new or old, so the automatic new groups subscription methods become meaningless. You should always set `gnus-check-new-newsgroups` to `nil` or `ask-server` if you set this variable to `nil` (see Section 1.6 [New Groups], page 5). This variable can also be a regular expression. If that's the case, remove all groups that do not match this regexp before saving. This can be useful in certain obscure situations that involve several servers where not all servers support `ask-server`.

The `gnus-startup-file` variable says where the startup files are. The default value is '`~/.newsrc`', with the Gnus (El Dingo) startup file being whatever that one is, with a '`.eld`' appended.

`gnus-save-newsrc-hook` is called before saving any of the newsrc files, while `gnus-save-quick-newsrc-hook` is called just before saving the '`.newsrc.eld`' file, and `gnus-save-standard-newsrc-hook` is called just before saving the '`.newsrc`' file. The latter two are commonly used to turn version control on or off. Version control is on by default when saving the startup files. If you want to turn backup creation off, say something like:

```
(defun turn-off-backup ()
  (set (make-local-variable 'backup-inhibited) t))

(add-hook 'gnus-save-quick-newsrc-hook 'turn-off-backup)
(add-hook 'gnus-save-standard-newsrc-hook 'turn-off-backup)
```

When Gnus starts, it will read the `gnus-site-init-file` ('`.../site-lisp/gnus`' by default) and `gnus-init-file` ('`~/.gnus`' by default) files. These are normal Emacs Lisp files and can be used to avoid cluttering your '`~/.emacs`' and '`site-init`' files with Gnus stuff. Gnus will also check for files with the same names as these, but with '`.elc`' and '`.el`' suffixes. In other words, if you have set `gnus-init-file` to '`~/.gnus`', it will look for '`~/.gnus.elc`', '`~/.gnus.el`', and finally '`~/.gnus`' (in this order).

1.9 Auto Save

Whenever you do something that changes the Gnus data (reading articles, catching up, killing/subscribing groups), the change is added to a special *dribble buffer*. This buffer is auto-saved the normal Emacs way. If your Emacs should crash before you have saved the '`.newsrc`' files, all changes you have made can be recovered from this file.

If Gnus detects this file at startup, it will ask the user whether to read it. The auto save file is deleted whenever the real startup file is saved.

If `gnus-use-dribble-file` is `nil`, Gnus won't create and maintain a dribble buffer. The default is `t`.

Gnus will put the dribble file(s) in `gnus-dribble-directory`. If this variable is `nil`, which it is by default, Gnus will dribble into the directory where the '`.newsrc`' file is located. (This is normally the user's home directory.) The dribble file will get the same file permissions as the `.newsrc` file.

If `gnus-always-read-dribble-file` is non-`nil`, Gnus will read the dribble file on startup without querying the user.

1.10 The Active File

When Gnus starts, or indeed whenever it tries to determine whether new articles have arrived, it reads the active file. This is a very large file that lists all the active groups and articles on the server.

Before examining the active file, Gnus deletes all lines that match the regexp `gnus-ignored-newsgroups`. This is done primarily to reject any groups with bogus names, but you can use this variable to make Gnus ignore hierarchies you aren't ever interested in.

However, this is not recommended. In fact, it's highly discouraged. Instead, see Section 1.6 [New Groups], page 5 for an overview of other variables that can be used instead.

The active file can be rather Huge, so if you have a slow network, you can set `gnus-read-active-file` to `nil` to prevent Gnus from reading the active file. This variable is `some` by default.

Gnus will try to make do by getting information just on the groups that you actually subscribe to.

Note that if you subscribe to lots and lots of groups, setting this variable to `nil` will probably make Gnus slower, not faster. At present, having this variable `nil` will slow Gnus down considerably, unless you read news over a 2400 baud modem.

This variable can also have the value `some`. Gnus will then attempt to read active info only on the subscribed groups. On some servers this is quite fast (on sparkling, brand new INN servers that support the `LIST ACTIVE group` command), on others this isn't fast at all. In any case, `some` should be faster than `nil`, and is certainly faster than `t` over slow lines.

Some news servers (old versions of Leafnode and old versions of INN, for instance) do not support the `LIST ACTIVE group`. For these servers, `nil` is probably the most efficient value for this variable.

If this variable is `nil`, Gnus will ask for group info in total lock-step, which isn't very fast. If it is `some` and you use an NNTP server, Gnus will pump out commands as fast as it can, and read all the replies in one swoop. This will normally result in better performance, but if the server does not support the aforementioned `LIST ACTIVE group` command, this isn't very nice to the server.

If you think that starting up Gnus takes too long, try all the three different values for this variable and see what works best for you.

In any case, if you use `some` or `nil`, you should definitely kill all groups that you aren't interested in to speed things up.

Note that this variable also affects active file retrieval from secondary select methods.

1.11 Startup Variables

`gnus-load-hook`
>   A hook run while Gnus is being loaded. Note that this hook will normally be run just once in each Emacs session, no matter how many times you start Gnus.

`gnus-before-startup-hook`
>   A hook run after starting up Gnus successfully.

`gnus-startup-hook`
>   A hook run as the very last thing after starting up Gnus

`gnus-started-hook`
>   A hook that is run as the very last thing after starting up Gnus successfully.

`gnus-setup-news-hook`
>   A hook that is run after reading the '`.newsrc`' file(s), but before generating the group buffer.

`gnus-check-bogus-newsgroups`

> If non-`nil`, Gnus will check for and delete all bogus groups at startup. A *bogus group* is a group that you have in your '.newsrc' file, but doesn't exist on the news server. Checking for bogus groups can take quite a while, so to save time and resources it's best to leave this option off, and do the checking for bogus groups once in a while from the group buffer instead (see Section 2.13 [Group Maintenance], page 28).

`gnus-inhibit-startup-message`

> If non-`nil`, the startup message won't be displayed. That way, your boss might not notice as easily that you are reading news instead of doing your job. Note that this variable is used before '.gnus.el' is loaded, so it should be set in .emacs instead.

`gnus-no-groups-message`

> Message displayed by Gnus when no groups are available.

`gnus-play-startup-jingle`

> If non-`nil`, play the Gnus jingle at startup.

`gnus-startup-jingle`

> Jingle to be played if the above variable is non-`nil`. The default is 'Tuxedomoon.Jingle4.au'.

2 The Group Buffer

The *group buffer* lists all (or parts) of the available groups. It is the first buffer shown when Gnus starts, and will never be killed as long as Gnus is active.

2.1 Group Buffer Format

2.1.1 Group Line Specification

The default format of the group buffer is nice and dull, but you can make it as exciting and ugly as you feel like.

Here's a couple of example group lines:

```
        25: news.announce.newusers
   *     0: alt.fan.andrea-dworkin
```

Quite simple, huh?

You can see that there are 25 unread articles in 'news.announce.newusers'. There are no unread articles, but some ticked articles, in 'alt.fan.andrea-dworkin' (see that little asterisk at the beginning of the line?).

You can change that format to whatever you want by fiddling with the `gnus-group-line-format` variable. This variable works along the lines of a `format` specification, which is pretty much the same as a `printf` specifications, for those of you who use (feh!) C. See Section 8.4 [Formatting Variables], page 190.

'`%M%S%5y: %(%g%)\n`' is the value that produced those lines above.

There should always be a colon on the line; the cursor always moves to the colon after performing an operation. Nothing else is required—not even the group name. All displayed text is just window dressing, and is never examined by Gnus. Gnus stores all real information it needs using text properties.

(Note that if you make a really strange, wonderful, spreadsheet-like layout, everybody will believe you are hard at work with the accounting instead of wasting time reading news.)

Here's a list of all available format characters:

'M'   An asterisk if the group only has marked articles.

'S'   Whether the group is subscribed.

'L'   Level of subscribedness.

'N'   Number of unread articles.

'I'   Number of dormant articles.

'T'   Number of ticked articles.

'R'   Number of read articles.

't'   Estimated total number of articles. (This is really *max-number* minus *min-number* plus 1.)

'y'   Number of unread, unticked, non-dormant articles.

'i'   Number of ticked and dormant articles.

'g'   Full group name.

'G'   Group name.

'D'   Newsgroup description.

'o'   'm' if moderated.

'O'   '(m)' if moderated.

's'   Select method.

'n'   Select from where.

'z'   A string that looks like '<%s:%n>' if a foreign select method is used.

'P'   Indentation based on the level of the topic (see Section 2.16 [Group Topics],
                page 29).

'c'   Short (collapsed) group name. The `gnus-group-uncollapsed-levels` variable
                says how many levels to leave at the end of the group name. The default is
                1—this will mean that group names like 'gnu.emacs.gnus' will be shortened
                to 'g.e.gnus'.

'm'   '%' (`gnus-new-mail-mark`) if there has arrived new mail to the group lately.

'p'   '#' (`gnus-process-mark`) if the group is process marked.

'd'   A string that says when you last read the group (see Section 2.17.3 [Group
                Timestamp], page 35).

'u'   User defined specifier. The next character in the format string should be a letter.
                Gnus will call the function `gnus-user-format-function-`'X', where 'X' is the
                letter following '%u'. The function will be passed a single dummy parameter as
                argument. The function should return a string, which will be inserted into the
                buffer just like information from any other specifier.

    All the "number-of" specs will be filled with an asterisk ('*') if no info is available—for
instance, if it is a non-activated foreign group, or a bogus native group.

2.1.2 Group Modeline Specification

    The mode line can be changed by setting `gnus-group-mode-line-format` (see Section 8.4.2 [Mode Line Formatting], page 191). It doesn't understand that many format
specifiers:

'S'   The native news server.

'M'   The native select method.

2.1.3 Group Highlighting

Highlighting in the group buffer is controlled by the `gnus-group-highlight` variable. This is an alist with elements that look like (*form* . *face*). If *form* evaluates to something non-`nil`, the *face* will be used on the line.

Here's an example value for this variable that might look nice if the background is dark:

```
(cond (window-system
        (setq custom-background-mode 'light)
        (defface my-group-face-1
 '((t (:foreground "Red" :bold t))) "First group face")
        (defface my-group-face-2
 '((t (:foreground "DarkSeaGreen4" :bold t))) "Second group face")
        (defface my-group-face-3
 '((t (:foreground "Green4" :bold t))) "Third group face")
        (defface my-group-face-4
 '((t (:foreground "SteelBlue" :bold t))) "Fourth group face")
        (defface my-group-face-5
 '((t (:foreground "Blue" :bold t))) "Fifth group face")))

(setq gnus-group-highlight
        '(((> unread 200) . my-group-face-1)
          ((and (< level 3) (zerop unread)) . my-group-face-2)
          ((< level 3) . my-group-face-3)
          ((zerop unread) . my-group-face-4)
          (t . my-group-face-5)))
```

Also see Section 8.6 [Faces and Fonts], page 197.

Variables that are dynamically bound when the forms are evaluated include:

`group`	The group name.
`unread`	The number of unread articles in the group.
`method`	The select method.
`mailp`	Whether the group is a mail group.
`level`	The level of the group.
`score`	The score of the group.
`ticked`	The number of ticked articles in the group.
`total`	The total number of articles in the group. Or rather, MAX-NUMBER minus MIN-NUMBER plus one.
`topic`	When using the topic minor mode, this variable is bound to the current topic being inserted.

When the forms are `eval`ed, point is at the beginning of the line of the group in question, so you can use many of the normal Gnus functions for snarfing info on the group.

`gnus-group-update-hook` is called when a group line is changed. It will not be called when `gnus-visual` is `nil`. This hook calls `gnus-group-highlight-line` by default.

2.2 Group Maneuvering

All movement commands understand the numeric prefix and will behave as expected, hopefully.

n   Go to the next group that has unread articles (`gnus-group-next-unread-group`).

p
⟨DEL⟩   Go to the previous group that has unread articles (`gnus-group-prev-unread-group`).

N   Go to the next group (`gnus-group-next-group`).

P   Go to the previous group (`gnus-group-prev-group`).

M-n   Go to the next unread group on the same (or lower) level (`gnus-group-next-unread-group-same-level`).

M-p   Go to the previous unread group on the same (or lower) level (`gnus-group-prev-unread-group-same-level`).

Three commands for jumping to groups:

j   Jump to a group (and make it visible if it isn't already) (`gnus-group-jump-to-group`). Killed groups can be jumped to, just like living groups.

,   Jump to the unread group with the lowest level (`gnus-group-best-unread-group`).

.   Jump to the first group with unread articles (`gnus-group-first-unread-group`).

If `gnus-group-goto-unread` is `nil`, all the movement commands will move to the next group, not the next unread group. Even the commands that say they move to the next unread group. The default is `t`.

2.3 Selecting a Group

⟨SPC⟩   Select the current group, switch to the summary buffer and display the first unread article (`gnus-group-read-group`). If there are no unread articles in the group, or if you give a non-numerical prefix to this command, Gnus will offer to fetch all the old articles in this group from the server. If you give a numerical prefix N, N determines the number of articles Gnus will fetch. If N is positive, Gnus fetches the N newest articles, if N is negative, Gnus fetches the `abs(N)` oldest articles.

Thus, `SPC` enters the group normally, `C-u SPC` offers old articles, `C-u 4 2 SPC` fetches the 42 newest articles, and `C-u - 4 2 SPC` fetches the 42 oldest ones.

When you are in the group (in the Summary buffer), you can type `M-g` to fetch new articles, or `C-u M-g` to also show the old ones.

⟨RET⟩

> Select the current group and switch to the summary buffer (`gnus-group-select-group`). Takes the same arguments as `gnus-group-read-group`—the only difference is that this command does not display the first unread article automatically upon group entry.

`M-`⟨RET⟩   This does the same as the command above, but tries to do it with the minimum amount of fuzz (`gnus-group-quick-select-group`). No scoring/killing will be performed, there will be no highlights and no expunging. This might be useful if you're in a real hurry and have to enter some humongous group. If you give a 0 prefix to this command (i.e., `0 M-`⟨RET⟩), Gnus won't even generate the summary buffer, which is useful if you want to toggle threading before generating the summary buffer (see Section 3.25.3 [Summary Generation Commands], page 88).

`M-`⟨SPC⟩   This is yet one more command that does the same as the ⟨RET⟩ command, but this one does it without expunging and hiding dormants (`gnus-group-visible-select-group`).

`M-C-`⟨RET⟩   Finally, this command selects the current group ephemerally without doing any processing of its contents (`gnus-group-select-group-ephemerally`). Even threading has been turned off. Everything you do in the group after selecting it in this manner will have no permanent effects.

The `gnus-large-newsgroup` variable says what Gnus should consider to be a big group. This is 200 by default. If the group has more (unread and/or ticked) articles than this, Gnus will query the user before entering the group. The user can then specify how many articles should be fetched from the server. If the user specifies a negative number (`-n`), the `n` oldest articles will be fetched. If it is positive, the `n` articles that have arrived most recently will be fetched.

`gnus-auto-select-first` control whether any articles are selected automatically when entering a group with the ⟨SPC⟩ command.

`nil`   Don't select any articles when entering the group. Just display the full summary buffer.

`t`   Select the first unread article when entering the group.

`best`   Select the highest scored article in the group when entering the group.

This variable can also be a function. In that case, that function will be called to place point on a subject line, and/or select some article. Useful functions include:

`gnus-summary-first-unread-subject`

> Place point on the subject line of the first unread article, but don't select the article.

`gnus-summary-first-unread-article`

> Select the first unread article.

`gnus-summary-best-unread-article`

> Select the highest-scored unread article.

If you want to prevent automatic selection in some group (say, in a binary group with Huge articles) you can set this variable to `nil` in `gnus-select-group-hook`, which is called when a group is selected.

2.4 Subscription Commands

S t
u   Toggle subscription to the current group (`gnus-group-unsubscribe-current-group`).

S s
U   Prompt for a group to subscribe, and then subscribe it. If it was subscribed already, unsubscribe it instead (`gnus-group-unsubscribe-group`).

S k
C-k   Kill the current group (`gnus-group-kill-group`).

S y
C-y   Yank the last killed group (`gnus-group-yank-group`).

C-x C-t   Transpose two groups (`gnus-group-transpose-groups`). This isn't really a subscription command, but you can use it instead of a kill-and-yank sequence sometimes.

S w
C-w   Kill all groups in the region (`gnus-group-kill-region`).

S z   Kill all zombie groups (`gnus-group-kill-all-zombies`).

S C-k   Kill all groups on a certain level (`gnus-group-kill-level`). These groups can't be yanked back after killing, so this command should be used with some caution. The only time where this command comes in really handy is when you have a '`.newsrc`' with lots of unsubscribed groups that you want to get rid off. *S C-k* on level 7 will kill off all unsubscribed groups that do not have message numbers in the '`.newsrc`' file.

Also see Section 2.6 [Group Levels], page 19.

2.5 Group Data

c   Mark all unticked articles in this group as read (`gnus-group-catchup-current`). `gnus-group-catchup-group-hook` is called when catching up a group from the group buffer.

C   Mark all articles in this group, even the ticked ones, as read (`gnus-group-catchup-current-all`).

M-c   Clear the data from the current group—nix out marks and the list of read articles (`gnus-group-clear-data`).

M-x gnus-group-clear-data-on-native-groups
                   If you have switched from one NNTP server to another, all your marks and read ranges have become worthless. You can use this command to clear out all data that you have on your native groups. Use with caution.

2.6 Group Levels

All groups have a level of *subscribedness*. For instance, if a group is on level 2, it is more subscribed than a group on level 5. You can ask Gnus to just list groups on a given level or lower (see Section 2.11 [Listing Groups], page 25), or to just check for new articles in groups on a given level or lower (see Section 2.17.1 [Scanning New Messages], page 34).

Remember: The higher the level of the group, the less important it is.

`S l`   Set the level of the current group. If a numeric prefix is given, the next *n* groups will have their levels set. The user will be prompted for a level.

Gnus considers groups from levels 1 to `gnus-level-subscribed` (inclusive) (default 5) to be subscribed, `gnus-level-subscribed` (exclusive) and `gnus-level-unsubscribed` (inclusive) (default 7) to be unsubscribed, `gnus-level-zombie` to be zombies (walking dead) (default 8) and `gnus-level-killed` to be killed (completely dead) (default 9). Gnus treats subscribed and unsubscribed groups exactly the same, but zombie and killed groups have no information on what articles you have read, etc, stored. This distinction between dead and living groups isn't done because it is nice or clever, it is done purely for reasons of efficiency.

It is recommended that you keep all your mail groups (if any) on quite low levels (e.g. 1 or 2).

Maybe the following description of the default behavior of Gnus helps to understand what these levels are all about. By default, Gnus shows you subscribed nonempty groups, but by hitting `L` you can have it show empty subscribed groups and unsubscribed groups, too. Type `l` to go back to showing nonempty subscribed groups again. Thus, unsubscribed groups are hidden, in a way.

Zombie and killed groups are similar to unsubscribed groups in that they are hidden by default. But they are different from subscribed and unsubscribed groups in that Gnus doesn't ask the news server for information (number of messages, number of unread messages) on zombie and killed groups. Normally, you use `C-k` to kill the groups you aren't interested in. If most groups are killed, Gnus is faster.

Why does Gnus distinguish between zombie and killed groups? Well, when a new group arrives on the server, Gnus by default makes it a zombie group. This means that you are normally not bothered with new groups, but you can type `A z` to get a list of all new groups. Subscribe the ones you like and kill the ones you don't want. (`A k` shows a list of killed groups.)

If you want to play with the level variables, you should show some care. Set them once, and don't touch them ever again. Better yet, don't touch them at all unless you know exactly what you're doing.

Two closely related variables are `gnus-level-default-subscribed` (default 3) and `gnus-level-default-unsubscribed` (default 6), which are the levels that new groups will be put on if they are (un)subscribed. These two variables should, of course, be inside the relevant valid ranges.

If `gnus-keep-same-level` is non-`nil`, some movement commands will only move to groups of the same level (or lower). In particular, going from the last article in one group to the next group will go to the next group of the same level (or lower). This might be handy if you want to read the most important groups before you read the rest.

If this variable is `best`, Gnus will make the next newsgroup the one with the best level.

All groups with a level less than or equal to `gnus-group-default-list-level` will be listed in the group buffer by default.

If `gnus-group-list-inactive-groups` is non-`nil`, non-active groups will be listed along with the unread groups. This variable is `t` by default. If it is `nil`, inactive groups won't be listed.

If `gnus-group-use-permanent-levels` is non-`nil`, once you give a level prefix to `g` or `l`, all subsequent commands will use this level as the "work" level.

Gnus will normally just activate (i. e., query the server about) groups on level `gnus-activate-level` or less. If you don't want to activate unsubscribed groups, for instance, you might set this variable to 5. The default is 6.

2.7 Group Score

You would normally keep important groups on high levels, but that scheme is somewhat restrictive. Don't you wish you could have Gnus sort the group buffer according to how often you read groups, perhaps? Within reason?

This is what *group score* is for. You can have Gnus assign a score to each group through the mechanism described below. You can then sort the group buffer based on this score. Alternatively, you can sort on score and then level. (Taken together, the level and the score is called the *rank* of the group. A group that is on level 4 and has a score of 1 has a higher rank than a group on level 5 that has a score of 300. (The level is the most significant part and the score is the least significant part.))

If you want groups you read often to get higher scores than groups you read seldom you can add the `gnus-summary-bubble-group` function to the `gnus-summary-exit-hook` hook. This will result (after sorting) in a bubbling sort of action. If you want to see that in action after each summary exit, you can add `gnus-group-sort-groups-by-rank` or `gnus-group-sort-groups-by-score` to the same hook, but that will slow things down somewhat.

2.8 Marking Groups

If you want to perform some command on several groups, and they appear subsequently in the group buffer, you would normally just give a numerical prefix to the command. Most group commands will then do your bidding on those groups.

However, if the groups are not in sequential order, you can still perform a command on several groups. You simply mark the groups first with the process mark and then execute the command.

`#`
`M m`   Set the mark on the current group (`gnus-group-mark-group`).

`M-#`
`M u`   Remove the mark from the current group (`gnus-group-unmark-group`).

`M U`   Remove the mark from all groups (`gnus-group-unmark-all-groups`).

`M w`   Mark all groups between point and mark (`gnus-group-mark-region`).

`M b`   Mark all groups in the buffer (`gnus-group-mark-buffer`).

`M r`   Mark all groups that match some regular expression (`gnus-group-mark-regexp`).

Also see Section 8.1 [Process/Prefix], page 189.

If you want to execute some command on all groups that have been marked with the process mark, you can use the `M-&` (`gnus-group-universal-argument`) command. It will prompt you for the command to be executed.

2.9 Foreign Groups

Below are some group mode commands for making and editing general foreign groups, as well as commands to ease the creation of a few special-purpose groups. All these commands insert the newly created groups under point—`gnus-subscribe-newsgroup-method` is not consulted.

`G m`   Make a new group (`gnus-group-make-group`). Gnus will prompt you for a name, a method and possibly an *address*. For an easier way to subscribe to NNTP groups, see Section 2.14 [Browse Foreign Server], page 28.

`G r`   Rename the current group to something else (`gnus-group-rename-group`). This is valid only on some groups—mail groups mostly. This command might very well be quite slow on some back ends.

`G c`   Customize the group parameters (`gnus-group-customize`).

`G e`   Enter a buffer where you can edit the select method of the current group (`gnus-group-edit-group-method`).

`G p`   Enter a buffer where you can edit the group parameters (`gnus-group-edit-group-parameters`).

`G E`   Enter a buffer where you can edit the group info (`gnus-group-edit-group`).

`G d`   Make a directory group (see Section 6.5.1 [Directory Groups], page 141). You will be prompted for a directory name (`gnus-group-make-directory-group`).

`G h`   Make the Gnus help group (`gnus-group-make-help-group`).

`G a`   Make a Gnus archive group (`gnus-group-make-archive-group`). By default a group pointing to the most recent articles will be created (`gnus-group-recent-archive-directory`), but given a prefix, a full group will be created from `gnus-group-archive-directory`.

`G k`   Make a kiboze group. You will be prompted for a name, for a regexp to match groups to be "included" in the kiboze group, and a series of strings to match on headers (`gnus-group-make-kiboze-group`). See Section 6.6.2 [Kibozed Groups], page 156.

`G D`   Read an arbitrary directory as if it were a newsgroup with the `nneething` back end (`gnus-group-enter-directory`). See Section 6.5.2 [Anything Groups], page 142.

`G f`   Make a group based on some file or other (`gnus-group-make-doc-group`). If you give a prefix to this command, you will be prompted for a file name and a file type. Currently supported types are `babyl`, `mbox`, `digest`, `mmdf`, `news`, `rnews`, `clari-briefs`, `rfc934`, `rfc822-forward`, `nsmail` and `forward`. If you run this command without a prefix, Gnus will guess at the file type. See Section 6.5.3 [Document Groups], page 143.

`G u`   Create one of the groups mentioned in `gnus-useful-groups` (`gnus-group-make-useful-group`).

`G w`   Make an ephemeral group based on a web search (`gnus-group-make-web-group`). If you give a prefix to this command, make a solid group instead. You will be prompted for the search engine type and the search string. Valid search engine types include `dejanews`, `altavista` and `reference`. See Section 6.4.1 [Web Searches], page 138.

              If you use the `dejanews` search engine, you can limit the search to a particular group by using a match string like '`~g alt.sysadmin.recovery shaving`'.

`G ⎗DEL⎖`   This function will delete the current group (`gnus-group-delete-group`). If given a prefix, this function will actually delete all the articles in the group, and forcibly remove the group itself from the face of the Earth. Use a prefix only if you are absolutely sure of what you are doing. This command can't be used on read-only groups (like `nntp` group), though.

`G V`   Make a new, fresh, empty `nnvirtual` group (`gnus-group-make-empty-virtual`). See Section 6.6.1 [Virtual Groups], page 155.

`G v`   Add the current group to an `nnvirtual` group (`gnus-group-add-to-virtual`). Uses the process/prefix convention.

See Chapter 6 [Select Methods], page 105, for more information on the various select methods.

If `gnus-activate-foreign-newsgroups` is a positive number, Gnus will check all foreign groups with this level or lower at startup. This might take quite a while, especially if you subscribe to lots of groups from different NNTP servers. Also see Section 2.6 [Group Levels], page 19; `gnus-activate-level` also affects activation of foreign newsgroups.

2.10 Group Parameters

The group parameters store information local to a particular group. Here's an example group parameter list:

```
((to-address . "ding@gnus.org")
 (auto-expire . t))
```

We see that each element consists of a "dotted pair"—the thing before the dot is the key, while the thing after the dot is the value. All the parameters have this form *except* local variable specs, which are not dotted pairs, but proper lists.

The following group parameters can be used:

`to-address`
              Address used by when doing followups and new posts.

                    (to-address .   "some@where.com")

This is primarily useful in mail groups that represent closed mailing lists—mailing lists where it's expected that everybody that writes to the mailing list is subscribed to it. Since using this parameter ensures that the mail only goes to the mailing list itself, it means that members won't receive two copies of your followups.

Using `to-address` will actually work whether the group is foreign or not. Let's say there's a group on the server that is called 'fa.4ad-l'. This is a real newsgroup, but the server has gotten the articles from a mail-to-news gateway. Posting directly to this group is therefore impossible—you have to send mail to the mailing list address instead.

Some parameters have corresponding customizable variables, each of which is an alist of regexps and values.

to-list   Address used when doing `a` in that group.

                    (to-list . "some@where.com")

It is totally ignored when doing a followup—except that if it is present in a news group, you'll get mail group semantics when doing `f`.

If you do an `a` command in a mail group and you have neither a `to-list` group parameter nor a `to-address` group parameter, then a `to-list` group parameter will be added automatically upon sending the message if `gnus-add-to-list` is set to `t`.

If you do an `a` command in a mail group and you don't have a `to-list` group parameter, one will be added automatically upon sending the message.

See also `gnus-parameter-to-list-alist`.

visible   If the group parameter list has the element (`visible . t`), that group will always be visible in the Group buffer, regardless of whether it has any unread articles.

broken-reply-to
             Elements like (`broken-reply-to . t`) signals that `Reply-To` headers in this group are to be ignored. This can be useful if you're reading a mailing list group where the listserv has inserted `Reply-To` headers that point back to the listserv itself. This is broken behavior. So there!

to-group   Elements like (`to-group . "some.group.name"`) means that all posts in that group will be sent to `some.group.name`.

newsgroup
             If you have (`newsgroup . t`) in the group parameter list, Gnus will treat all responses as if they were responses to news articles. This can be useful if you have a mail group that's really a mirror of a news group.

gcc-self   If (`gcc-self . t`) is present in the group parameter list, newly composed messages will be Gcc'd to the current group. If (`gcc-self . none`) is present, no `Gcc:` header will be generated, if (`gcc-self . "string"`) is present, this string will be inserted literally as a `gcc` header. This parameter takes precedence over

any default `Gcc` rules as described later (see Section 5.4 [Archived Messages], page 100).

auto-expire
> If the group parameter has an element that looks like (`auto-expire . t`), all articles read will be marked as expirable. For an alternative approach, see Section 6.3.9 [Expiring Mail], page 128.

total-expire
> If the group parameter has an element that looks like (`total-expire . t`), all read articles will be put through the expiry process, even if they are not marked as expirable. Use with caution. Unread, ticked and dormant articles are not eligible for expiry.
>
> See also `gnus-total-expirable-newsgroups`.

expiry-wait
> If the group parameter has an element that looks like (`expiry-wait . 10`), this value will override any `nnmail-expiry-wait` and `nnmail-expiry-wait-function` when expiring expirable messages. The value can either be a number of days (not necessarily an integer) or the symbols `never` or `immediate`.

score-file
> Elements that look like (`score-file . "file"`) will make 'file' into the current score file for the group in question. All interactive score entries will be put into this file.

adapt-file
> Elements that look like (`adapt-file . "file"`) will make 'file' into the current adaptive file for the group in question. All adaptive score entries will be put into this file.

admin-address
> When unsubscribing from a mailing list you should never send the unsubscription notice to the mailing list itself. Instead, you'd send messages to the administrative address. This parameter allows you to put the admin address somewhere convenient.

display
> Elements that look like (`display . MODE`) say which articles to display on entering the group. Valid values are:
>
> all
> > Display all articles, both read and unread.
>
> default
> > Display the default visible articles, which normally includes unread and ticked articles.

comment
> Elements that look like (`comment . "This is a comment"`) are arbitrary comments on the group. They are currently ignored by Gnus, but provide a place for you to store information on particular groups.

charset
> Elements that look like (`charset . iso-8859-1`) will make `iso-8859-1` the default charset; that is, the charset that will be used for all articles that do not specify a charset.
>
> See also `gnus-group-charset-alist`.

ignored-charsets
>   Elements that look like (`ignored-charsets x-known iso-8859-1`) will make
>   `iso-8859-1` and `x-unknown` ignored; that is, the default charset will be used
>   for decoding articles.
>
>   See also `gnus-group-ignored-charsets-alist`.

posting-style
>   You can store additional posting style information for this group only here (see
>   Section 5.5 [Posting Styles], page 102). The format is that of an entry in the
>   `gnus-posting-styles` alist, except that there's no regexp matching the group
>   name (of course). Style elements in this group parameter will take precedence
>   over the ones found in `gnus-posting-styles`.
>
>   For instance, if you want a funky name and signature in this group only, instead
>   of hacking `gnus-posting-styles`, you could put something like this in the
>   group parameters:
>
> ```
> (posting-style
>   (name "Funky Name")
>   (signature "Funky Signature"))
> ```

banner   An item like (`banner . "regex"`) causes any part of an article that matches
          the regular expression "regex" to be stripped. Instead of "regex", you can also
          use the symbol `signature` which strips the last signature or any of the elements
          of the alist `gnus-article-banner-alist`.

(*variable form*)
>   You can use the group parameters to set variables local to the group you are
>   entering. If you want to turn threading off in 'news.answers', you could put
>   (`gnus-show-threads nil`) in the group parameters of that group. `gnus-show-`
>   `threads` will be made into a local variable in the summary buffer you enter,
>   and the form `nil` will be `eval`ed there.
>
>   This can also be used as a group-specific hook function, if you like. If you want
>   to hear a beep when you enter a group, you could put something like (`dummy-`
>   `variable (ding)`) in the parameters of that group. `dummy-variable` will be
>   set to the result of the (`ding`) form, but who cares?

Use the `G p` or the `G c` command to edit group parameters of a group. (`G p` presents
you with a Lisp-based interface, `G c` presents you with a Customize-like interface. The
latter helps avoid silly Lisp errors.) You might also be interested in reading about topic
parameters (see Section 2.16.5 [Topic Parameters], page 32).

2.11 Listing Groups

These commands all list various slices of the groups available.

`l`
`A s`   List all groups that have unread articles (`gnus-group-list-groups`). If the
          numeric prefix is used, this command will list only groups of level ARG and
          lower. By default, it only lists groups of level five (i. e., `gnus-group-default-`
          `list-level`) or lower (i.e., just subscribed groups).

`L`

`A u`   List all groups, whether they have unread articles or not (`gnus-group-list-all-groups`). If the numeric prefix is used, this command will list only groups of level ARG and lower. By default, it lists groups of level seven or lower (i.e., just subscribed and unsubscribed groups).

`A l`   List all unread groups on a specific level (`gnus-group-list-level`). If given a prefix, also list the groups with no unread articles.

`A k`   List all killed groups (`gnus-group-list-killed`). If given a prefix argument, really list all groups that are available, but aren't currently (un)subscribed. This could entail reading the active file from the server.

`A z`   List all zombie groups (`gnus-group-list-zombies`).

`A m`   List all unread, subscribed groups with names that match a regexp (`gnus-group-list-matching`).

`A M`   List groups that match a regexp (`gnus-group-list-all-matching`).

`A A`   List absolutely all groups in the active file(s) of the server(s) you are connected to (`gnus-group-list-active`). This might very well take quite a while. It might actually be a better idea to do a `A M` to list all matching, and just give '.' as the thing to match on. Also note that this command may list groups that don't exist (yet)—these will be listed as if they were killed groups. Take the output with some grains of salt.

`A a`   List all groups that have names that match a regexp (`gnus-group-apropos`).

`A d`   List all groups that have names or descriptions that match a regexp (`gnus-group-description-apropos`).

`A c`   List all groups with cached articles (`gnus-group-list-cached`).

`A ?`   List all groups with dormant articles (`gnus-group-list-dormant`).

Groups that match the `gnus-permanently-visible-groups` regexp will always be shown, whether they have unread articles or not. You can also add the `visible` element to the group parameters in question to get the same effect.

Groups that have just ticked articles in it are normally listed in the group buffer. If `gnus-list-groups-with-ticked-articles` is `nil`, these groups will be treated just like totally empty groups. It is `t` by default.

2.12 Sorting Groups

The `C-c C-s` (`gnus-group-sort-groups`) command sorts the group buffer according to the function(s) given by the `gnus-group-sort-function` variable. Available sorting functions include:

`gnus-group-sort-by-alphabet`
         Sort the group names alphabetically. This is the default.

`gnus-group-sort-by-real-name`
         Sort the group alphabetically on the real (unprefixed) group names.

`gnus-group-sort-by-level`

> Sort by group level.

`gnus-group-sort-by-score`

> Sort by group score. See Section 2.7 [Group Score], page 20.

`gnus-group-sort-by-rank`

> Sort by group score and then the group level. The level and the score are, when taken together, the group's *rank*. See Section 2.7 [Group Score], page 20.

`gnus-group-sort-by-unread`

> Sort by number of unread articles.

`gnus-group-sort-by-method`

> Sort alphabetically on the select method.

`gnus-group-sort-function` can also be a list of sorting functions. In that case, the most significant sort key function must be the last one.

There are also a number of commands for sorting directly according to some sorting criteria:

G S a   Sort the group buffer alphabetically by group name (`gnus-group-sort-groups-by-alphabet`).

G S u   Sort the group buffer by the number of unread articles (`gnus-group-sort-groups-by-unread`).

G S l   Sort the group buffer by group level (`gnus-group-sort-groups-by-level`).

G S v   Sort the group buffer by group score (`gnus-group-sort-groups-by-score`). See Section 2.7 [Group Score], page 20.

G S r   Sort the group buffer by group rank (`gnus-group-sort-groups-by-rank`). See Section 2.7 [Group Score], page 20.

G S m   Sort the group buffer alphabetically by back end name (`gnus-group-sort-groups-by-method`).

All the commands below obey the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

When given a symbolic prefix (see Section 8.3 [Symbolic Prefixes], page 190), all these commands will sort in reverse order.

You can also sort a subset of the groups:

G P a   Sort the groups alphabetically by group name (`gnus-group-sort-selected-groups-by-alphabet`).

G P u   Sort the groups by the number of unread articles (`gnus-group-sort-selected-groups-by-unread`).

G P l   Sort the groups by group level (`gnus-group-sort-selected-groups-by-level`).

G P v   Sort the groups by group score (`gnus-group-sort-selected-groups-by-score`). See Section 2.7 [Group Score], page 20.

`G P r`   Sort the groups by group rank (`gnus-group-sort-selected-groups-by-rank`). See Section 2.7 [Group Score], page 20.

`G P m`   Sort the groups alphabetically by back end name (`gnus-group-sort-selected-groups-by-method`).

2.13 Group Maintenance

`b`   Find bogus groups and delete them (`gnus-group-check-bogus-groups`).

`F`   Find new groups and process them (`gnus-group-find-new-groups`). With 1 `C-u`, use the `ask-server` method to query the server for new groups. With 2 `C-u`'s, use most complete method possible to query the server for new groups, and subscribe the new groups as zombies.

`C-c C-x`   Run all expirable articles in the current group through the expiry process (if any) (`gnus-group-expire-articles`).

`C-c M-C-x`   Run all articles in all groups through the expiry process (`gnus-group-expire-all-groups`).

2.14 Browse Foreign Server

`B`   You will be queried for a select method and a server name. Gnus will then attempt to contact this server and let you browse the groups there (`gnus-group-browse-foreign-server`).

A new buffer with a list of available groups will appear. This buffer will use the `gnus-browse-mode`. This buffer looks a bit (well, a lot) like a normal group buffer.

Here's a list of keystrokes available in the browse mode:

`n`   Go to the next group (`gnus-group-next-group`).

`p`   Go to the previous group (`gnus-group-prev-group`).

⟨SPC⟩   Enter the current group and display the first article (`gnus-browse-read-group`).

⟨RET⟩   Enter the current group (`gnus-browse-select-group`).

`u`   Unsubscribe to the current group, or, as will be the case here, subscribe to it (`gnus-browse-unsubscribe-current-group`).

`l`
`q`   Exit browse mode (`gnus-browse-exit`).

`?`   Describe browse mode briefly (well, there's not much to describe, is there) (`gnus-browse-describe-briefly`).

2.15 Exiting Gnus

Yes, Gnus is ex(c)iting.

z   Suspend Gnus (`gnus-group-suspend`). This doesn't really exit Gnus, but it
            kills all buffers except the Group buffer. I'm not sure why this is a gain, but
            then who am I to judge?

q   Quit Gnus (`gnus-group-exit`).

Q   Quit Gnus without saving the '`.newsrc`' files (`gnus-group-quit`). The dribble
            file will be saved, though (see Section 1.9 [Auto Save], page 9).

`gnus-suspend-gnus-hook` is called when you suspend Gnus and `gnus-exit-gnus-hook`
is called when you quit Gnus, while `gnus-after-exiting-gnus-hook` is called as the final
item when exiting Gnus.

If you wish to completely unload Gnus and all its adherents, you can use the `gnus-unload`
command. This command is also very handy when trying to customize meta-variables.

Note:
    Miss Lisa Cannifax, while sitting in English class, felt her feet go numbly heavy
    and herself fall into a hazy trance as the boy sitting behind her drew repeated
    lines with his pencil across the back of her plastic chair.

2.16 Group Topics

If you read lots and lots of groups, it might be convenient to group them hierarchically
according to topics. You put your Emacs groups over here, your sex groups over there, and
the rest (what, two groups or so?) you put in some misc section that you never bother
with anyway. You can even group the Emacs sex groups as a sub-topic to either the Emacs
groups or the sex groups—or both! Go wild!

Here's an example:
```
Gnus
  Emacs -- I wuw it!
     3: comp.emacs
     2: alt.religion.emacs
   Naughty Emacs
   452: alt.sex.emacs
     0: comp.talk.emacs.recovery
  Misc
     8: comp.binaries.fractals
    13: comp.sources.unix
```
To get this *fab* functionality you simply turn on (ooh!) the `gnus-topic` minor mode—
type `t` in the group buffer. (This is a toggling command.)

Go ahead, just try it. I'll still be here when you get back. La de dum... Nice tune, that...
la la la... What, you're back? Yes, and now press `l`. There. All your groups are now listed
under '`misc`'. Doesn't that make you feel all warm and fuzzy? Hot and bothered?

If you want this permanently enabled, you should add that minor mode to the hook for
the group mode:
```
(add-hook 'gnus-group-mode-hook 'gnus-topic-mode)
```

2.16.1 Topic Variables

Now, if you select a topic, it will fold/unfold that topic, which is really neat, I think.

The topic lines themselves are created according to the `gnus-topic-line-format` variable (see Section 8.4 [Formatting Variables], page 190). Valid elements are:

'i'   Indentation.

'n'   Topic name.

'v'   Visibility.

'l'   Level.

'g'   Number of groups in the topic.

'a'   Number of unread articles in the topic.

'A'   Number of unread articles in the topic and all its subtopics.

Each sub-topic (and the groups in the sub-topics) will be indented with `gnus-topic-indent-level` times the topic level number of spaces. The default is 2.

`gnus-topic-mode-hook` is called in topic minor mode buffers.

The `gnus-topic-display-empty-topics` says whether to display even topics that have no unread articles in them. The default is `t`.

2.16.2 Topic Commands

When the topic minor mode is turned on, a new `T` submap will be available. In addition, a few of the standard keys change their definitions slightly.

T n   Prompt for a new topic name and create it (`gnus-topic-create-topic`).

T m   Move the current group to some other topic (`gnus-topic-move-group`). This command uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

T j   Go to a topic (`gnus-topic-jump-to-topic`).

T c   Copy the current group to some other topic (`gnus-topic-copy-group`). This command uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

T h   Hide the current topic (`gnus-topic-hide-topic`). If given a prefix, hide the topic permanently.

T s   Show the current topic (`gnus-topic-show-topic`). If given a prefix, show the topic permanently.

T D   Remove a group from the current topic (`gnus-topic-remove-group`). This command is mainly useful if you have the same group in several topics and wish to remove it from one of the topics. You may also remove a group from all topics, but in that case, Gnus will add it to the root topic the next time you start Gnus. In fact, all new groups (which, naturally, don't belong to any topic) will show up in the root topic.

This command uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

T M   Move all groups that match some regular expression to a topic (`gnus-topic-move-matching`).

T C   Copy all groups that match some regular expression to a topic (`gnus-topic-copy-matching`).

T H   Toggle hiding empty topics (`gnus-topic-toggle-display-empty-topics`).

T #   Mark all groups in the current topic with the process mark (`gnus-topic-mark-topic`).

T M-#   Remove the process mark from all groups in the current topic (`gnus-topic-unmark-topic`).

T TAB

TAB   "Indent" the current topic so that it becomes a sub-topic of the previous topic (`gnus-topic-indent`). If given a prefix, "un-indent" the topic instead.

M-TAB   "Un-indent" the current topic so that it becomes a sub-topic of the parent of its current parent (`gnus-topic-unindent`).

⟨RET⟩

⟨SPC⟩   Either select a group or fold a topic (`gnus-topic-select-group`). When you perform this command on a group, you'll enter the group, as usual. When done on a topic line, the topic will be folded (if it was visible) or unfolded (if it was folded already). So it's basically a toggling command on topics. In addition, if you give a numerical prefix, group on that level (and lower) will be displayed.

C-c C-x   Run all expirable articles in the current group or topic through the expiry process (if any) (`gnus-topic-expire-articles`).

C-k   Kill a group or topic (`gnus-topic-kill-group`). All groups in the topic will be removed along with the topic.

C-y   Yank the previously killed group or topic (`gnus-topic-yank-group`). Note that all topics will be yanked before all groups.

T r   Rename a topic (`gnus-topic-rename`).

T ⟨DEL⟩   Delete an empty topic (`gnus-topic-delete`).

A T   List all groups that Gnus knows about in a topics-ified way (`gnus-topic-list-active`).

G p   Edit the topic parameters (`gnus-topic-edit-parameters`). See Section 2.16.5 [Topic Parameters], page 32.

2.16.3 Topic Sorting

You can sort the groups in each topic individually with the following commands:

T S a   Sort the current topic alphabetically by group name (`gnus-topic-sort-groups-by-alphabet`).

T S u   Sort the current topic by the number of unread articles (`gnus-topic-sort-`
              `groups-by-unread`).

T S l   Sort the current topic by group level (`gnus-topic-sort-groups-by-level`).

T S v   Sort the current topic by group score (`gnus-topic-sort-groups-by-score`).
              See Section 2.7 [Group Score], page 20.

T S r   Sort the current topic by group rank (`gnus-topic-sort-groups-by-rank`). See
              Section 2.7 [Group Score], page 20.

T S m   Sort the current topic alphabetically by back end name (`gnus-topic-sort-`
              `groups-by-method`).

See Section 2.12 [Sorting Groups], page 26, for more information about group sorting.

2.16.4 Topic Topology

So, let's have a look at an example group buffer:

```
Gnus
  Emacs -- I wuw it!
     3: comp.emacs
     2: alt.religion.emacs
    Naughty Emacs
     452: alt.sex.emacs
       0: comp.talk.emacs.recovery
  Misc
     8: comp.binaries.fractals
    13: comp.sources.unix
```

So, here we have one top-level topic ('`Gnus`'), two topics under that, and one sub-topic under one of the sub-topics. (There is always just one (1) top-level topic). This topology can be expressed as follows:

```
(("Gnus" visible)
 (("Emacs -- I wuw it!" visible)
  (("Naughty Emacs" visible)))
 (("Misc" visible)))
```

This is in fact how the variable `gnus-topic-topology` would look for the display above. That variable is saved in the '`.newsrc.eld`' file, and shouldn't be messed with manually— unless you really want to. Since this variable is read from the '`.newsrc.eld`' file, setting it in any other startup files will have no effect.

This topology shows what topics are sub-topics of what topics (right), and which topics are visible. Two settings are currently allowed—`visible` and `invisible`.

2.16.5 Topic Parameters

All groups in a topic will inherit group parameters from the parent (and ancestor) topic parameters. All valid group parameters are valid topic parameters (see Section 2.10 [Group Parameters], page 22).

In addition, the following parameters are only valid as topic parameters:

subscribe

> When subscribing new groups by topic (see Section 1.6.2 [Subscription Methods], page 6), the `subscribe` topic parameter says what groups go in what topic. Its value should be a regexp to match the groups that should go in that topic.

Group parameters (of course) override topic parameters, and topic parameters in subtopics override topic parameters in super-topics. You know. Normal inheritance rules. (*Rules* is here a noun, not a verb, although you may feel free to disagree with me here.)

```
    Gnus
      Emacs
         3: comp.emacs
         2: alt.religion.emacs
       452: alt.sex.emacs
        Relief
         452: alt.sex.emacs
           0: comp.talk.emacs.recovery
      Misc
         8: comp.binaries.fractals
        13: comp.sources.unix
       452: alt.sex.emacs
```

The 'Emacs' topic has the topic parameter (`score-file . "emacs.SCORE"`); the 'Relief' topic has the topic parameter (`score-file . "relief.SCORE"`); and the 'Misc' topic has the topic parameter (`score-file . "emacs.SCORE"`). In addition, 'alt.religion.emacs' has the group parameter (`score-file . "religion.SCORE"`).

Now, when you enter 'alt.sex.emacs' in the 'Relief' topic, you will get the 'relief.SCORE' home score file. If you enter the same group in the 'Emacs' topic, you'll get the 'emacs.SCORE' home score file. If you enter the group 'alt.religion.emacs', you'll get the 'religion.SCORE' home score file.

This seems rather simple and self-evident, doesn't it? Well, yes. But there are some problems, especially with the `total-expiry` parameter. Say you have a mail group in two topics; one with `total-expiry` and one without. What happens when you do *M-x gnus-expire-all-expirable-groups*? Gnus has no way of telling which one of these topics you mean to expire articles from, so anything may happen. In fact, I hereby declare that it is *undefined* what happens. You just have to be careful if you do stuff like that.

2.17 Misc Group Stuff

^

> Enter the server buffer (`gnus-group-enter-server-mode`). See Section 6.1 [The Server Buffer], page 105.

a

> Post an article to a group (`gnus-group-post-news`). If given a prefix, the current group name will be used as the default.

m

> Mail a message somewhere (`gnus-group-mail`).

Variables for the group buffer:

gnus-group-mode-hook

> is called after the group buffer has been created.

`gnus-group-prepare-hook`
>   is called after the group buffer is generated. It may be used to modify the buffer
>   in some strange, unnatural way.

`gnus-group-prepared-hook`
>   is called as the very last thing after the group buffer has been generated. It
>   may be used to move point around, for instance.

`gnus-permanently-visible-groups`
>   Groups matching this regexp will always be listed in the group buffer, whether
>   they are empty or not.

`gnus-group-name-charset-method-alist`
>   An alist of method and the charset for group names. It is used to show non-
>   ASCII group names.
>
>   For example:
>
> ```
> (setq gnus-group-name-charset-method-alist
>     '(((nntp "news.com.cn") . cn-gb-2312)))
> ```

`gnus-group-name-charset-group-alist`
>   An alist of regexp of group name and the charset for group names. It is used
>   to show non-ASCII group names.
>
>   For example:
>
> ```
> (setq gnus-group-name-charset-group-alist
>     '(("\\.com\\.cn:" . cn-gb-2312)))
> ```

2.17.1 Scanning New Messages

`g`   Check the server(s) for new articles. If the numerical prefix is used, this com-
                 mand will check only groups of level *arg* and lower (`gnus-group-get-new-
                 news`). If given a non-numerical prefix, this command will force a total re-
                 reading of the active file(s) from the back end(s).

`M-g`   Check whether new articles have arrived in the current group (`gnus-group-
                 get-new-news-this-group`). `gnus-goto-next-group-when-activating` says
                 whether this command is to move point to the next group or not. It is `t` by
                 default.

`C-c M-g`   Activate absolutely all groups (`gnus-activate-all-groups`).

`R`   Restart Gnus (`gnus-group-restart`). This saves the '.newsrc' file(s), closes
                 the connection to all servers, clears up all run-time Gnus variables, and then
                 starts Gnus all over again.

`gnus-get-new-news-hook` is run just before checking for new news.

`gnus-after-getting-new-news-hook` is run after checking for new news.

2.17.2 Group Information

H f   Try to fetch the FAQ for the current group (`gnus-group-fetch-faq`). Gnus will try to get the FAQ from `gnus-group-faq-directory`, which is usually a directory on a remote machine. This variable can also be a list of directories. In that case, giving a prefix to this command will allow you to choose between the various sites. `ange-ftp` (or `efs`) will be used for fetching the file.

                  If fetching from the first site is unsuccessful, Gnus will attempt to go through `gnus-group-faq-directory` and try to open them one by one.

H d
C-c C-d   Describe the current group (`gnus-group-describe-group`). If given a prefix, force Gnus to re-read the description from the server.

M-d   Describe all groups (`gnus-group-describe-all-groups`). If given a prefix, force Gnus to re-read the description file from the server.

H v
V   Display current Gnus version numbers (`gnus-version`).

?   Give a very short help message (`gnus-group-describe-briefly`).

C-c C-i   Go to the Gnus info node (`gnus-info-find-node`).

2.17.3 Group Timestamp

It can be convenient to let Gnus keep track of when you last read a group. To set the ball rolling, you should add `gnus-group-set-timestamp` to `gnus-select-group-hook`:

```
(add-hook 'gnus-select-group-hook 'gnus-group-set-timestamp)
```

After doing this, each time you enter a group, it'll be recorded.

This information can be displayed in various ways—the easiest is to use the '`%d`' spec in the group line format:

```
(setq gnus-group-line-format
      "%M\%S\%p\%P\%5y: %(%-40,40g%) %d\n")
```

This will result in lines looking like:

```
*        0: mail.ding                                    19961002T012943
         0: custom                                       19961002T012713
```

As you can see, the date is displayed in compact ISO 8601 format. This may be a bit too much, so to just display the date, you could say something like:

```
(setq gnus-group-line-format
      "%M\%S\%p\%P\%5y: %(%-40,40g%) %6,6~(cut 2)d\n")
```

2.17.4 File Commands

r   Re-read the init file (`gnus-init-file`, which defaults to '`~/.gnus`') (`gnus-group-read-init-file`).

s   Save the '`.newsrc.eld`' file (and '`.newsrc`' if wanted) (`gnus-group-save-newsrc`). If given a prefix, force saving the file(s) whether Gnus thinks it is necessary or not.

3 The Summary Buffer

A line for each article is displayed in the summary buffer. You can move around, read articles, post articles and reply to articles.

The most common way to a summary buffer is to select a group from the group buffer (see Section 2.3 [Selecting a Group], page 16).

You can have as many summary buffers open as you wish.

3.1 Summary Buffer Format

Gnus will use the value of the `gnus-extract-address-components` variable as a function for getting the name and address parts of a `From` header. Two pre-defined functions exist: `gnus-extract-address-components`, which is the default, quite fast, and too simplistic solution; and `mail-extract-address-components`, which works very nicely, but is slower. The default function will return the wrong answer in 5% of the cases. If this is unacceptable to you, use the other function instead:

```
(setq gnus-extract-address-components
      'mail-extract-address-components)
```

`gnus-summary-same-subject` is a string indicating that the current article has the same subject as the previous. This string will be used with those specs that require it. The default is `""`.

3.1.1 Summary Buffer Lines

You can change the format of the lines in the summary buffer by changing the `gnus-summary-line-format` variable. It works along the same lines as a normal `format` string, with some extensions (see Section 8.4 [Formatting Variables], page 190).

The default string is '`%U%R%z%I%(%[%4L: %-20,20n%]%) %s\n`'.

The following format specification characters are understood:

'N'   Article number.

'S'   Subject string. List identifiers stripped, `gnus-list-identifies`. See Section 3.16.3 [Article Hiding], page 71.

's'   Subject if the article is the root of the thread or the previous article had a different subject, `gnus-summary-same-subject` otherwise. (`gnus-summary-same-subject` defaults to `""`.)

'F'   Full `From` header.

'n'   The name (from the `From` header).

'f'   The name, code `To` header or the `Newsgroups` header (see Section 3.1.2 [To From Newsgroups], page 39).

'a'   The name (from the `From` header). This differs from the `n` spec in that it uses the function designated by the `gnus-extract-address-components` variable, which is slower, but may be more thorough.

'A'   The address (from the `From` header). This works the same way as the `a` spec.

'L'   Number of lines in the article.

'c'   Number of characters in the article. This specifier is not supported in some
               methods (like nnfolder).

'I'   Indentation based on thread level (see Section 3.8.1 [Customizing Threading],
               page 53).

'T'   Nothing if the article is a root and lots of spaces if it isn't (it pushes everything
               after it off the screen).

'['   Opening bracket, which is normally '[', but can also be '<' for adopted articles
               (see Section 3.8.1 [Customizing Threading], page 53).

']'   Closing bracket, which is normally ']', but can also be '>' for adopted articles.

'>'   One space for each thread level.

'<'   Twenty minus thread level spaces.

'U'   Unread.

'R'   This misleadingly named specifier is the *secondary mark*. This mark will say
               whether the article has been replied to, has been cached, or has been saved.

'i'   Score as a number (see Chapter 7 [Scoring], page 167).

'z'   Zcore, '+' if above the default level and '−' if below the default level. If the
               difference between `gnus-summary-default-score` and the score is less than
               `gnus-summary-zcore-fuzz`, this spec will not be used.

'V'   Total thread score.

'x'   `Xref`.

'D'   `Date`.

'd'   The `Date` in DD-MMM format.

'o'   The `Date` in *YYYYMMDD*T*HHMMSS* format.

'M'   `Message-ID`.

'r'   `References`.

't'   Number of articles in the current sub-thread. Using this spec will slow down
               summary buffer generation somewhat.

'e'   An '=' (`gnus-not-empty-thread-mark`) will be displayed if the article has any
               children.

'P'   The line number.

'O'   Download mark.

'u'   User defined specifier. The next character in the format string should be a
               letter. Gnus will call the function `gnus-user-format-function-`'X', where 'X'
               is the letter following '%u'. The function will be passed the current header as
               argument. The function should return a string, which will be inserted into the
               summary just like information from any other summary specifier.

The '`%U`' (status), '`%R`' (replied) and '`%z`' (zcore) specs have to be handled with care. For reasons of efficiency, Gnus will compute what column these characters will end up in, and "hard-code" that. This means that it is invalid to have these specs after a variable-length spec. Well, you might not be arrested, but your summary buffer will look strange, which is bad enough.

The smart choice is to have these specs as far to the left as possible. (Isn't that the case with everything, though? But I digress.)

This restriction may disappear in later versions of Gnus.

3.1.2 To From Newsgroups

In some groups (particularly in archive groups), the `From` header isn't very interesting, since all the articles there are written by you. To display the information in the `To` or `Newsgroups` headers instead, you need to decide three things: What information to gather; where to display it; and when to display it.

1. The reading of extra header information is controlled by the `gnus-extra-headers`. This is a list of header symbols. For instance:

   ```
   (setq gnus-extra-headers
         '(To Newsgroups X-Newsreader))
   ```

   This will result in Gnus trying to obtain these three headers, and storing it in header structures for later easy retrieval.

2. The value of these extra headers can be accessed via the `gnus-extra-header` function. Here's a format line spec that will access the `X-Newsreader` header:

   ```
   "%~(form (gnus-extra-header 'X-Newsreader))@"
   ```

3. The `gnus-ignored-from-addresses` variable says when the '`%f`' summary line spec returns the `To`, `Newsreader` or `From` header. If this regexp matches the contents of the `From` header, the value of the `To` or `Newsreader` headers are used instead.

A related variable is `nnmail-extra-headers`, which controls when to include extra headers when generating overview (NOV) files. If you have old overview files, you should regenerate them after changing this variable.

You also have to instruct Gnus to display the data by changing the `%n` spec to the `%f` spec in the `gnus-summary-line-format` variable.

In summary, you'd typically do something like the following:

```
(setq gnus-extra-headers
      '(To Newsgroups))
(setq nnmail-extra-headers gnus-extra-headers)
(setq gnus-summary-line-format
      "%U%R%z%I%(%[%4L: %-20,20f%]%) %s\n")
(setq gnus-ignored-from-addresses
      "Your Name Here")
```

Now, this is mostly useful for mail groups, where you have control over the NOV files that are created. However, if you can persuade your nntp admin to add:

```
Newsgroups:full
```

to the end of her '`overview.fmt`' file, then you can use that just as you would the extra headers from the mail groups.

3.1.3 Summary Buffer Mode Line

You can also change the format of the summary mode bar (see Section 8.4.2 [Mode Line Formatting], page 191). Set `gnus-summary-mode-line-format` to whatever you like. The default is '`Gnus: %%b [%A] %Z`'.

Here are the elements you can play with:

'`G`'   Group name.

'`p`'   Unprefixed group name.

'`A`'   Current article number.

'`z`'   Current article score.

'`V`'   Gnus version.

'`U`'   Number of unread articles in this group.

'`e`'   Number of unread articles in this group that aren't displayed in the summary buffer.

'`Z`'   A string with the number of unread and unselected articles represented either as '`<%U(+%e) more>`' if there are both unread and unselected articles, and just as '`<%U more>`' if there are just unread articles and no unselected ones.

'`g`'   Shortish group name. For instance, '`rec.arts.anime`' will be shortened to '`r.a.anime`'.

'`S`'   Subject of the current article.

'`u`'   User-defined spec (see Section 8.4.4 [User-Defined Specs], page 192).

'`s`'   Name of the current score file (see Chapter 7 [Scoring], page 167).

'`d`'   Number of dormant articles (see Section 3.6.1 [Unread Articles], page 47).

'`t`'   Number of ticked articles (see Section 3.6.1 [Unread Articles], page 47).

'`r`'   Number of articles that have been marked as read in this session.

'`E`'   Number of articles expunged by the score files.

3.1.4 Summary Highlighting

`gnus-visual-mark-article-hook`
          This hook is run after selecting an article. It is meant to be used for highlighting the article in some way. It is not run if `gnus-visual` is `nil`.

`gnus-summary-update-hook`
          This hook is called when a summary line is changed. It is not run if `gnus-visual` is `nil`.

`gnus-summary-selected-face`
          This is the face (or *font* as some people call it) used to highlight the current article in the summary buffer.

`gnus-summary-highlight`

>   Summary lines are highlighted according to this variable, which is a list where
>   the elements are of the format (*form* . *face*). If you would, for instance, like
>   ticked articles to be italic and high-scored articles to be bold, you could set this
>   variable to something like

```
(((eq mark gnus-ticked-mark) . italic)
 ((> score default) . bold))
```

>   As you may have guessed, if *form* returns a non-`nil` value, *face* will be applied
>   to the line.

3.2 Summary Maneuvering

All the straight movement commands understand the numeric prefix and behave pretty
much as you'd expect.

None of these commands select articles.

`G M-n`

`M-n`   Go to the next summary line of an unread article (`gnus-summary-next-unread-subject`).

`G M-p`

`M-p`   Go to the previous summary line of an unread article (`gnus-summary-prev-unread-subject`).

`G g`   Ask for an article number and then go to the summary line of that article
without displaying the article (`gnus-summary-goto-subject`).

If Gnus asks you to press a key to confirm going to the next group, you can use the `C-n`
and `C-p` keys to move around the group buffer, searching for the next group to read without
actually returning to the group buffer.

Variables related to summary movement:

`gnus-auto-select-next`

>   If you issue one of the movement commands (like `n`) and there are no more
>   unread articles after the current one, Gnus will offer to go to the next group.
>   If this variable is `t` and the next group is empty, Gnus will exit summary mode
>   and return to the group buffer. If this variable is neither `t` nor `nil`, Gnus will
>   select the next group, no matter whether it has any unread articles or not.
>   As a special case, if this variable is `quietly`, Gnus will select the next group
>   without asking for confirmation. If this variable is `almost-quietly`, the same
>   will happen only if you are located on the last article in the group. Finally, if
>   this variable is `slightly-quietly`, the `Z n` command will go to the next group
>   without confirmation. Also see Section 2.6 [Group Levels], page 19.

`gnus-auto-select-same`

>   If non-`nil`, all the movement commands will try to go to the next article with
>   the same subject as the current. (*Same* here might mean *roughly equal*. See
>   `gnus-summary-gather-subject-limit` for details (see Section 3.8.1 [Customiz-
>   ing Threading], page 53).) If there are no more articles with the same subject,
>   go to the first unread article.

This variable is not particularly useful if you use a threaded display.

`gnus-summary-check-current`

If non-`nil`, all the "unread" movement commands will not proceed to the next (or previous) article if the current article is unread. Instead, they will choose the current article.

`gnus-auto-center-summary`

If non-`nil`, Gnus will keep the point in the summary buffer centered at all times. This makes things quite tidy, but if you have a slow network connection, or simply do not like this un-Emacsism, you can set this variable to `nil` to get the normal Emacs scrolling action. This will also inhibit horizontal recentering of the summary buffer, which might make it more inconvenient to read extremely long threads.

This variable can also be a number. In that case, center the window at the given number of lines from the top.

3.3 Choosing Articles

3.3.1 Choosing Commands

None of the following movement commands understand the numeric prefix, and they all select and display an article.

If you want to fetch new articles or redisplay the group, see Section 3.26 [Exiting the Summary Buffer], page 89.

⟨SPC⟩   Select the current article, or, if that one's read already, the next unread article (`gnus-summary-next-page`).

`G n`
`n`   Go to next unread article (`gnus-summary-next-unread-article`).

`G p`
`p`   Go to previous unread article (`gnus-summary-prev-unread-article`).

`G N`
`N`   Go to the next article (`gnus-summary-next-article`).

`G P`
`P`   Go to the previous article (`gnus-summary-prev-article`).

`G C-n`   Go to the next article with the same subject (`gnus-summary-next-same-subject`).

`G C-p`   Go to the previous article with the same subject (`gnus-summary-prev-same-subject`).

`G f`
`.`   Go to the first unread article (`gnus-summary-first-unread-article`).

`G b`
`,`   Go to the article with the highest score (`gnus-summary-best-unread-article`).

G l

l   Go to the previous article read (`gnus-summary-goto-last-article`).

G o   Pop an article off the summary history and go to this article (`gnus-summary-pop-article`). This command differs from the command above in that you can pop as many previous articles off the history as you like, while *l* toggles the two last read articles. For a somewhat related issue (if you use these commands a lot), see Section 3.13 [Article Backlog], page 62.

G j

j   Ask for an article number or `Message-ID`, and then go to that article (`gnus-summary-goto-article`).

3.3.2 Choosing Variables

Some variables relevant for moving and selecting articles:

`gnus-auto-extend-newsgroup`
          All the movement commands will try to go to the previous (or next) article, even if that article isn't displayed in the Summary buffer if this variable is non-`nil`. Gnus will then fetch the article from the server and display it in the article buffer.

`gnus-select-article-hook`
          This hook is called whenever an article is selected. By default it exposes any threads hidden under the selected article.

`gnus-mark-article-hook`
          This hook is called whenever an article is selected. It is intended to be used for marking articles as read. The default value is `gnus-summary-mark-read-and-unread-as-read`, and will change the mark of almost any article you read to `gnus-unread-mark`. The only articles not affected by this function are ticked, dormant, and expirable articles. If you'd instead like to just have unread articles marked as read, you can use `gnus-summary-mark-unread-as-read` instead. It will leave marks like `gnus-low-score-mark`, `gnus-del-mark` (and so on) alone.

3.4 Scrolling the Article

⟨SPC⟩   Pressing ⟨SPC⟩ will scroll the current article forward one page, or, if you have come to the end of the current article, will choose the next article (`gnus-summary-next-page`).

⟨DEL⟩   Scroll the current article back one page (`gnus-summary-prev-page`).

⟨RET⟩   Scroll the current article one line forward (`gnus-summary-scroll-up`).

M-⟨RET⟩   Scroll the current article one line backward (`gnus-summary-scroll-down`).

A g

g   (Re)fetch the current article (`gnus-summary-show-article`). If given a prefix, fetch the current article, but don't run any of the article treatment functions. This will give you a "raw" article, just the way it came from the server.

If given a numerical prefix, you can do semi-manual charset stuff. `C-u 0 g cn-gb-2312` (RET) will decode the message as if it were encoded in the `cn-gb-2312` charset. If you have

```
(setq gnus-summary-show-article-charset-alist
      '((1 . cn-gb-2312)
        (2 . big5)))
```

then you can say `C-u 1 g` to get the same effect.

`A <`
`<`   Scroll to the beginning of the article (`gnus-summary-beginning-of-article`).

`A >`
`>`   Scroll to the end of the article (`gnus-summary-end-of-article`).

`A s`
`s`   Perform an isearch in the article buffer (`gnus-summary-isearch-article`).

`h`   Select the article buffer (`gnus-summary-select-article-buffer`).

3.5 Reply, Followup and Post

3.5.1 Summary Mail Commands

Commands for composing a mail message:

`S r`
`r`   Mail a reply to the author of the current article (`gnus-summary-reply`).

`S R`
`R`   Mail a reply to the author of the current article and include the original message (`gnus-summary-reply-with-original`). This command uses the process/prefix convention.

`S w`   Mail a wide reply to the author of the current article (`gnus-summary-wide-reply`). A *wide reply* is a reply that goes out to all people listed in the `To`, `From` (or `Reply-to`) and `Cc` headers.

`S W`   Mail a wide reply to the current article and include the original message (`gnus-summary-wide-reply-with-original`). This command uses the process/prefix convention.

`S o m`
`C-c C-f`   Forward the current article to some other person (`gnus-summary-mail-forward`). If no prefix is given, the message is forwarded according to the value of (`message-forward-as-mime`) and (`message-forward-show-mml`); if the prefix is 1, decode the message and forward directly inline; if the prefix is 2, forward message as an rfc822 MIME section; if the prefix is 3, decode message and forward as an rfc822 MIME section; if the prefix is 4, forward message directly inline; otherwise, the message is forwarded as no prefix given but use the flipped value of (`message-forward-as-mime`). By default, the message is decoded and forwarded as an rfc822 MIME section.

S m

m   Send a mail to some other person (`gnus-summary-mail-other-window`).

S D b   If you have sent a mail, but the mail was bounced back to you for some reason
            (wrong address, transient failure), you can use this command to resend that
            bounced mail (`gnus-summary-resend-bounced-mail`). You will be popped
            into a mail buffer where you can edit the headers before sending the mail off
            again. If you give a prefix to this command, and the bounced mail is a reply to
            some other mail, Gnus will try to fetch that mail and display it for easy perusal
            of its headers. This might very well fail, though.

S D r   Not to be confused with the previous command, `gnus-summary-resend-`
            `message` will prompt you for an address to send the current message off
            to, and then send it to that place. The headers of the message won't be
            altered—but lots of headers that say `Resent-To`, `Resent-From` and so on will
            be added. This means that you actually send a mail to someone that has a
            `To` header that (probably) points to yourself. This will confuse people. So,
            natcherly you'll only do that if you're really eVIl.

            This command is mainly used if you have several accounts and want to ship a
            mail to a different account of yours. (If you're both `root` and `postmaster` and
            get a mail for `postmaster` to the `root` account, you may want to resend it to
            `postmaster`. Ordnung mu sein!

            This command understands the process/prefix convention (see Section 8.1 [Pro-
            cess/Prefix], page 189).

S O m   Digest the current series (see Section 3.15 [Decoding Articles], page 65) and for-
            ward the result using mail (`gnus-uu-digest-mail-forward`). This command
            uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

S M-c   Send a complaint about excessive crossposting to the author of the current
            article (`gnus-summary-mail-crosspost-complaint`).

            This command is provided as a way to fight back against the current crosspost-
            ing pandemic that's sweeping Usenet. It will compose a reply using the `gnus-`
            `crosspost-complaint` variable as a preamble. This command understands the
            process/prefix convention (see Section 8.1 [Process/Prefix], page 189) and will
            prompt you before sending each mail.

Also see ⟨undefined⟩ [(message)Header Commands], page ⟨undefined⟩ for more informa-
tion.

3.5.2 Summary Post Commands

Commands for posting a news article:

S p

a   Post an article to the current group (`gnus-summary-post-news`).

S f

f   Post a followup to the current article (`gnus-summary-followup`).

S F
F   Post a followup to the current article and include the original message (**gnus-summary-followup-with-original**). This command uses the process/prefix convention.

S n   Post a followup to the current article via news, even if you got the message through mail (**gnus-summary-followup-to-mail**).

S N   Post a followup to the current article via news, even if you got the message through mail and include the original message (**gnus-summary-followup-to-mail-with-original**). This command uses the process/prefix convention.

S o p   Forward the current article to a newsgroup (**gnus-summary-post-forward**). If no prefix is given, the message is forwarded according to the value of (**message-forward-as-mime**) and (**message-forward-show-mml**); if the prefix is 1, decode the message and forward directly inline; if the prefix is 2, forward message as an rfc822 MIME section; if the prefix is 3, decode message and forward as an rfc822 MIME section; if the prefix is 4, forward message directly inline; otherwise, the message is forwarded as no prefix given but use the flipped value of (**message-forward-as-mime**). By default, the message is decoded and forwarded as an rfc822 MIME section.

S O p   Digest the current series and forward the result to a newsgroup (**gnus-uu-digest-mail-forward**). This command uses the process/prefix convention.

S u   Uuencode a file, split it into parts, and post it as a series (**gnus-uu-post-news**). (see Section 3.15.5.3 [Uuencoding and Posting], page 68).

Also see ⟨undefined⟩ [(message)Header Commands], page ⟨undefined⟩ for more information.

3.5.3 Summary Message Commands

S y   Yank the current article into an already existing Message composition buffer (**gnus-summary-yank-message**). This command prompts for what message buffer you want to yank into, and understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

3.5.4 Canceling Articles

Have you ever written something, and then decided that you really, really, really wish you hadn't posted that?

Well, you can't cancel mail, but you can cancel posts.

Find the article you wish to cancel (you can only cancel your own articles, so don't try any funny stuff). Then press *C* or *S c* (**gnus-summary-cancel-article**). Your article will be canceled—machines all over the world will be deleting your article. This command uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

Be aware, however, that not all sites honor cancels, so your article may live on here and there, while most sites will delete the article in question.

Gnus will use the "current" select method when canceling. If you want to use the standard posting method, use the 'a' symbolic prefix (see Section 8.3 [Symbolic Prefixes], page 190).

If you discover that you have made some mistakes and want to do some corrections, you can post a *superseding* article that will replace your original article.

Go to the original article and press `S s` (`gnus-summary-supersede-article`). You will be put in a buffer where you can edit the article all you want before sending it off the usual way.

The same goes for superseding as for canceling, only more so: Some sites do not honor superseding. On those sites, it will appear that you have posted almost the same article twice.

If you have just posted the article, and change your mind right away, there is a trick you can use to cancel/supersede the article without waiting for the article to appear on your site first. You simply return to the post buffer (which is called `*sent ...*`). There you will find the article you just posted, with all the headers intact. Change the `Message-ID` header to a `Cancel` or `Supersedes` header by substituting one of those words for the word `Message-ID`. Then just press `C-c C-c` to send the article as you would do normally. The previous article will be canceled/superseded.

Just remember, kids: There is no 'c' in 'supersede'.

3.6 Marking Articles

There are several marks you can set on an article.

You have marks that decide the *readedness* (whoo, neato-keano neologism ohoy!) of the article. Alphabetic marks generally mean *read*, while non-alphabetic characters generally mean *unread*.

In addition, you also have marks that do not affect readedness.

3.6.1 Unread Articles

The following marks mark articles as (kinda) unread, in one form or other.

'!'   Marked as ticked (`gnus-ticked-mark`).

            Ticked articles are articles that will remain visible always. If you see an article
            that you find interesting, or you want to put off reading it, or replying to it,
            until sometime later, you'd typically tick it. However, articles can be expired,
            so if you want to keep an article forever, you'll have to make it persistent (see
            Section 3.12 [Persistent Articles], page 61).

'?'   Marked as dormant (`gnus-dormant-mark`).

            Dormant articles will only appear in the summary buffer if there are followups
            to it. If you want to see them even if they don't have followups, you can use
            the `/ D` command (see Section 3.7 [Limiting], page 51).

'SPC'   Marked as unread (`gnus-unread-mark`).

            Unread articles are articles that haven't been read at all yet.

3.6.2 Read Articles

All the following marks mark articles as read.

'r'   These are articles that the user has marked as read with the `d` command man-
               ually, more or less (`gnus-del-mark`).

'R'   Articles that have actually been read (`gnus-read-mark`).

'O'   Articles that were marked as read in previous sessions and are now *old* (`gnus-
               ancient-mark`).

'K'   Marked as killed (`gnus-killed-mark`).

'X'   Marked as killed by kill files (`gnus-kill-file-mark`).

'Y'   Marked as read by having too low a score (`gnus-low-score-mark`).

'C'   Marked as read by a catchup (`gnus-catchup-mark`).

'G'   Canceled article (`gnus-canceled-mark`)

'F'   SOUPed article (`gnus-souped-mark`). See Section 6.5.4 [SOUP], page 145.

'Q'   Sparsely reffed article (`gnus-sparse-mark`). See Section 3.8.1 [Customizing
               Threading], page 53.

'M'   Article marked as read by duplicate suppression (`gnus-duplicated-mark`). See
               Section 3.28 [Duplicate Suppression], page 91.

All these marks just mean that the article is marked as read, really. They are interpreted
differently when doing adaptive scoring, though.

One more special mark, though:

'E'   Marked as expirable (`gnus-expirable-mark`).

               Marking articles as *expirable* (or have them marked as such automatically)
               doesn't make much sense in normal groups—a user doesn't control expiring of
               news articles, but in mail groups, for instance, articles marked as *expirable* can
               be deleted by Gnus at any time.

3.6.3 Other Marks

There are some marks that have nothing to do with whether the article is read or not.

- You can set a bookmark in the current article. Say you are reading a long thesis on
  cats' urinary tracts, and have to go home for dinner before you've finished reading
  the thesis. You can then set a bookmark in the article, and Gnus will jump to this
  bookmark the next time it encounters the article. See Section 3.6.4 [Setting Marks],
  page 49.

- All articles that you have replied to or made a followup to (i.e., have answered) will be
  marked with an 'A' in the second column (`gnus-replied-mark`).

- Articles stored in the article cache will be marked with an '*' in the second column
  (`gnus-cached-mark`). See Section 3.11 [Article Caching], page 61.

- Articles "saved" (in some manner or other; not necessarily religiously) are marked with
  an 'S' in the second column (`gnus-saved-mark`).

- If the '%e' spec is used, the presence of threads or not will be marked with `gnus-not-empty-thread-mark` and `gnus-empty-thread-mark` in the third column, respectively.
- Finally we have the *process mark* (`gnus-process-mark`). A variety of commands react to the presence of the process mark. For instance, `X u` (`gnus-uu-decode-uu`) will uudecode and view all articles that have been marked with the process mark. Articles marked with the process mark have a '`#`' in the second column.

You might have noticed that most of these "non-readedness" marks appear in the second column by default. So if you have a cached, saved, replied article that you have process-marked, what will that look like?

Nothing much. The precedence rules go as follows: process -> cache -> replied -> saved. So if the article is in the cache and is replied, you'll only see the cache mark and not the replied mark.

3.6.4 Setting Marks

All the marking commands understand the numeric prefix.

`M c`
`M-u`   Clear all readedness-marks from the current article (`gnus-summary-clear-mark-forward`). In other words, mark the article as unread.

`M t`
`!`   Tick the current article (`gnus-summary-tick-article-forward`). See Section 3.11 [Article Caching], page 61.

`M ?`
`?`   Mark the current article as dormant (`gnus-summary-mark-as-dormant`). See Section 3.11 [Article Caching], page 61.

`M d`
`d`   Mark the current article as read (`gnus-summary-mark-as-read-forward`).

`D`   Mark the current article as read and move point to the previous line (`gnus-summary-mark-as-read-backward`).

`M k`
`k`   Mark all articles that have the same subject as the current one as read, and then select the next unread article (`gnus-summary-kill-same-subject-and-select`).

`M K`
`C-k`   Mark all articles that have the same subject as the current one as read (`gnus-summary-kill-same-subject`).

`M C`   Mark all unread articles as read (`gnus-summary-catchup`).

`M C-c`   Mark all articles in the group as read—even the ticked and dormant articles (`gnus-summary-catchup-all`).

`M H`   Catchup the current group to point (`gnus-summary-catchup-to-here`).

`C-w`   Mark all articles between point and mark as read (`gnus-summary-mark-region-as-read`).

M V k	Kill all articles with scores below the default score (or below the numeric prefix) (`gnus-summary-kill-below`).
M e	
E	Mark the current article as expirable (`gnus-summary-mark-as-expirable`).
M b	Set a bookmark in the current article (`gnus-summary-set-bookmark`).
M B	Remove the bookmark from the current article (`gnus-summary-remove-bookmark`).
M V c	Clear all marks from articles with scores over the default score (or over the numeric prefix) (`gnus-summary-clear-above`).
M V u	Tick all articles with scores over the default score (or over the numeric prefix) (`gnus-summary-tick-above`).
M V m	Prompt for a mark, and mark all articles with scores over the default score (or over the numeric prefix) with this mark (`gnus-summary-clear-above`).

The `gnus-summary-goto-unread` variable controls what action should be taken after setting a mark. If non-`nil`, point will move to the next/previous unread article. If `nil`, point will just move one line up or down. As a special case, if this variable is `never`, all the marking commands as well as other commands (like $\overline{\text{SPC}}$) will move to the next article, whether it is unread or not. The default is `t`.

3.6.5 Generic Marking Commands

Some people would like the command that ticks an article (*!*) go to the next article. Others would like it to go to the next unread article. Yet others would like it to stay on the current article. And even though I haven't heard of anybody wanting it to go to the previous (unread) article, I'm sure there are people that want that as well.

Multiply these five behaviors by five different marking commands, and you get a potentially complex set of variable to control what each command should do.

To sidestep that mess, Gnus provides commands that do all these different things. They can be found on the *M M* map in the summary buffer. Type *M M C-h* to see them all—there are too many of them to list in this manual.

While you can use these commands directly, most users would prefer altering the summary mode keymap. For instance, if you would like the *!* command to go to the next article instead of the next unread article, you could say something like:

```
(add-hook 'gnus-summary-mode-hook 'my-alter-summary-map)
(defun my-alter-summary-map ()
  (local-set-key "!" 'gnus-summary-put-mark-as-ticked-next))
```

or

```
(defun my-alter-summary-map ()
  (local-set-key "!" "MM!n"))
```

3.6.6 Setting Process Marks

M P p
Mark the current article with the process mark (`gnus-summary-mark-as-processable`).

M P u
M-#   Remove the process mark, if any, from the current article (`gnus-summary-unmark-as-processable`).

M P U   Remove the process mark from all articles (`gnus-summary-unmark-all-processable`).

M P i   Invert the list of process marked articles (`gnus-uu-invert-processable`).

M P R   Mark articles that have a `Subject` header that matches a regular expression (`gnus-uu-mark-by-regexp`).

M P G   Unmark articles that have a `Subject` header that matches a regular expression (`gnus-uu-unmark-by-regexp`).

M P r   Mark articles in region (`gnus-uu-mark-region`).

M P t   Mark all articles in the current (sub)thread (`gnus-uu-mark-thread`).

M P T   Unmark all articles in the current (sub)thread (`gnus-uu-unmark-thread`).

M P v   Mark all articles that have a score above the prefix argument (`gnus-uu-mark-over`).

M P s   Mark all articles in the current series (`gnus-uu-mark-series`).

M P S   Mark all series that have already had some articles marked (`gnus-uu-mark-sparse`).

M P a   Mark all articles in series order (`gnus-uu-mark-series`).

M P b   Mark all articles in the buffer in the order they appear (`gnus-uu-mark-buffer`).

M P k   Push the current process mark set onto the stack and unmark all articles (`gnus-summary-kill-process-mark`).

M P y   Pop the previous process mark set from the stack and restore it (`gnus-summary-yank-process-mark`).

M P w   Push the current process mark set onto the stack (`gnus-summary-save-process-mark`).

Also see the `&` command in see Section 3.25.2 [Searching for Articles], page 88 for how to set process marks based on article body contents.

3.7 Limiting

It can be convenient to limit the summary buffer to just show some subset of the articles currently in the group. The effect most limit commands have is to remove a few (or many) articles from the summary buffer.

All limiting commands work on subsets of the articles already fetched from the servers. None of these commands query the server for additional articles.

/ /
/ s   Limit the summary buffer to articles that match some subject (`gnus-summary-limit-to-subject`).

/ a   Limit the summary buffer to articles that match some author (`gnus-summary-limit-to-author`).

/ x   Limit the summary buffer to articles that match one of the "extra" headers (see Section 3.1.2 [To From Newsgroups], page 39) (`gnus-summary-limit-to-extra`).

/ u
x   Limit the summary buffer to articles not marked as read (`gnus-summary-limit-to-unread`). If given a prefix, limit the buffer to articles strictly unread. This means that ticked and dormant articles will also be excluded.

/ m   Ask for a mark and then limit to all articles that have been marked with that mark (`gnus-summary-limit-to-marks`).

/ t   Ask for a number and then limit the summary buffer to articles older than (or equal to) that number of days (`gnus-summary-limit-to-age`). If given a prefix, limit to articles younger than that number of days.

/ n   Limit the summary buffer to the current article (`gnus-summary-limit-to-articles`). Uses the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

/ w   Pop the previous limit off the stack and restore it (`gnus-summary-pop-limit`). If given a prefix, pop all limits off the stack.

/ v   Limit the summary buffer to articles that have a score at or above some score (`gnus-summary-limit-to-score`).

/ E
M S   Include all expunged articles in the limit (`gnus-summary-limit-include-expunged`).

/ D   Include all dormant articles in the limit (`gnus-summary-limit-include-dormant`).

/ *   Include all cached articles in the limit (`gnus-summary-limit-include-cached`).

/ d   Exclude all dormant articles from the limit (`gnus-summary-limit-exclude-dormant`).

/ M   Exclude all marked articles (`gnus-summary-limit-exclude-marks`).

/ T   Include all the articles in the current thread in the limit.

/ c   Exclude all dormant articles that have no children from the limit (`gnus-summary-limit-exclude-childless-dormant`).

/ C   Mark all excluded unread articles as read (`gnus-summary-limit-mark-excluded-as-read`). If given a prefix, also mark excluded ticked and dormant articles as read.

3.8 Threading

Gnus threads articles by default. *To thread* is to put responses to articles directly after the articles they respond to—in a hierarchical fashion.

Threading is done by looking at the `References` headers of the articles. In a perfect world, this would be enough to build pretty trees, but unfortunately, the `References` header is often broken or simply missing. Weird news propagation exacerbates the problem, so one has to employ other heuristics to get pleasing results. A plethora of approaches exists, as detailed in horrible detail in see Section 3.8.1 [Customizing Threading], page 53.

First, a quick overview of the concepts:

root   The top-most article in a thread; the first article in the thread.

thread   A tree-like article structure.

sub-thread
             A small(er) section of this tree-like structure.

loose threads
             Threads often lose their roots due to article expiry, or due to the root already having been read in a previous session, and not displayed in the summary buffer. We then typically have many sub-threads that really belong to one thread, but are without connecting roots. These are called loose threads.

thread gathering
             An attempt to gather loose threads into bigger threads.

sparse threads
             A thread where the missing articles have been "guessed" at, and are displayed as empty lines in the summary buffer.

3.8.1 Customizing Threading

3.8.1.1 Loose Threads

`gnus-summary-make-false-root`
             If non-`nil`, Gnus will gather all loose subtrees into one big tree and create a dummy root at the top. (Wait a minute. Root at the top? Yup.) Loose subtrees occur when the real root has expired, or you've read or killed the root in a previous session.

             When there is no real root of a thread, Gnus will have to fudge something. This variable says what fudging method Gnus should use. There are four possible values:

             adopt   Gnus will make the first of the orphaned articles the parent. This parent will adopt all the other articles. The adopted articles will be marked as such by pointy brackets ('<>') instead of the standard square brackets ('[]'). This is the default method.

dummy   Gnus will create a dummy summary line that will pretend to be the
               parent. This dummy line does not correspond to any real article,
               so selecting it will just select the first real article after the dummy
               article. `gnus-summary-dummy-line-format` is used to specify the
               format of the dummy roots. It accepts only one format spec: 'S',
               which is the subject of the article. See Section 8.4 [Formatting
               Variables], page 190.

empty   Gnus won't actually make any article the parent, but simply leave
               the subject field of all orphans except the first empty. (Actually,
               it will use `gnus-summary-same-subject` as the subject (see Sec-
               tion 3.1 [Summary Buffer Format], page 37).)

none   Don't make any article parent at all. Just gather the threads and
               display them after one another.

nil   Don't gather loose threads.

`gnus-summary-gather-subject-limit`
               Loose threads are gathered by comparing subjects of articles. If this variable
               is `nil`, Gnus requires an exact match between the subjects of the loose threads
               before gathering them into one big super-thread. This might be too strict a
               requirement, what with the presence of stupid newsreaders that chop off long
               subject lines. If you think so, set this variable to, say, 20 to require that only
               the first 20 characters of the subjects have to match. If you set this variable to
               a really low number, you'll find that Gnus will gather everything in sight into
               one thread, which isn't very helpful.

               If you set this variable to the special value `fuzzy`, Gnus will use a fuzzy
               string comparison algorithm on the subjects (see Section 8.17 [Fuzzy Matching],
               page 208).

`gnus-simplify-subject-fuzzy-regexp`
               This can either be a regular expression or list of regular expressions that match
               strings that will be removed from subjects if fuzzy subject simplification is used.

`gnus-simplify-ignored-prefixes`
               If you set `gnus-summary-gather-subject-limit` to something as low as 10,
               you might consider setting this variable to something sensible:

```
(setq gnus-simplify-ignored-prefixes
      (concat
       "\\'\\[?\\("
       (mapconcat
        'identity
        '("looking"
          "wanted" "followup" "summary\\( of\\)?"
          "help" "query" "problem" "question"
          "answer" "reference" "announce"
          "How can I" "How to" "Comparison of"
          ;; ...
          )
```

```
                         "\\|")
                        "\\)\\s *\\("
                        (mapconcat 'identity
                                   '("for" "for reference" "with" "about")
                                   "\\|")
                        "\\)?\\]?:?[ \t]*"))
```

All words that match this regexp will be removed before comparing two subjects.

gnus-simplify-subject-functions

> If non-**nil**, this variable overrides **gnus-summary-gather-subject-limit**. This variable should be a list of functions to apply to the **Subject** string iteratively to arrive at the simplified version of the string.
>
> Useful functions to put in this list include:
>
> **gnus-simplify-subject-re**
>> Strip the leading 'Re:'.
>
> **gnus-simplify-subject-fuzzy**
>> Simplify fuzzily.
>
> **gnus-simplify-whitespace**
>> Remove excessive whitespace.
>
> You may also write your own functions, of course.

gnus-summary-gather-exclude-subject

> Since loose thread gathering is done on subjects only, that might lead to many false hits, especially with certain common subjects like '' and '**(none)**'. To make the situation slightly better, you can use the regexp **gnus-summary-gather-exclude-subject** to say what subjects should be excluded from the gathering process.
> The default is '**^ *$\\|^(none)$**'.

gnus-summary-thread-gathering-function

> Gnus gathers threads by looking at **Subject** headers. This means that totally unrelated articles may end up in the same "thread", which is confusing. An alternate approach is to look at all the **Message-IDs** in all the **References** headers to find matches. This will ensure that no gathered threads ever include unrelated articles, but it also means that people who have posted with broken newsreaders won't be gathered properly. The choice is yours—plague or cholera:
>
> **gnus-gather-threads-by-subject**
>> This function is the default gathering function and looks at **Subjects** exclusively.
>
> **gnus-gather-threads-by-references**
>> This function looks at **References** headers exclusively.
>
> If you want to test gathering by **References**, you could say something like:
>
> ```
> (setq gnus-summary-thread-gathering-function
>       'gnus-gather-threads-by-references)
> ```

3.8.1.2 Filling In Threads

`gnus-fetch-old-headers`

>   If non-`nil`, Gnus will attempt to build old threads by fetching more old
>   headers—headers to articles marked as read. If you would like to display as
>   few summary lines as possible, but still connect as many loose threads as
>   possible, you should set this variable to `some` or a number. If you set it to
>   a number, no more than that number of extra old headers will be fetched.
>   In either case, fetching old headers only works if the back end you are using
>   carries overview files—this would normally be `nntp`, `nnspool` and `nnml`. Also
>   remember that if the root of the thread has been expired by the server, there's
>   not much Gnus can do about that.
>
>   This variable can also be set to `invisible`. This won't have any visible effects,
>   but is useful if you use the `A T` command a lot (see Section 3.21 [Finding the
>   Parent], page 81).

`gnus-build-sparse-threads`

>   Fetching old headers can be slow. A low-rent similar effect can be gotten by
>   setting this variable to `some`. Gnus will then look at the complete `References`
>   headers of all articles and try to string together articles that belong in the same
>   thread. This will leave *gaps* in the threading display where Gnus guesses that
>   an article is missing from the thread. (These gaps appear like normal summary
>   lines. If you select a gap, Gnus will try to fetch the article in question.) If this
>   variable is `t`, Gnus will display all these "gaps" without regard for whether they
>   are useful for completing the thread or not. Finally, if this variable is `more`,
>   Gnus won't cut off sparse leaf nodes that don't lead anywhere. This variable is
>   `nil` by default.

3.8.1.3 More Threading

`gnus-show-threads`

>   If this variable is `nil`, no threading will be done, and all of the rest of the vari-
>   ables here will have no effect. Turning threading off will speed group selection
>   up a bit, but it is sure to make reading slower and more awkward.

`gnus-thread-hide-subtree`

>   If non-`nil`, all threads will be hidden when the summary buffer is generated.

`gnus-thread-expunge-below`

>   All threads that have a total score (as defined by `gnus-thread-score-`
>   `function`) less than this number will be expunged. This variable is `nil` by
>   default, which means that no threads are expunged.

`gnus-thread-hide-killed`

>   if you kill a thread and this variable is non-`nil`, the subtree will be hidden.

`gnus-thread-ignore-subject`

>   Sometimes somebody changes the subject in the middle of a thread. If this
>   variable is non-`nil`, the subject change is ignored. If it is `nil`, which is the
>   default, a change in the subject will result in a new thread.

gnus-thread-indent-level

> This is a number that says how much each sub-thread should be indented. The default is 4.

gnus-sort-gathered-threads-function

> Sometimes, particularly with mailing lists, the order in which mails arrive locally is not necessarily the same as the order in which they arrived on the mailing list. Consequently, when sorting sub-threads using the default `gnus-thread-sort-by-number`, responses can end up appearing before the article to which they are responding to. Setting this variable to an alternate value (e.g. `gnus-thread-sort-by-date`), in a group's parameters or in an appropriate hook (e.g. `gnus-summary-generate-hook`) can produce a more logical sub-thread ordering in such instances.

3.8.1.4 Low-Level Threading

gnus-parse-headers-hook

> Hook run before parsing any headers.

gnus-alter-header-function

> If non-`nil`, this function will be called to allow alteration of article header structures. The function is called with one parameter, the article header vector, which it may alter in any way. For instance, if you have a mail-to-news gateway which alters the `Message-ID`s in systematic ways (by adding prefixes and such), you can use this variable to un-scramble the `Message-ID`s so that they are more meaningful. Here's one example:

```
(setq gnus-alter-header-function 'my-alter-message-id)

(defun my-alter-message-id (header)
  (let ((id (mail-header-id header)))
    (when (string-match
            "\\(<[^<>@]*\\)\\.?cygnus\\..*@\\([^<>@]*>\\)" id)
      (mail-header-set-id
       (concat (match-string 1 id) "@" (match-string 2 id))
       header)))))
```

3.8.2 Thread Commands

T k
M-C-k   Mark all articles in the current (sub-)thread as read (**gnus-summary-kill-thread**). If the prefix argument is positive, remove all marks instead. If the prefix argument is negative, tick articles instead.

T l
M-C-l   Lower the score of the current (sub-)thread (**gnus-summary-lower-thread**).

T i   Increase the score of the current (sub-)thread (**gnus-summary-raise-thread**).

T #   Set the process mark on the current (sub-)thread (**gnus-uu-mark-thread**).

T M-#   Remove the process mark from the current (sub-)thread (`gnus-uu-unmark-thread`).

T T   Toggle threading (`gnus-summary-toggle-threads`).

T s   Expose the (sub-)thread hidden under the current article, if any (`gnus-summary-show-thread`).

T h   Hide the current (sub-)thread (`gnus-summary-hide-thread`).

T S   Expose all hidden threads (`gnus-summary-show-all-threads`).

T H   Hide all threads (`gnus-summary-hide-all-threads`).

T t   Re-thread the current article's thread (`gnus-summary-rethread-current`). This works even when the summary buffer is otherwise unthreaded.

T ^   Make the current article the child of the marked (or previous) article (`gnus-summary-reparent-thread`).

The following commands are thread movement commands. They all understand the numeric prefix.

T n
M-C-n
M-down   Go to the next thread (`gnus-summary-next-thread`).

T p
M-C-p
M-up   Go to the previous thread (`gnus-summary-prev-thread`).

T d   Descend the thread (`gnus-summary-down-thread`).

T u   Ascend the thread (`gnus-summary-up-thread`).

T o   Go to the top of the thread (`gnus-summary-top-thread`).

If you ignore subject while threading, you'll naturally end up with threads that have several different subjects in them. If you then issue a command like *T k* (`gnus-summary-kill-thread`) you might not wish to kill the entire thread, but just those parts of the thread that have the same subject as the current article. If you like this idea, you can fiddle with `gnus-thread-operation-ignore-subject`. If it is non-`nil` (which it is by default), subjects will be ignored when doing thread commands. If this variable is `nil`, articles in the same thread with different subjects will not be included in the operation in question. If this variable is `fuzzy`, only articles that have subjects fuzzily equal will be included (see Section 8.17 [Fuzzy Matching], page 208).

3.9 Sorting

If you are using a threaded summary display, you can sort the threads by setting `gnus-thread-sort-functions`, which can be either a single function, a list of functions, or a list containing functions and (`not some-function`) elements.

By default, sorting is done on article numbers. Ready-made sorting predicate functions include `gnus-thread-sort-by-number`, `gnus-thread-sort-by-author`, `gnus-thread-sort-by-subject`, `gnus-thread-sort-by-date`, `gnus-thread-sort-by-score`, and `gnus-thread-sort-by-total-score`.

Each function takes two threads and returns non-`nil` if the first thread should be sorted before the other. Note that sorting really is normally done by looking only at the roots of each thread.

If you use more than one function, the primary sort key should be the last function in the list. You should probably always include `gnus-thread-sort-by-number` in the list of sorting functions—preferably first. This will ensure that threads that are equal with respect to the other sort criteria will be displayed in ascending article order.

If you would like to sort by reverse score, then by subject, and finally by number, you could do something like:

```
(setq gnus-thread-sort-functions
      '(gnus-thread-sort-by-number
        gnus-thread-sort-by-subject
        (not gnus-thread-sort-by-total-score)))
```

The threads that have highest score will be displayed first in the summary buffer. When threads have the same score, they will be sorted alphabetically. The threads that have the same score and the same subject will be sorted by number, which is (normally) the sequence in which the articles arrived.

If you want to sort by score and then reverse arrival order, you could say something like:

```
(setq gnus-thread-sort-functions
      '((lambda (t1 t2)
          (not (gnus-thread-sort-by-number t1 t2)))
        gnus-thread-sort-by-score))
```

The function in the `gnus-thread-score-function` variable (default `+`) is used for calculating the total score of a thread. Useful functions might be `max`, `min`, or squared means, or whatever tickles your fancy.

If you are using an unthreaded display for some strange reason or other, you have to fiddle with the `gnus-article-sort-functions` variable. It is very similar to the `gnus-thread-sort-functions`, except that it uses slightly different functions for article comparison. Available sorting predicate functions are `gnus-article-sort-by-number`, `gnus-article-sort-by-author`, `gnus-article-sort-by-subject`, `gnus-article-sort-by-date`, and `gnus-article-sort-by-score`.

If you want to sort an unthreaded summary display by subject, you could say something like:

```
(setq gnus-article-sort-functions
      '(gnus-article-sort-by-number
        gnus-article-sort-by-subject))
```

3.10 Asynchronous Article Fetching

If you read your news from an NNTP server that's far away, the network latencies may make reading articles a chore. You have to wait for a while after pressing `n` to go to the next article before the article appears. Why can't Gnus just go ahead and fetch the article while you are reading the previous one? Why not, indeed.

First, some caveats. There are some pitfalls to using asynchronous article fetching, especially the way Gnus does it.

Let's say you are reading article 1, which is short, and article 2 is quite long, and you are not interested in reading that. Gnus does not know this, so it goes ahead and fetches article 2. You decide to read article 3, but since Gnus is in the process of fetching article 2, the connection is blocked.

To avoid these situations, Gnus will open two (count 'em two) connections to the server. Some people may think this isn't a very nice thing to do, but I don't see any real alternatives. Setting up that extra connection takes some time, so Gnus startup will be slower.

Gnus will fetch more articles than you will read. This will mean that the link between your machine and the NNTP server will become more loaded than if you didn't use article pre-fetch. The server itself will also become more loaded—both with the extra article requests, and the extra connection.

Ok, so now you know that you shouldn't really use this thing... unless you really want to.

Here's how: Set `gnus-asynchronous` to `t`. The rest should happen automatically.

You can control how many articles are to be pre-fetched by setting `gnus-use-article-prefetch`. This is 30 by default, which means that when you read an article in the group, the back end will pre-fetch the next 30 articles. If this variable is `t`, the back end will pre-fetch all the articles it can without bound. If it is `nil`, no pre-fetching will be done.

There are probably some articles that you don't want to pre-fetch—read articles, for instance. The `gnus-async-prefetch-article-p` variable controls whether an article is to be pre-fetched. This function should return non-`nil` when the article in question is to be pre-fetched. The default is `gnus-async-read-p`, which returns `nil` on read articles. The function is called with an article data structure as the only parameter.

If, for instance, you wish to pre-fetch only unread articles shorter than 100 lines, you could say something like:

```
(defun my-async-short-unread-p (data)
  "Return non-nil for short, unread articles."
  (and (gnus-data-unread-p data)
       (< (mail-header-lines (gnus-data-header data))
          100)))

(setq gnus-async-prefetch-article-p 'my-async-short-unread-p)
```

These functions will be called many, many times, so they should preferably be short and sweet to avoid slowing down Gnus too much. It's probably a good idea to byte-compile things like this.

Articles have to be removed from the asynch buffer sooner or later. The `gnus-prefetched-article-deletion-strategy` says when to remove articles. This is a list that may contain the following elements:

`read`   Remove articles when they are read.

`exit`   Remove articles when exiting the group.

The default value is (`read exit`).

3.11 Article Caching

If you have an *extremely* slow NNTP connection, you may consider turning article caching on. Each article will then be stored locally under your home directory. As you may surmise, this could potentially use *huge* amounts of disk space, as well as eat up all your inodes so fast it will make your head swim. In vodka.

Used carefully, though, it could be just an easier way to save articles.

To turn caching on, set `gnus-use-cache` to `t`. By default, all articles ticked or marked as dormant will then be copied over to your local cache (`gnus-cache-directory`). Whether this cache is flat or hierarchical is controlled by the `gnus-use-long-file-name` variable, as usual.

When re-selecting a ticked or dormant article, it will be fetched from the cache instead of from the server. As articles in your cache will never expire, this might serve as a method of saving articles while still keeping them where they belong. Just mark all articles you want to save as dormant, and don't worry.

When an article is marked as read, is it removed from the cache.

The entering/removal of articles from the cache is controlled by the `gnus-cache-enter-articles` and `gnus-cache-remove-articles` variables. Both are lists of symbols. The first is (`ticked dormant`) by default, meaning that ticked and dormant articles will be put in the cache. The latter is (`read`) by default, meaning that articles marked as read are removed from the cache. Possibly symbols in these two lists are `ticked`, `dormant`, `unread` and `read`.

So where does the massive article-fetching and storing come into the picture? The `gnus-jog-cache` command will go through all subscribed newsgroups, request all unread articles, score them, and store them in the cache. You should only ever, ever ever ever, use this command if 1) your connection to the NNTP server is really, really, really slow and 2) you have a really, really, really huge disk. Seriously. One way to cut down on the number of articles downloaded is to score unwanted articles down and have them marked as read. They will not then be downloaded by this command.

It is likely that you do not want caching on all groups. For instance, if your `nnml` mail is located under your home directory, it makes no sense to cache it somewhere else under your home directory. Unless you feel that it's neat to use twice as much space.

To limit the caching, you could set `gnus-cacheable-groups` to a regexp of groups to cache, '`^nntp`' for instance, or set the `gnus-uncacheable-groups` regexp to '`^nnml`', for instance. Both variables are `nil` by default. If a group matches both variables, the group is not cached.

The cache stores information on what articles it contains in its active file (`gnus-cache-active-file`). If this file (or any other parts of the cache) becomes all messed up for some reason or other, Gnus offers two functions that will try to set things right. `M-x gnus-cache-generate-nov-databases` will (re)build all the NOV files, and `gnus-cache-generate-active` will (re)generate the active file.

3.12 Persistent Articles

Closely related to article caching, we have *persistent articles*. In fact, it's just a different way of looking at caching, and much more useful in my opinion.

Say you're reading a newsgroup, and you happen on to some valuable gem that you want to keep and treasure forever. You'd normally just save it (using one of the many saving commands) in some file. The problem with that is that it's just, well, yucky. Ideally you'd prefer just having the article remain in the group where you found it forever; untouched by the expiry going on at the news server.

This is what a *persistent article* is—an article that just won't be deleted. It's implemented using the normal cache functions, but you use two explicit commands for managing persistent articles:

*   Make the current article persistent (`gnus-cache-enter-article`).

`M-*`   Remove the current article from the persistent articles (`gnus-cache-remove-article`). This will normally delete the article.

Both these commands understand the process/prefix convention.

To avoid having all ticked articles (and stuff) entered into the cache, you should set `gnus-use-cache` to `passive` if you're just interested in persistent articles:

        (setq gnus-use-cache 'passive)

3.13 Article Backlog

If you have a slow connection, but the idea of using caching seems unappealing to you (and it is, really), you can help the situation some by switching on the *backlog*. This is where Gnus will buffer already read articles so that it doesn't have to re-fetch articles you've already read. This only helps if you are in the habit of re-selecting articles you've recently read, of course. If you never do that, turning the backlog on will slow Gnus down a little bit, and increase memory usage some.

If you set `gnus-keep-backlog` to a number *n*, Gnus will store at most *n* old articles in a buffer for later re-fetching. If this variable is non-`nil` and is not a number, Gnus will store *all* read articles, which means that your Emacs will grow without bound before exploding and taking your machine down with you. I put that in there just to keep y'all on your toes.

This variable is `nil` by default.

3.14 Saving Articles

Gnus can save articles in a number of ways. Below is the documentation for saving articles in a fairly straight-forward fashion (i.e., little processing of the article is done before it is saved). For a different approach (uudecoding, unsharing) you should use `gnus-uu` (see Section 3.15 [Decoding Articles], page 65).

If `gnus-save-all-headers` is non-`nil`, Gnus will not delete unwanted headers before saving the article.

If the preceding variable is `nil`, all headers that match the `gnus-saved-headers` regexp will be kept, while the rest will be deleted before saving.

`O o`
`o`   Save the current article using the default article saver (`gnus-summary-save-article`).

`O m`   Save the current article in mail format (`gnus-summary-save-article-mail`).

`O r`   Save the current article in rmail format (`gnus-summary-save-article-rmail`).

`O f`   Save the current article in plain file format (`gnus-summary-save-article-file`).

`O F`   Write the current article in plain file format, overwriting any previous file contents (`gnus-summary-write-article-file`).

`O b`   Save the current article body in plain file format (`gnus-summary-save-article-body-file`).

`O h`   Save the current article in mh folder format (`gnus-summary-save-article-folder`).

`O v`   Save the current article in a VM folder (`gnus-summary-save-article-vm`).

`O p`   Save the current article in a pipe. Uhm, like, what I mean is—Pipe the current article to a process (`gnus-summary-pipe-output`).

All these commands use the process/prefix convention (see Section 8.1 [Process/Prefix], page 189). If you save bunches of articles using these functions, you might get tired of being prompted for files to save each and every article in. The prompting action is controlled by the `gnus-prompt-before-saving` variable, which is `always` by default, giving you that excessive prompting action you know and loathe. If you set this variable to `t` instead, you'll be prompted just once for each series of articles you save. If you like to really have Gnus do all your thinking for you, you can even set this variable to `nil`, which means that you will never be prompted for files to save articles in. Gnus will simply save all the articles in the default files.

You can customize the `gnus-default-article-saver` variable to make Gnus do what you want it to. You can use any of the six ready-made functions below, or you can create your own.

`gnus-summary-save-in-rmail`
         This is the default format, *babyl*. Uses the function in the `gnus-rmail-save-name` variable to get a file name to save the article in. The default is `gnus-plain-save-name`.

`gnus-summary-save-in-mail`
         Save in a Unix mail (mbox) file. Uses the function in the `gnus-mail-save-name` variable to get a file name to save the article in. The default is `gnus-plain-save-name`.

`gnus-summary-save-in-file`
         Append the article straight to an ordinary file. Uses the function in the `gnus-file-save-name` variable to get a file name to save the article in. The default is `gnus-numeric-save-name`.

`gnus-summary-save-body-in-file`
         Append the article body to an ordinary file. Uses the function in the `gnus-file-save-name` variable to get a file name to save the article in. The default is `gnus-numeric-save-name`.

`gnus-summary-save-in-folder`
>   Save the article to an MH folder using `rcvstore` from the MH library. Uses
>   the function in the `gnus-folder-save-name` variable to get a file name to save
>   the article in. The default is `gnus-folder-save-name`, but you can also use
>   `gnus-Folder-save-name`, which creates capitalized names.

`gnus-summary-save-in-vm`
>   Save the article in a VM folder. You have to have the VM mail reader to use
>   this setting.

All of these functions, except for the last one, will save the article in the `gnus-article-save-directory`, which is initialized from the `SAVEDIR` environment variable. This is '`~/News/`' by default.

As you can see above, the functions use different functions to find a suitable name of a file to save the article in. Below is a list of available functions that generate names:

`gnus-Numeric-save-name`
>   File names like '`~/News/Alt.andrea-dworkin/45`'.

`gnus-numeric-save-name`
>   File names like '`~/News/alt.andrea-dworkin/45`'.

`gnus-Plain-save-name`
>   File names like '`~/News/Alt.andrea-dworkin`'.

`gnus-plain-save-name`
>   File names like '`~/News/alt.andrea-dworkin`'.

You can have Gnus suggest where to save articles by plonking a regexp into the `gnus-split-methods` alist. For instance, if you would like to save articles related to Gnus in the file '`gnus-stuff`', and articles related to VM in `vm-stuff`, you could set this variable to something like:

```
(("^Subject:.*gnus\\|^Newsgroups:.*gnus" "gnus-stuff")
 ("^Subject:.*vm\\|^Xref:.*vm" "vm-stuff")
 (my-choosing-function "../other-dir/my-stuff")
 ((equal gnus-newsgroup-name "mail.misc") "mail-stuff"))
```

We see that this is a list where each element is a list that has two elements—the *match* and the *file*. The match can either be a string (in which case it is used as a regexp to match on the article head); it can be a symbol (which will be called as a function with the group name as a parameter); or it can be a list (which will be `eval`ed). If any of these actions have a non-`nil` result, the *file* will be used as a default prompt. In addition, the result of the operation itself will be used if the function or form called returns a string or a list of strings.

You basically end up with a list of file names that might be used when saving the current article. (All "matches" will be used.) You will then be prompted for what you really want to use as a name, with file name completion over the results from applying this variable.

This variable is `((gnus-article-archive-name))` by default, which means that Gnus will look at the articles it saves for an `Archive-name` line and use that as a suggestion for the file name.

Here's an example function to clean up file names somewhat. If you have lots of mail groups called things like '`nnml:mail.whatever`', you may want to chop off the beginning of these group names before creating the file name to save to. The following will do just that:

```
(defun my-save-name (group)
  (when (string-match "^nnml:mail." group)
    (substring group (match-end 0))))

(setq gnus-split-methods
      '((gnus-article-archive-name)
        (my-save-name)))
```

Finally, you have the `gnus-use-long-file-name` variable. If it is `nil`, all the preceding functions will replace all periods ('.') in the group names with slashes ('/')—which means that the functions will generate hierarchies of directories instead of having all the files in the top level directory ('`~/News/alt/andrea-dworkin`' instead of '`~/News/alt.andrea-dworkin`'.) This variable is `t` by default on most systems. However, for historical reasons, this is `nil` on Xenix and usg-unix-v machines by default.

This function also affects kill and score file names. If this variable is a list, and the list contains the element `not-score`, long file names will not be used for score files, if it contains the element `not-save`, long file names will not be used for saving, and if it contains the element `not-kill`, long file names will not be used for kill files.

If you'd like to save articles in a hierarchy that looks something like a spool, you could

```
(setq gnus-use-long-file-name '(not-save)) ; to get a hierarchy
(setq gnus-default-article-saver
      'gnus-summary-save-in-file) ; no encoding
```

Then just save with *o*. You'd then read this hierarchy with ephemeral `nneething` groups—*G D* in the group buffer, and the top level directory as the argument ('`~/News/`'). Then just walk around to the groups/directories with `nneething`.

3.15 Decoding Articles

Sometime users post articles (or series of articles) that have been encoded in some way or other. Gnus can decode them for you.

All these functions use the process/prefix convention (see Section 8.1 [Process/Prefix], page 189) for finding out what articles to work on, with the extension that a "single article" means "a single series". Gnus can find out by itself what articles belong to a series, decode all the articles and unpack/view/save the resulting file(s).

Gnus guesses what articles are in the series according to the following simplish rule: The subjects must be (nearly) identical, except for the last two numbers of the line. (Spaces are largely ignored, however.)

For example: If you choose a subject called '`cat.gif (2/3)`', Gnus will find all the articles that match the regexp '`^cat.gif ([0-9]+/[0-9]+).*$`'.

Subjects that are non-standard, like '`cat.gif (2/3) Part 6 of a series`', will not be properly recognized by any of the automatic viewing commands, and you have to mark the articles manually with `#`.

3.15.1 Uuencoded Articles

X u   Uudecodes the current series (`gnus-uu-decode-uu`).

X U   Uudecodes and saves the current series (`gnus-uu-decode-uu-and-save`).

X v u   Uudecodes and views the current series (`gnus-uu-decode-uu-view`).

X v U   Uudecodes, views and saves the current series (`gnus-uu-decode-uu-and-save-view`).

Remember that these all react to the presence of articles marked with the process mark. If, for instance, you'd like to decode and save an entire newsgroup, you'd typically do *M P a* (`gnus-uu-mark-all`) and then *X U* (`gnus-uu-decode-uu-and-save`).

All this is very much different from how `gnus-uu` worked with GNUS 4.1, where you had explicit keystrokes for everything under the sun. This version of `gnus-uu` generally assumes that you mark articles in some way (see Section 3.6.6 [Setting Process Marks], page 51) and then press *X u*.

Note: When trying to decode articles that have names matching `gnus-uu-notify-files`, which is hard-coded to '`[Cc][Ii][Nn][Dd][Yy][0-9]+.\\(gif\\|jpg\\)`', gnus-uu will automatically post an article on '`comp.unix.wizards`' saying that you have just viewed the file in question. This feature can't be turned off.

3.15.2 Shell Archives

Shell archives ("shar files") used to be a popular way to distribute sources, but it isn't used all that much today. In any case, we have some commands to deal with these:

X s   Unshars the current series (`gnus-uu-decode-unshar`).

X S   Unshars and saves the current series (`gnus-uu-decode-unshar-and-save`).

X v s   Unshars and views the current series (`gnus-uu-decode-unshar-view`).

X v S   Unshars, views and saves the current series (`gnus-uu-decode-unshar-and-save-view`).

3.15.3 PostScript Files

X p   Unpack the current PostScript series (`gnus-uu-decode-postscript`).

X P   Unpack and save the current PostScript series (`gnus-uu-decode-postscript-and-save`).

X v p   View the current PostScript series (`gnus-uu-decode-postscript-view`).

X v P   View and save the current PostScript series (`gnus-uu-decode-postscript-and-save-view`).

3.15.4 Other Files

X o   Save the current series (`gnus-uu-decode-save`).

X b   Unbinhex the current series (`gnus-uu-decode-binhex`). This doesn't really work yet.

3.15.5 Decoding Variables

Adjective, not verb.

3.15.5.1 Rule Variables

Gnus uses *rule variables* to decide how to view a file. All these variables are of the form

```
(list '(regexp1 command2)
      '(regexp2 command2)
      ...)
```

`gnus-uu-user-view-rules`

This variable is consulted first when viewing files. If you wish to use, for instance, `sox` to convert an '`.au`' sound file, you could say something like:

```
(setq gnus-uu-user-view-rules
      (list '("\\\\.au$" "sox %s -t .aiff > /dev/audio")))
```

`gnus-uu-user-view-rules-end`

This variable is consulted if Gnus couldn't make any matches from the user and default view rules.

`gnus-uu-user-archive-rules`

This variable can be used to say what commands should be used to unpack archives.

3.15.5.2 Other Decode Variables

`gnus-uu-grabbed-file-functions`

All functions in this list will be called right after each file has been successfully decoded—so that you can move or view files right away, and don't have to wait for all files to be decoded before you can do anything. Ready-made functions you can put in this list are:

`gnus-uu-grab-view`

View the file.

`gnus-uu-grab-move`

Move the file (if you're using a saving function.)

`gnus-uu-be-dangerous`

Specifies what to do if unusual situations arise during decoding. If `nil`, be as conservative as possible. If `t`, ignore things that didn't work, and overwrite existing files. Otherwise, ask each time.

`gnus-uu-ignore-files-by-name`

Files with name matching this regular expression won't be viewed.

`gnus-uu-ignore-files-by-type`

Files with a MIME type matching this variable won't be viewed. Note that Gnus tries to guess what type the file is based on the name. `gnus-uu` is not a MIME package (yet), so this is slightly kludgey.

`gnus-uu-tmp-dir`
> Where `gnus-uu` does its work.

`gnus-uu-do-not-unpack-archives`
> Non-`nil` means that `gnus-uu` won't peek inside archives looking for files to display.

`gnus-uu-view-and-save`
> Non-`nil` means that the user will always be asked to save a file after viewing it.

`gnus-uu-ignore-default-view-rules`
> Non-`nil` means that `gnus-uu` will ignore the default viewing rules.

`gnus-uu-ignore-default-archive-rules`
> Non-`nil` means that `gnus-uu` will ignore the default archive unpacking commands.

`gnus-uu-kill-carriage-return`
> Non-`nil` means that `gnus-uu` will strip all carriage returns from articles.

`gnus-uu-unmark-articles-not-decoded`
> Non-`nil` means that `gnus-uu` will mark unsuccessfully decoded articles as unread.

`gnus-uu-correct-stripped-uucode`
> Non-`nil` means that `gnus-uu` will *try* to fix uuencoded files that have had trailing spaces deleted.

`gnus-uu-pre-uudecode-hook`
> Hook run before sending a message to `uudecode`.

`gnus-uu-view-with-metamail`
> Non-`nil` means that `gnus-uu` will ignore the viewing commands defined by the rule variables and just fudge a MIME content type based on the file name. The result will be fed to `metamail` for viewing.

`gnus-uu-save-in-digest`
> Non-`nil` means that `gnus-uu`, when asked to save without decoding, will save in digests. If this variable is `nil`, `gnus-uu` will just save everything in a file without any embellishments. The digesting almost conforms to RFC 1153—no easy way to specify any meaningful volume and issue numbers were found, so I simply dropped them.

3.15.5.3 Uuencoding and Posting

`gnus-uu-post-include-before-composing`
> Non-`nil` means that `gnus-uu` will ask for a file to encode before you compose the article. If this variable is `t`, you can either include an encoded file with `C-c C-i` or have one included for you when you post the article.

`gnus-uu-post-length`
> Maximum length of an article. The encoded file will be split into how many articles it takes to post the entire file.

`gnus-uu-post-threaded`
>   Non-`nil` means that `gnus-uu` will post the encoded file in a thread. This may not be smart, as no other decoder I have seen is able to follow threads when collecting uuencoded articles. (Well, I have seen one package that does that— `gnus-uu`, but somehow, I don't think that counts...) Default is `nil`.

`gnus-uu-post-separate-description`
>   Non-`nil` means that the description will be posted in a separate article. The first article will typically be numbered (0/x). If this variable is `nil`, the description the user enters will be included at the beginning of the first article, which will be numbered (1/x). Default is `t`.

3.15.6 Viewing Files

After decoding, if the file is some sort of archive, Gnus will attempt to unpack the archive and see if any of the files in the archive can be viewed. For instance, if you have a gzipped tar file 'pics.tar.gz' containing the files 'pic1.jpg' and 'pic2.gif', Gnus will uncompress and de-tar the main file, and then view the two pictures. This unpacking process is recursive, so if the archive contains archives of archives, it'll all be unpacked.

Finally, Gnus will normally insert a *pseudo-article* for each extracted file into the summary buffer. If you go to these "articles", you will be prompted for a command to run (usually Gnus will make a suggestion), and then the command will be run.

If `gnus-view-pseudo-asynchronously` is `nil`, Emacs will wait until the viewing is done before proceeding.

If `gnus-view-pseudos` is `automatic`, Gnus will not insert the pseudo-articles into the summary buffer, but view them immediately. If this variable is `not-confirm`, the user won't even be asked for a confirmation before viewing is done.

If `gnus-view-pseudos-separately` is non-`nil`, one pseudo-article will be created for each file to be viewed. If `nil`, all files that use the same viewing command will be given as a list of parameters to that command.

If `gnus-insert-pseudo-articles` is non-`nil`, insert pseudo-articles when decoding. It is `t` by default.

So; there you are, reading your *pseudo-articles* in your *virtual newsgroup* from the *virtual server*; and you think: Why isn't anything real anymore? How did we get here?

3.16 Article Treatment

Reading through this huge manual, you may have quite forgotten that the object of newsreaders is to actually, like, read what people have written. Reading articles. Unfortunately, people are quite bad at writing, so there are tons of functions and variables to make reading these articles easier.

3.16.1 Article Highlighting

Not only do you want your article buffer to look like fruit salad, but you want it to look like technicolor fruit salad.

W H a   Do much highlighting of the current article (`gnus-article-highlight`). This
function highlights header, cited text, the signature, and adds buttons to the
body and the head.

W H h   Highlight the headers (`gnus-article-highlight-headers`). The highlighting
will be done according to the `gnus-header-face-alist` variable, which is a list
where each element has the form (*regexp name content*). *regexp* is a regular
expression for matching the header, *name* is the face used for highlighting the
header name (see Section 8.6 [Faces and Fonts], page 197) and *content* is the
face for highlighting the header value. The first match made will be used. Note
that *regexp* shouldn't have '`^`' prepended—Gnus will add one.

W H c   Highlight cited text (`gnus-article-highlight-citation`).

         Some variables to customize the citation highlights:

         `gnus-cite-parse-max-size`
                If the article size if bigger than this variable (which is 25000 by
                default), no citation highlighting will be performed.

         `gnus-cite-prefix-regexp`
                Regexp matching the longest possible citation prefix on a line.

         `gnus-cite-max-prefix`
                Maximum possible length for a citation prefix (default 20).

         `gnus-cite-face-list`
                List of faces used for highlighting citations (see Section 8.6 [Faces
                and Fonts], page 197). When there are citations from multiple
                articles in the same message, Gnus will try to give each citation
                from each article its own face. This should make it easier to see
                who wrote what.

         `gnus-supercite-regexp`
                Regexp matching normal Supercite attribution lines.

         `gnus-supercite-secondary-regexp`
                Regexp matching mangled Supercite attribution lines.

         `gnus-cite-minimum-match-count`
                Minimum number of identical prefixes we have to see before we
                believe that it's a citation.

         `gnus-cite-attribution-prefix`
                Regexp matching the beginning of an attribution line.

         `gnus-cite-attribution-suffix`
                Regexp matching the end of an attribution line.

         `gnus-cite-attribution-face`
                Face used for attribution lines. It is merged with the face for the
                cited text belonging to the attribution.

W H s   Highlight the signature (`gnus-article-highlight-signature`). Everything
after `gnus-signature-separator` (see Section 3.16.7 [Article Signature],

page 77) in an article will be considered a signature and will be highlighted with `gnus-signature-face`, which is `italic` by default.

See Section 4.3 [Customizing Articles], page 95, for how to highlight articles automatically.

3.16.2 Article Fontisizing

People commonly add emphasis to words in news articles by writing things like '`_this_`' or '`*this*`' or '`/this/`'. Gnus can make this look nicer by running the article through the *W e* (`gnus-article-emphasize`) command.

How the emphasis is computed is controlled by the `gnus-emphasis-alist` variable. This is an alist where the first element is a regular expression to be matched. The second is a number that says what regular expression grouping is used to find the entire emphasized word. The third is a number that says what regexp grouping should be displayed and highlighted. (The text between these two groupings will be hidden.) The fourth is the face used for highlighting.

```
(setq gnus-article-emphasis
      '(("_\\(\\w+\\)_" 0 1 gnus-emphasis-underline)
        ("\\*\\(\\w+\\)\\*" 0 1 gnus-emphasis-bold)))
```

By default, there are seven rules, and they use the following faces: `gnus-emphasis-bold`, `gnus-emphasis-italic`, `gnus-emphasis-underline`, `gnus-emphasis-bold-italic`, `gnus-emphasis-underline-italic`, `gnus-emphasis-underline-bold`, and `gnus-emphasis-underline-bold-italic`.

If you want to change these faces, you can either use *M-x customize*, or you can use `copy-face`. For instance, if you want to make `gnus-emphasis-italic` use a red face instead, you could say something like:

```
(copy-face 'red 'gnus-emphasis-italic)
```

If you want to highlight arbitrary words, you can use the `gnus-group-highlight-words-alist` variable, which uses the same syntax as `gnus-emphasis-alist`. The `highlight-words` group parameter (see Section 2.10 [Group Parameters], page 22) can also be used.

See Section 4.3 [Customizing Articles], page 95, for how to fontize articles automatically.

3.16.3 Article Hiding

Or rather, hiding certain things in each article. There usually is much too much cruft in most articles.

W W a   Do quite a lot of hiding on the article buffer (*gnus-article-hide*). In particular, this function will hide headers, PGP, cited text and the signature.

W W h   Hide headers (`gnus-article-hide-headers`). See Section 4.1 [Hiding Headers], page 93.

W W b   Hide headers that aren't particularly interesting (`gnus-article-hide-boring-headers`). See Section 4.1 [Hiding Headers], page 93.

W W s   Hide signature (`gnus-article-hide-signature`). See Section 3.16.7 [Article
           Signature], page 77.

W W l   Strip list identifiers specified in `gnus-list-identifiers`. These are strings
           some mailing list servers add to the beginning of all `Subject` headers—for
           example, '`[zebra 4711]`'. Any leading '`Re: `' is skipped before stripping. `gnus-`
           `list-identifiers` may not contain `\\(..\\)`.

           gnus-list-identifiers
                    A regular expression that matches list identifiers to be removed
                    from subject. This can also be a list of regular expressions.

W W p   Hide PGP signatures (`gnus-article-hide-pgp`). The `gnus-article-hide-`
           `pgp-hook` hook will be run after a PGP signature has been hidden. For example,
           to automatically verify articles that have signatures in them do:

                    ;;; Hide pgp cruft if any.

                    (setq gnus-treat-strip-pgp t)

                    ;;; After hiding pgp, verify the message;
                    ;;; only happens if pgp signature is found.

                    (add-hook 'gnus-article-hide-pgp-hook
                              (lambda ()
                                (save-excursion
                                  (set-buffer gnus-original-article-buffer)
                                  (mc-verify))))

W W P   Hide PEM (privacy enhanced messages) cruft (`gnus-article-hide-pem`).

W W B   Strip the banner specified by the `banner` group parameter (`gnus-article-`
           `strip-banner`). This is mainly used to hide those annoying banners and/or
           signatures that some mailing lists and moderated groups adds to all the mes-
           sages. The way to use this function is to add the `banner` group parameter
           (see Section 2.10 [Group Parameters], page 22) to the group you want banners
           stripped from. The parameter either be a string, which will be interpreted as
           a regular expression matching text to be removed, or the symbol `signature`,
           meaning that the (last) signature should be removed, or other symbol, meaning
           that the corresponding regular expression in `gnus-article-banner-alist` is
           used.

W W c   Hide citation (`gnus-article-hide-citation`). Some variables for customizing
           the hiding:

           gnus-cited-opened-text-button-line-format
           gnus-cited-closed-text-button-line-format
                    Gnus adds buttons to show where the cited text has been hidden,
                    and to allow toggle hiding the text. The format of the variable is
                    specified by these format-like variable (see Section 8.4 [Formatting
                    Variables], page 190). These specs are valid:

                    'b'   Starting point of the hidden text.

'e'   Ending point of the hidden text.

'l'   Number of characters in the hidden region.

'n'   Number of lines of hidden text.

gnus-cited-lines-visible
                  The number of lines at the beginning of the cited text to leave
                  shown. This can also be a cons cell with the number of lines at the
                  top and bottom of the text, respectively, to remain visible.

W W C-c

                  Hide citation (gnus-article-hide-citation-maybe) depending on the follow-
                  ing two variables:

                  gnus-cite-hide-percentage
                                    If the cited text is of a bigger percentage than this variable (default
                                    50), hide the cited text.

                  gnus-cite-hide-absolute
                                    The cited text must have at least this length (default 10) before it
                                    is hidden.

W W C   Hide cited text in articles that aren't roots (gnus-article-hide-citation-
                in-followups). This isn't very useful as an interactive command, but might
                be a handy function to stick have happen automatically (see Section 4.3 [Cus-
                tomizing Articles], page 95).

    All these "hiding" commands are toggles, but if you give a negative prefix to these
commands, they will show what they have previously hidden. If you give a positive prefix,
they will always hide.

    Also see Section 3.16.1 [Article Highlighting], page 69 for further variables for citation
customization.

    See Section 4.3 [Customizing Articles], page 95, for how to hide article elements auto-
matically.

3.16.4 Article Washing

    We call this "article washing" for a really good reason. Namely, the A key was taken, so
we had to use the W key instead.

    Washing is defined by us as "changing something from something to something else",
but normally results in something looking better. Cleaner, perhaps.

    See Section 4.3 [Customizing Articles], page 95, if you want to change how Gnus displays
articles by default.

C-u g   This is not really washing, it's sort of the opposite of washing. If you type this,
                you see the article exactly as it exists on disk or on the server.

W l   Remove   page   breaks   from   the   current   article   (gnus-summary-stop-
                page-breaking).   See Section 4.5 [Misc Article], page 97, for page
                delimiters.

W r   Do a Caesar rotate (rot13) on the article buffer (`gnus-summary-caesar-message`). Unreadable articles that tell you to read them with Caesar rotate or rot13. (Typically offensive jokes and such.)

It's commonly called *rot13* because each letter is rotated 13 positions in the alphabet, e. g. 'B' (letter #2) -> 'O' (letter #15). It is sometimes referred to as "Caesar rotate" because Caesar is rumored to have employed this form of, uh, somewhat weak encryption.

W t

t   Toggle whether to display all headers in the article buffer (`gnus-summary-toggle-header`).

W v   Toggle whether to display all headers in the article buffer permanently (`gnus-summary-verbose-header`).

W o   Treat overstrike (`gnus-article-treat-overstrike`).

W d   Treat M******** sm*rtq**t*s according to `gnus-article-dumbquotes-map` (`gnus-article-treat-dumbquotes`). Note that this function guesses whether a character is a sm*rtq**t* or not, so it should only be used interactively.

In reality, this function is translates a subset of the subset of the `cp1252` (or `Windows-1252`) character set that isn't in ISO Latin-1, including the quote characters `\222` and `\264`. Messages in this character set often have a MIME header saying that they are Latin-1.

W w   Do word wrap (`gnus-article-fill-cited-article`).

You can give the command a numerical prefix to specify the width to use when filling.

W Q   Fill long lines (`gnus-article-fill-long-lines`).

W C   Capitalize the first word in each sentence (`gnus-article-capitalize-sentences`).

W c   Translate CRLF pairs (i. e., '^M's on the end of the lines) into LF (this takes care of DOS line endings), and then translate any remaining CRs into LF (this takes care of Mac line endings) (`gnus-article-remove-cr`).

W q   Treat quoted-printable (`gnus-article-de-quoted-unreadable`). Quoted-Printable is one common MIME encoding employed when sending non-ASCII (i. e., 8-bit) articles. It typically makes strings like 'dj vu' look like 'd=E9j=E0 vu', which doesn't look very readable to me. Note that the this is usually done automatically by Gnus if the message in question has a `Content-Transfer-Encoding` header that says that this encoding has been done.

W 6   Treat base64 (`gnus-article-de-base64-unreadable`). Base64 is one common MIME encoding employed when sending non-ASCII (i. e., 8-bit) articles. Note that the this is usually done automatically by Gnus if the message in question has a `Content-Transfer-Encoding` header that says that this encoding has been done.

`W Z`   Treat HZ or HZP (`gnus-article-decode-HZ`). HZ (or HZP) is one common encoding employed when sending Chinese articles. It typically makes strings look like '`~{<:Ky2;S{#,NpJ)l6HK!#~}`'.

`W h`   Treat HTML (`gnus-article-wash-html`). Note that the this is usually done automatically by Gnus if the message in question has a `Content-Type` header that says that this type has been done.

`W f`   Look for and display any X-Face headers (`gnus-article-display-x-face`). The command executed by this function is given by the `gnus-article-x-face-command` variable. If this variable is a string, this string will be executed in a sub-shell. If it is a function, this function will be called with the face as the argument. If the `gnus-article-x-face-too-ugly` (which is a regexp) matches the `From` header, the face will not be shown. The default action under Emacs is to fork off the `display` program[1] to view the face. Under XEmacs or Emacs 21+ with suitable image support, the default action is to display the face before the `From` header. (It's nicer if XEmacs has been compiled with X-Face support— that will make display somewhat faster. If there's no native X-Face support, Gnus will try to convert the `X-Face` header using external programs from the `pbmplus` package and friends.[2]) If you want to have this function in the display hook, it should probably come last.

`W b`   Add clickable buttons to the article (`gnus-article-add-buttons`). See Section 3.16.5 [Article Buttons], page 76.

`W B`   Add clickable buttons to the article headers (`gnus-article-add-buttons-to-head`).

`W W H`   Strip headers like the `X-No-Archive` header from the beginning of article bodies (`gnus-article-strip-headers-from-body`).

`W E l`   Remove all blank lines from the beginning of the article (`gnus-article-strip-leading-blank-lines`).

`W E m`   Replace all blank lines with empty lines and then all multiple empty lines with a single empty line. (`gnus-article-strip-multiple-blank-lines`).

`W E t`   Remove all blank lines at the end of the article (`gnus-article-remove-trailing-blank-lines`).

`W E a`   Do all the three commands above (`gnus-article-strip-blank-lines`).

`W E A`   Remove all blank lines (`gnus-article-strip-all-blank-lines`).

`W E s`   Remove all white space from the beginning of all lines of the article body (`gnus-article-strip-leading-space`).

`W E e`   Remove all white space from the end of all lines of the article body (`gnus-article-strip-trailing-space`).

See Section 4.3 [Customizing Articles], page 95, for how to wash articles automatically.

---

[1]   `display` is from the ImageMagick package. For the `uncompface` and `icontopbm` programs look for a package like 'compface' or 'faces-xface' on a GNU/Linux system.

[2]   On a GNU/Linux system look for packages with names like `netpbm` or `libgr-progs`.

3.16.5 Article Buttons

People often include references to other stuff in articles, and it would be nice if Gnus could just fetch whatever it is that people talk about with the minimum of fuzz when you hit ⟨RET⟩ or use the middle mouse button on these references.

Gnus adds *buttons* to certain standard references by default: Well-formed URLs, mail addresses and Message-IDs. This is controlled by two variables, one that handles article bodies and one that handles article heads:

`gnus-button-alist`
>   This is an alist where each entry has this form:
>
>   >   `(REGEXP BUTTON-PAR USE-P FUNCTION DATA-PAR)`
>
>   *regexp* All text that match this regular expression will be considered an external reference. Here's a typical regexp that matches embedded URLs: '`<URL:\\([^\n\r>]*\\)>`'.
>
>   *button-par*
>   >   Gnus has to know which parts of the matches is to be highlighted. This is a number that says what sub-expression of the regexp is to be highlighted. If you want it all highlighted, you use 0 here.
>
>   *use-p*   This form will be `eval`ed, and if the result is non-`nil`, this is considered a match. This is useful if you want extra sifting to avoid false matches.
>
>   *function* This function will be called when you click on this button.
>
>   *data-par* As with *button-par*, this is a sub-expression number, but this one says which part of the match is to be sent as data to *function*.
>
>   So the full entry for buttonizing URLs is then
>
>   >   `("<URL:\\([^\n\r>]*\\)>" 0 t gnus-button-url 1)`

`gnus-header-button-alist`
>   This is just like the other alist, except that it is applied to the article head only, and that each entry has an additional element that is used to say what headers to apply the buttonize coding to:
>
>   >   `(HEADER REGEXP BUTTON-PAR USE-P FUNCTION DATA-PAR)`
>
>   *header* is a regular expression.

`gnus-button-url-regexp`
>   A regular expression that matches embedded URLs. It is used in the default values of the variables above.

`gnus-article-button-face`
>   Face used on buttons.

`gnus-article-mouse-face`
>   Face used when the mouse cursor is over a button.

See Section 4.3 [Customizing Articles], page 95, for how to buttonize articles automatically.

3.16.6 Article Date

The date is most likely generated in some obscure timezone you've never heard of, so it's quite nice to be able to find out what the time was when the article was sent.

W T u   Display the date in UT (aka. GMT, aka ZULU) (`gnus-article-date-ut`).

W T i   Display the date in international format, aka. ISO 8601 (`gnus-article-date-iso8601`).

W T l   Display the date in the local timezone (`gnus-article-date-local`).

W T s   Display the date using a user-defined format (`gnus-article-date-user`). The format is specified by the `gnus-article-time-format` variable, and is a string that's passed to `format-time-string`. See the documentation of that variable for a list of possible format specs.

W T e   Say how much time has elapsed between the article was posted and now (`gnus-article-date-lapsed`). It looks something like:

```
X-Sent: 6 weeks, 4 days, 1 hour, 3 minutes, 8 seconds ago
```

The value of `gnus-article-date-lapsed-new-header` determines whether this header will just be added below the old Date one, or will replace it.

An advantage of using Gnus to read mail is that it converts simple bugs into wonderful absurdities.

If you want to have this line updated continually, you can put

```
(gnus-start-date-timer)
```

in your '`.gnus.el`' file, or you can run it off of some hook. If you want to stop the timer, you can use the `gnus-stop-date-timer` command.

W T o   Display the original date (`gnus-article-date-original`). This can be useful if you normally use some other conversion function and are worried that it might be doing something totally wrong. Say, claiming that the article was posted in 1854. Although something like that is *totally* impossible. Don't you trust me? *titter*

See Section 4.3 [Customizing Articles], page 95, for how to display the date in your preferred format automatically.

3.16.7 Article Signature

Each article is divided into two parts—the head and the body. The body can be divided into a signature part and a text part. The variable that says what is to be considered a signature is `gnus-signature-separator`. This is normally the standard '`^-- $`' as mandated by son-of-RFC 1036. However, many people use non-standard signature separators, so this variable can also be a list of regular expressions to be tested, one by one. (Searches are done from the end of the body towards the beginning.) One likely value is:

```
(setq gnus-signature-separator
       '("^-- $"          ; The standard
         "^-- *$"         ; A common mangling
         "^-------*$"     ; Many people just use a looong
```

```
                                ; line of dashes.   Shame!
          "^ *--------*$" ; Double-shame!
          "^_____*$"   ; Underscores are also popular
          "^========*$")) ; Pervert!
```

The more permissive you are, the more likely it is that you'll get false positives.

`gnus-signature-limit` provides a limit to what is considered a signature when displaying articles.

1. If it is an integer, no signature may be longer (in characters) than that integer.
2. If it is a floating point number, no signature may be longer (in lines) than that number.
3. If it is a function, the function will be called without any parameters, and if it returns `nil`, there is no signature in the buffer.
4. If it is a string, it will be used as a regexp. If it matches, the text in question is not a signature.

This variable can also be a list where the elements may be of the types listed above. Here's an example:

```
(setq gnus-signature-limit
      '(200.0 "^---*Forwarded article"))
```

This means that if there are more than 200 lines after the signature separator, or the text after the signature separator is matched by the regular expression '`^---*Forwarded article`', then it isn't a signature after all.

3.16.8 Article Miscellania

`A t`   Translate the article from one language to another (`gnus-article-babel`).

3.17 MIME Commands

The following commands all understand the numerical prefix. For instance, `3 b` means "view the third MIME part".

`b`
`K v`   View the MIME part.

`K o`   Save the MIME part.

`K c`   Copy the MIME part.

`K e`   View the MIME part externally.

`K i`   View the MIME part internally.

`K |`   Pipe the MIME part to an external command.

The rest of these MIME commands do not use the numerical prefix in the same manner:

`K b`   Make all the MIME parts have buttons in front of them. This is mostly useful if you wish to save (or perform other actions) on inlined parts.

`K m`   Some multipart messages are transmitted with missing or faulty headers. This command will attempt to "repair" these messages so that they can be viewed in a more pleasant manner (`gnus-summary-repair-multipart`).

X m   Save all parts matching a MIME type to a directory (`gnus-summary-save-parts`). Understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

M-t   Toggle the buttonized display of the article buffer (`gnus-summary-toggle-display-buttonized`).

W M w   Decode RFC 2047-encoded words in the article headers (`gnus-article-decode-mime-words`).

W M c   Decode encoded article bodies as well as charsets (`gnus-article-decode-charset`).

              This command looks in the `Content-Type` header to determine the charset. If there is no such header in the article, you can give it a prefix, which will prompt for the charset to decode as. In regional groups where people post using some common encoding (but do not include MIME headers), you can set the `charset` group/topic parameter to the required charset (see Section 2.10 [Group Parameters], page 22).

W M v   View all the MIME parts in the current article (`gnus-mime-view-all-parts`).

   Relevant variables:

`gnus-ignored-mime-types`
              This is a list of regexps. MIME types that match a regexp from this list will be completely ignored by Gnus. The default value is `nil`.

              To have all Vcards be ignored, you'd say something like this:

```
(setq gnus-ignored-mime-types
      '("text/x-vcard"))
```

`gnus-unbuttonized-mime-types`
              This is a list of regexps. MIME types that match a regexp from this list won't have MIME buttons inserted unless they aren't displayed. The default value is `(".*/.*")`.

`gnus-article-mime-part-function`
              For each MIME part, this function will be called with the MIME handle as the parameter. The function is meant to be used to allow users to gather information from the article (e. g., add Vcard info to the bbdb database) or to do actions based on parts (e. g., automatically save all jpegs into some directory).

              Here's an example function the does the latter:

```
(defun my-save-all-jpeg-parts (handle)
  (when (equal (car (mm-handle-type handle)) "image/jpeg")
    (with-temp-buffer
      (insert (mm-get-part handle))
      (write-region (point-min) (point-max)
                    (read-file-name "Save jpeg to: ")))))
(setq gnus-article-mime-part-function
      'my-save-all-jpeg-parts)
```

`gnus-mime-multipart-functions`
              Alist of MIME multipart types and functions to handle them.

3.18 Charsets

People use different charsets, and we have MIME to let us know what charsets they use. Or rather, we wish we had. Many people use newsreaders and mailers that do not understand or use MIME, and just send out messages without saying what character sets they use. To help a bit with this, some local news hierarchies have policies that say what character set is the default. For instance, the 'fj' hierarchy uses `iso-2022-jp-2`.

This knowledge is encoded in the `gnus-group-charset-alist` variable, which is an alist of regexps (to match group names) and default charsets to be used when reading these groups.

In addition, some people do use soi-disant MIME-aware agents that aren't. These blithely mark messages as being in `iso-8859-1` even if they really are in `koi-8`. To help here, the `gnus-newsgroup-ignored-charsets` variable can be used. The charsets that are listed here will be ignored. The variable can be set on a group-by-group basis using the group parameters (see Section 2.10 [Group Parameters], page 22). The default value is `(unknown-8bit)`, which is something some agents insist on having in there.

When posting, `gnus-group-posting-charset-alist` is used to determine which charsets should not be encoded using the MIME encodings. For instance, some hierarchies discourage using quoted-printable header encoding.

This variable is an alist of regexps and permitted unencoded charsets for posting. Each element of the alist has the form (*test header body-list*), where:

test   is either a regular expression matching the newsgroup header or a variable to query,

header   is the charset which may be left unencoded in the header (`nil` means encode all charsets),

body-list   is a list of charsets which may be encoded using 8bit content-transfer encoding in the body, or one of the special values `nil` (always encode using quoted-printable) or `t` (always use 8bit).

Other charset tricks that may be useful, although not Gnus-specific:

If there are several MIME charsets that encode the same Emacs charset, you can choose what charset to use by saying the following:

```
(put-charset-property 'cyrillic-iso8859-5
                      'preferred-coding-system 'koi8-r)
```

This means that Russian will be encoded using `koi8-r` instead of the default `iso-8859-5` MIME charset.

If you want to read messages in `koi8-u`, you can cheat and say

```
(define-coding-system-alias 'koi8-u 'koi8-r)
```

This will almost do the right thing.

And finally, to read charsets like `windows-1251`, you can say something like

```
(codepage-setup 1251)
(define-coding-system-alias 'windows-1251 'cp1251)
```

while if you use a non-Latin-1 language environment you could see the Latin-1 subset of `windows-1252` using:

```
(define-coding-system-alias 'windows-1252 'latin-1)
```

3.19 Article Commands

A P   Generate and print a PostScript image of the article buffer (`gnus-summary-print-article`). `gnus-ps-print-hook` will be run just before printing the buffer.

3.20 Summary Sorting

You can have the summary buffer sorted in various ways, even though I can't really see why you'd want that.

C-c C-s C-n

        Sort by article number (`gnus-summary-sort-by-number`).

C-c C-s C-a

        Sort by author (`gnus-summary-sort-by-author`).

C-c C-s C-s

        Sort by subject (`gnus-summary-sort-by-subject`).

C-c C-s C-d

        Sort by date (`gnus-summary-sort-by-date`).

C-c C-s C-l

        Sort by lines (`gnus-summary-sort-by-lines`).

C-c C-s C-c

        Sort by article length (`gnus-summary-sort-by-chars`).

C-c C-s C-i

        Sort by score (`gnus-summary-sort-by-score`).

These functions will work both when you use threading and when you don't use threading. In the latter case, all summary lines will be sorted, line by line. In the former case, sorting will be done on a root-by-root basis, which might not be what you were looking for. To toggle whether to use threading, type *T T* (see Section 3.8.2 [Thread Commands], page 57).

3.21 Finding the Parent

^   If you'd like to read the parent of the current article, and it is not displayed in the summary buffer, you might still be able to. That is, if the current group is fetched by NNTP, the parent hasn't expired and the `References` in the current article are not mangled, you can just press *^* or *A r* (`gnus-summary-refer-parent-article`). If everything goes well, you'll get the parent. If the parent is already displayed in the summary buffer, point will just move to this article.

        If given a positive numerical prefix, fetch that many articles back into the ancestry. If given a negative numerical prefix, fetch just that ancestor. So if you say *3 ^*, Gnus will fetch the parent, the grandparent and the grandgrandparent of the current article. If you say *-3 ^*, Gnus will only fetch the grandgrandparent of the current article.

`A R` (Summary)

> Fetch all articles mentioned in the `References` header of the article (`gnus-summary-refer-references`).

`A T` (Summary)

> Display the full thread where the current article appears (`gnus-summary-refer-thread`). This command has to fetch all the headers in the current group to work, so it usually takes a while. If you do it often, you may consider setting `gnus-fetch-old-headers` to `invisible` (see Section 3.8.1.2 [Filling In Threads], page 56). This won't have any visible effects normally, but it'll make this command work a whole lot faster. Of course, it'll make group entry somewhat slow.
>
> The `gnus-refer-thread-limit` variable says how many old (i. e., articles before the first displayed in the current group) headers to fetch when doing this command. The default is 200. If `t`, all the available headers will be fetched. This variable can be overridden by giving the `A T` command a numerical prefix.

`M-^` (Summary)

> You can also ask the NNTP server for an arbitrary article, no matter what group it belongs to. `M-^` (`gnus-summary-refer-article`) will ask you for a `Message-ID`, which is one of those long, hard-to-read thingies that look something like '`<38o6up$6f2@hymir.ifi.uio.no>`'. You have to get it all exactly right. No fuzzy searches, I'm afraid.

The current select method will be used when fetching by `Message-ID` from non-news select method, but you can override this by giving this command a prefix.

If the group you are reading is located on a back end that does not support fetching by `Message-ID` very well (like `nnspool`), you can set `gnus-refer-article-method` to an NNTP method. It would, perhaps, be best if the NNTP server you consult is the one updating the spool you are reading from, but that's not really necessary.

It can also be a list of select methods, as well as the special symbol `current`, which means to use the current select method. If it is a list, Gnus will try all the methods in the list until it finds a match.

Here's an example setting that will first try the current method, and then ask Deja if that fails:

```
(setq gnus-refer-article-method
      '(current
        (nnweb "refer" (nnweb-type dejanews))))
```

Most of the mail back ends support fetching by `Message-ID`, but do not do a particularly excellent job at it. That is, `nnmbox` and `nnbabyl` are able to locate articles from any groups, while `nnml` and `nnfolder` are only able to locate articles that have been posted to the current group. (Anything else would be too time consuming.) `nnmh` does not support this at all.

3.22 Alternative Approaches

Different people like to read news using different methods. This being Gnus, we offer a small selection of minor modes for the summary buffers.

3.22.1 Pick and Read

Some newsreaders (like `nn` and, uhm, `Netnews` on VM/CMS) use a two-phased reading interface. The user first marks in a summary buffer the articles she wants to read. Then she starts reading the articles with just an article buffer displayed.

Gnus provides a summary buffer minor mode that allows this—`gnus-pick-mode`. This basically means that a few process mark commands become one-keystroke commands to allow easy marking, and it provides one additional command for switching to the summary buffer.

Here are the available keystrokes when using pick mode:

.   Pick the article or thread on the current line (`gnus-pick-article-or-thread`). If the variable `gnus-thread-hide-subtree` is true, then this key selects the entire thread when used at the first article of the thread. Otherwise, it selects just the article. If given a numerical prefix, go to that thread or article and pick it. (The line number is normally displayed at the beginning of the summary pick lines.)

⟨SPC⟩   Scroll the summary buffer up one page (`gnus-pick-next-page`). If at the end of the buffer, start reading the picked articles.

u   Unpick the thread or article (`gnus-pick-unmark-article-or-thread`). If the variable `gnus-thread-hide-subtree` is true, then this key unpicks the thread if used at the first article of the thread. Otherwise it unpicks just the article. You can give this key a numerical prefix to unpick the thread or article at that line.

⟨RET⟩   Start reading the picked articles (`gnus-pick-start-reading`). If given a prefix, mark all unpicked articles as read first. If `gnus-pick-display-summary` is non-`nil`, the summary buffer will still be visible when you are reading.

All the normal summary mode commands are still available in the pick-mode, with the exception of `u`. However `!` is available which is mapped to the same function `gnus-summary-tick-article-forward`.

If this sounds like a good idea to you, you could say:

```
(add-hook 'gnus-summary-mode-hook 'gnus-pick-mode)
```

`gnus-pick-mode-hook` is run in pick minor mode buffers.

If `gnus-mark-unpicked-articles-as-read` is non-`nil`, mark all unpicked articles as read. The default is `nil`.

The summary line format in pick mode is slightly different from the standard format. At the beginning of each line the line number is displayed. The pick mode line format is controlled by the `gnus-summary-pick-line-format` variable (see Section 8.4 [Formatting Variables], page 190). It accepts the same format specs that `gnus-summary-line-format` does (see Section 3.1.1 [Summary Buffer Lines], page 37).

3.22.2 Binary Groups

If you spend much time in binary groups, you may grow tired of hitting `X u`, `n`, ⟨RET⟩ all the time. `M-x gnus-binary-mode` is a minor mode for summary buffers that makes all

ordinary Gnus article selection functions uudecode series of articles and display the result instead of just displaying the articles the normal way.

The only way, in fact, to see the actual articles is the *g* command, when you have turned on this mode (`gnus-binary-show-article`).

`gnus-binary-mode-hook` is called in binary minor mode buffers.

3.23 Tree Display

If you don't like the normal Gnus summary display, you might try setting `gnus-use-trees` to `t`. This will create (by default) an additional *tree buffer*. You can execute all summary mode commands in the tree buffer.

There are a few variables to customize the tree display, of course:

`gnus-tree-mode-hook`
>   A hook called in all tree mode buffers.

`gnus-tree-mode-line-format`
>   A format string for the mode bar in the tree mode buffers (see Section 8.4.2
>   [Mode Line Formatting], page 191). The default is 'Gnus: %%b %S %Z'. For a
>   list of valid specs, see Section 3.1.3 [Summary Buffer Mode Line], page 40.

`gnus-selected-tree-face`
>   Face used for highlighting the selected article in the tree buffer. The default is
>   `modeline`.

`gnus-tree-line-format`
>   A format string for the tree nodes. The name is a bit of a misnomer, though—it
>   doesn't define a line, but just the node. The default value is '%(%[%3,3n%]%)',
>   which displays the first three characters of the name of the poster. It is vital
>   that all nodes are of the same length, so you *must* use '%4,4n'-like specifiers.
>
>   Valid specs are:
>
>   'n'   The name of the poster.
>
>   'f'   The `From` header.
>
>   'N'   The number of the article.
>
>   '['   The opening bracket.
>
>   ']'   The closing bracket.
>
>   's'   The subject.
>
>   See Section 8.4 [Formatting Variables], page 190.
>
>   Variables related to the display are:
>
>   `gnus-tree-brackets`
>   >   This is used for differentiating between "real" articles and "sparse"
>   >   articles. The format is ((*real-open* . *real-close*) (*sparse-open* .
>   >   *sparse-close*) (*dummy-open* . *dummy-close*)), and the default is
>   >   ((?[. ?]) (?(. ?)) (?{ . ?}) (?< . ?>)).

gnus-tree-parent-child-edges

> This is a list that contains the characters used for connecting parent nodes to their children. The default is (?- ?\\ ?|).

gnus-tree-minimize-window

> If this variable is non-nil, Gnus will try to keep the tree buffer as small as possible to allow more room for the other Gnus windows. If this variable is a number, the tree buffer will never be higher than that number. The default is t. Note that if you have several windows displayed side-by-side in a frame and the tree buffer is one of these, minimizing the tree window will also resize all other windows displayed next to it.

gnus-generate-tree-function

> The function that actually generates the thread tree. Two predefined functions are available: gnus-generate-horizontal-tree and gnus-generate-vertical-tree (which is the default).

Here's an example from a horizontal tree buffer:

```
{***}-(***)-[odd]-[Gun]
    |        \[Jan]
    |        \[odd]-[Eri]
    |        \(***)-[Eri]
    |              \[odd]-[Paa]
    \[Bjo]
    \[Gun]
    \[Gun]-[Jor]
```

Here's the same thread displayed in a vertical tree buffer:

```
{***}
  |------------------------\-----\-----\
(***)                            [Bjo] [Gun] [Gun]
  |--\-----\-----\                            |
[odd] [Jan] [odd] (***)                     [Jor]
  |          |      |--\
[Gun]      [Eri] [Eri] [odd]
                          |
                        [Paa]
```

If you're using horizontal trees, it might be nice to display the trees side-by-side with the summary buffer. You could add something like the following to your '.gnus.el' file:

```
(setq gnus-use-trees t
      gnus-generate-tree-function 'gnus-generate-horizontal-tree
      gnus-tree-minimize-window nil)
(gnus-add-configuration
 '(article
   (vertical 1.0
             (horizontal 0.25
                         (summary 0.75 point)
                         (tree 1.0))
             (article 1.0))))
```

See Section 8.5 [Windows Configuration], page 193.

3.24 Mail Group Commands

Some commands only make sense in mail groups. If these commands are invalid in the current group, they will raise a hell and let you know.

All these commands (except the expiry and edit commands) use the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

B e　　　　Expire all expirable articles in the group (`gnus-summary-expire-articles`).

B M-C-e　　Delete all the expirable articles in the group (`gnus-summary-expire-articles-now`). This means that **all** articles eligible for expiry in the current group will disappear forever into that big '`/dev/null`' in the sky.

B (DEL)　　Delete the mail article. This is "delete" as in "delete it from your disk forever and ever, never to return again." Use with caution. (`gnus-summary-delete-article`).

B m　　　　Move the article from one mail group to another (`gnus-summary-move-article`). Marks will be preserved if *gnus-preserve-marks* is non-`nil` (which is the default).

B c　　　　Copy the article from one group (mail group or not) to a mail group (`gnus-summary-copy-article`). Marks will be preserved if *gnus-preserve-marks* is non-`nil` (which is the default).

B B　　　　Crosspost the current article to some other group (`gnus-summary-crosspost-article`). This will create a new copy of the article in the other group, and the Xref headers of the article will be properly updated.

B i　　　　Import an arbitrary file into the current mail newsgroup (`gnus-summary-import-article`). You will be prompted for a file name, a `From` header and a `Subject` header.

B r　　　　Respool the mail article (`gnus-summary-respool-article`). `gnus-summary-respool-default-method` will be used as the default select method when respooling. This variable is `nil` by default, which means that the current group select method will be used instead. Marks will be preserved if *gnus-preserve-marks* is non-`nil` (which is the default).

B w
e　　　　　Edit the current article (`gnus-summary-edit-article`). To finish editing and make the changes permanent, type *C-c C-c* (*gnus-summary-edit-article-done*). If you give a prefix to the *C-c C-c* command, Gnus won't re-highlight the article.

B q　　　　If you want to re-spool an article, you might be curious as to what group the article will end up in before you do the re-spooling. This command will tell you (`gnus-summary-respool-query`).

B t　　　　Similarly, this command will display all fancy splitting patterns used when repooling, if any (`gnus-summary-respool-trace`).

B p　　　　Some people have a tendency to send you "courtesy" copies when they follow up to articles you have posted. These usually have a `Newsgroups` header in

them, but not always. This command (`gnus-summary-article-posted-p`) will try to fetch the current article from your news server (or rather, from `gnus-refer-article-method` or `gnus-select-method`) and will report back whether it found the article or not. Even if it says that it didn't find the article, it may have been posted anyway—mail propagation is much faster than news propagation, and the news copy may just not have arrived yet.

If you move (or copy) articles regularly, you might wish to have Gnus suggest where to put the articles. `gnus-move-split-methods` is a variable that uses the same syntax as `gnus-split-methods` (see Section 3.14 [Saving Articles], page 62). You may customize that variable to create suggestions you find reasonable. (Note that `gnus-move-split-methods` uses group names where `gnus-split-methods` uses file names.)

```
(setq gnus-move-split-methods
      '(("^From:.*Lars Magne" "nnml:junk")
        ("^Subject:.*gnus" "nnfolder:important")
        (".*" "nnml:misc")))
```

3.25 Various Summary Stuff

`gnus-summary-mode-hook`
> This hook is called when creating a summary mode buffer.

`gnus-summary-generate-hook`
> This is called as the last thing before doing the threading and the generation of the summary buffer. It's quite convenient for customizing the threading variables based on what data the newsgroup has. This hook is called from the summary buffer after most summary buffer variables have been set.

`gnus-summary-prepare-hook`
> It is called after the summary buffer has been generated. You might use it to, for instance, highlight lines or modify the look of the buffer in some other ungodly manner. I don't care.

`gnus-summary-prepared-hook`
> A hook called as the very last thing after the summary buffer has been generated.

`gnus-summary-ignore-duplicates`
> When Gnus discovers two articles that have the same `Message-ID`, it has to do something drastic. No articles are allowed to have the same `Message-ID`, but this may happen when reading mail from some sources. Gnus allows you to customize what happens with this variable. If it is `nil` (which is the default), Gnus will rename the `Message-ID` (for display purposes only) and display the article as any other article. If this variable is `t`, it won't display the article—it'll be as if it never existed.

`gnus-alter-articles-to-read-function`
> This function, which takes two parameters (the group name and the list of articles to be selected), is called to allow the user to alter the list of articles to be selected.

For instance, the following function adds the list of cached articles to the list in one particular group:

```
(defun my-add-cached-articles (group articles)
  (if (string= group "some.group")
      (append gnus-newsgroup-cached articles)
    articles))
```

3.25.1 Summary Group Information

H f   Try to fetch the FAQ (list of frequently asked questions) for the current group (`gnus-summary-fetch-faq`). Gnus will try to get the FAQ from `gnus-group-faq-directory`, which is usually a directory on a remote machine. This variable can also be a list of directories. In that case, giving a prefix to this command will allow you to choose between the various sites. `ange-ftp` or `efs` will probably be used for fetching the file.

H d   Give a brief description of the current group (`gnus-summary-describe-group`). If given a prefix, force rereading the description from the server.

H h   Give an extremely brief description of the most important summary keystrokes (`gnus-summary-describe-briefly`).

H i   Go to the Gnus info node (`gnus-info-find-node`).

3.25.2 Searching for Articles

M-s   Search through all subsequent articles for a regexp (`gnus-summary-search-article-forward`).

M-r   Search through all previous articles for a regexp (`gnus-summary-search-article-backward`).

&   This command will prompt you for a header, a regular expression to match on this field, and a command to be executed if the match is made (`gnus-summary-execute-command`). If the header is an empty string, the match is done on the entire article. If given a prefix, search backward instead.

       For instance, *&* RET *some.*string* *#* will put the process mark on all articles that have heads or bodies that match '`some.*string`'.

M-&   Perform any operation on all articles that have been marked with the process mark (`gnus-summary-universal-argument`).

3.25.3 Summary Generation Commands

Y g   Regenerate the current summary buffer (`gnus-summary-prepare`).

Y c   Pull all cached articles (for the current group) into the summary buffer (`gnus-summary-insert-cached-articles`).

3.25.4 Really Various Summary Commands

`A D`
`C-d`   If the current article is a collection of other articles (for instance, a digest),
             you might use this command to enter a group based on the that article (`gnus-`
             `summary-enter-digest-group`). Gnus will try to guess what article type is
             currently displayed unless you give a prefix to this command, which forces
             a "digest" interpretation. Basically, whenever you see a message that is a
             collection of other messages of some format, you `C-d` and read these messages
             in a more convenient fashion.

`M-C-d`   This command is very similar to the one above, but lets you gather several
             documents into one biiig group (`gnus-summary-read-document`). It does this
             by opening several `nndoc` groups for each document, and then opening an
             `nnvirtual` group on top of these `nndoc` groups. This command understands
             the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

`C-t`   Toggle truncation of summary lines (`gnus-summary-toggle-truncation`).
             This will probably confuse the line centering function in the summary buffer,
             so it's not a good idea to have truncation switched off while reading articles.

`=`   Expand the summary buffer window (`gnus-summary-expand-window`). If given
             a prefix, force an `article` window configuration.

`M-C-e`   Edit the group parameters (see Section 2.10 [Group Parameters], page 22) of
             the current group (`gnus-summary-edit-parameters`).

`M-C-a`   Customize the group parameters (see Section 2.10 [Group Parameters], page 22)
             of the current group (`gnus-summary-customize-parameters`).

3.26 Exiting the Summary Buffer

   Exiting from the summary buffer will normally update all info on the group and return
you to the group buffer.

`Z Z`
`q`   Exit the current group and update all information on the group (`gnus-`
             `summary-exit`). `gnus-summary-prepare-exit-hook` is called before doing
             much of the exiting, which calls `gnus-summary-expire-articles` by
             default. `gnus-summary-exit-hook` is called after finishing the exit process.
             `gnus-group-no-more-groups-hook` is run when returning to group mode
             having no more (unread) groups.

`Z E`
`Q`   Exit the current group without updating any information on the group (`gnus-`
             `summary-exit-no-update`).

`Z c`
`c`   Mark all unticked articles in the group as read and then exit (`gnus-summary-`
             `catchup-and-exit`).

Z C   Mark all articles, even the ticked ones, as read and then exit (`gnus-summary-catchup-all-and-exit`).

Z n   Mark all articles as read and go to the next group (`gnus-summary-catchup-and-goto-next-group`).

Z R   Exit this group, and then enter it again (`gnus-summary-reselect-current-group`). If given a prefix, select all articles, both read and unread.

Z G

M-g   Exit the group, check for new articles in the group, and select the group (`gnus-summary-rescan-group`). If given a prefix, select all articles, both read and unread.

Z N   Exit the group and go to the next group (`gnus-summary-next-group`).

Z P   Exit the group and go to the previous group (`gnus-summary-prev-group`).

Z s   Save the current number of read/marked articles in the dribble buffer and then save the dribble buffer (`gnus-summary-save-newsrc`). If given a prefix, also save the '`.newsrc`' file(s). Using this command will make exit without updating (the `Q` command) worthless.

`gnus-exit-group-hook` is called when you exit the current group with an "updating" exit. For instance `Q` (`gnus-summary-exit-no-update`) does not call this hook.

If you're in the habit of exiting groups, and then changing your mind about it, you might set `gnus-kill-summary-on-exit` to `nil`. If you do that, Gnus won't kill the summary buffer when you exit it. (Quelle surprise!) Instead it will change the name of the buffer to something like '`*Dead Summary ... *`' and install a minor mode called `gnus-dead-summary-mode`. Now, if you switch back to this buffer, you'll find that all keys are mapped to a function called `gnus-summary-wake-up-the-dead`. So tapping any keys in a dead summary buffer will result in a live, normal summary buffer.

There will never be more than one dead summary buffer at any one time.

The data on the current group will be updated (which articles you have read, which articles you have replied to, etc.) when you exit the summary buffer. If the `gnus-use-cross-reference` variable is `t` (which is the default), articles that are cross-referenced to this group and are marked as read, will also be marked as read in the other subscribed groups they were cross-posted to. If this variable is neither `nil` nor `t`, the article will be marked as read in both subscribed and unsubscribed groups (see Section 3.27 [Crosspost Handling], page 90).

3.27 Crosspost Handling

Marking cross-posted articles as read ensures that you'll never have to read the same article more than once. Unless, of course, somebody has posted it to several groups separately. Posting the same article to several groups (not cross-posting) is called *spamming*, and you are by law required to send nasty-grams to anyone who perpetrates such a heinous crime. You may want to try NoCeM handling to filter out spam (see Section 8.12 [NoCeM], page 200).

Remember: Cross-posting is kinda ok, but posting the same article separately to several groups is not. Massive cross-posting (aka. *velveeta*) is to be avoided at all costs, and you can even use the `gnus-summary-mail-crosspost-complaint` command to complain about excessive crossposting (see Section 3.5.1 [Summary Mail Commands], page 44).

One thing that may cause Gnus to not do the cross-posting thing correctly is if you use an NNTP server that supports XOVER (which is very nice, because it speeds things up considerably) which does not include the `Xref` header in its NOV lines. This is Evil, but all too common, alas, alack. Gnus tries to Do The Right Thing even with XOVER by registering the `Xref` lines of all articles you actually read, but if you kill the articles, or just mark them as read without reading them, Gnus will not get a chance to snoop the `Xref` lines out of these articles, and will be unable to use the cross reference mechanism.

To check whether your NNTP server includes the `Xref` header in its overview files, try '`telnet your.nntp.server nntp`', '`MODE READER`' on `inn` servers, and then say '`LIST overview.fmt`'. This may not work, but if it does, and the last line you get does not read '`Xref:full`', then you should shout and whine at your news admin until she includes the `Xref` header in the overview files.

If you want Gnus to get the `Xref`s right all the time, you have to set `gnus-nov-is-evil` to `t`, which slows things down considerably.

C'est la vie.

For an alternative approach, see Section 3.28 [Duplicate Suppression], page 91.

3.28 Duplicate Suppression

By default, Gnus tries to make sure that you don't have to read the same article more than once by utilizing the crossposting mechanism (see Section 3.27 [Crosspost Handling], page 90). However, that simple and efficient approach may not work satisfactory for some users for various reasons.

1. The NNTP server may fail to generate the `Xref` header. This is evil and not very common.

2. The NNTP server may fail to include the `Xref` header in the '`.overview`' data bases. This is evil and all too common, alas.

3. You may be reading the same group (or several related groups) from different NNTP servers.

4. You may be getting mail that duplicates articles posted to groups.

I'm sure there are other situations where `Xref` handling fails as well, but these four are the most common situations.

If, and only if, `Xref` handling fails for you, then you may consider switching on *duplicate suppression*. If you do so, Gnus will remember the `Message-ID`s of all articles you have read or otherwise marked as read, and then, as if by magic, mark them as read all subsequent times you see them—in *all* groups. Using this mechanism is quite likely to be somewhat inefficient, but not overly so. It's certainly preferable to reading the same articles more than once.

Duplicate suppression is not a very subtle instrument. It's more like a sledge hammer than anything else. It works in a very simple fashion—if you have marked an article as

read, it adds this Message-ID to a cache. The next time it sees this Message-ID, it will
mark the article as read with the 'M' mark. It doesn't care what group it saw the article in.

gnus-suppress-duplicates
>   If non-`nil`, suppress duplicates.

gnus-save-duplicate-list
>   If non-`nil`, save the list of duplicates to a file. This will make startup and
>   shutdown take longer, so the default is `nil`. However, this means that only
>   duplicate articles read in a single Gnus session are suppressed.

gnus-duplicate-list-length
>   This variable says how many `Message-ID`s to keep in the duplicate suppression
>   list. The default is 10000.

gnus-duplicate-file
>   The name of the file to store the duplicate suppression list in. The default is
>   '`~/News/suppression`'.

If you have a tendency to stop and start Gnus often, setting `gnus-save-duplicate-list`
to `t` is probably a good idea. If you leave Gnus running for weeks on end, you may have it
`nil`. On the other hand, saving the list makes startup and shutdown much slower, so that
means that if you stop and start Gnus often, you should set `gnus-save-duplicate-list`
to `nil`. Uhm. I'll leave this up to you to figure out, I think.

4 The Article Buffer

The articles are displayed in the article buffer, of which there is only one. All the summary buffers share the same article buffer unless you tell Gnus otherwise.

4.1 Hiding Headers

The top section of each article is the *head*. (The rest is the *body*, but you may have guessed that already.)

There is a lot of useful information in the head: the name of the person who wrote the article, the date it was written and the subject of the article. That's well and nice, but there's also lots of information most people do not want to see—what systems the article has passed through before reaching you, the `Message-ID`, the `References`, etc. ad nauseum—and you'll probably want to get rid of some of those lines. If you want to keep all those lines in the article buffer, you can set `gnus-show-all-headers` to `t`.

Gnus provides you with two variables for sifting headers:

`gnus-visible-headers`

        If this variable is non-`nil`, it should be a regular expression that says what headers you wish to keep in the article buffer. All headers that do not match this variable will be hidden.

        For instance, if you only want to see the name of the person who wrote the article and the subject, you'd say:

            `(setq gnus-visible-headers "^From:\\|^Subject:")`

        This variable can also be a list of regexps to match headers to remain visible.

`gnus-ignored-headers`

        This variable is the reverse of `gnus-visible-headers`. If this variable is set (and `gnus-visible-headers` is `nil`), it should be a regular expression that matches all lines that you want to hide. All lines that do not match this variable will remain visible.

        For instance, if you just want to get rid of the `References` line and the `Xref` line, you might say:

            `(setq gnus-ignored-headers "^References:\\|^Xref:")`

        This variable can also be a list of regexps to match headers to be removed.

        Note that if `gnus-visible-headers` is non-`nil`, this variable will have no effect.

Gnus can also sort the headers for you. (It does this by default.) You can control the sorting by setting the `gnus-sorted-header-list` variable. It is a list of regular expressions that says in what order the headers are to be displayed.

For instance, if you want the name of the author of the article first, and then the subject, you might say something like:

    `(setq gnus-sorted-header-list '("^From:" "^Subject:"))`

Any headers that are to remain visible, but are not listed in this variable, will be displayed in random order after all the headers listed in this variable.

You can hide further boring headers by setting `gnus-treat-hide-boring-header` to `head`. What this function does depends on the `gnus-boring-article-headers` variable. It's a list, but this list doesn't actually contain header names. Instead is lists various *boring conditions* that Gnus can check and remove from sight.

These conditions are:

empty   Remove all empty headers.

followup-to
> Remove the `Followup-To` header if it is identical to the `Newsgroups` header.

reply-to   Remove the `Reply-To` header if it lists the same address as the `From` header.

newsgroups
> Remove the `Newsgroups` header if it only contains the current group name.

date   Remove the `Date` header if the article is less than three days old.

long-to   Remove the `To` header if it is very long.

many-to   Remove all `To` headers if there are more than one.

To include the four three elements, you could say something like;

```
(setq gnus-boring-article-headers
      '(empty followup-to reply-to))
```

This is also the default value for this variable.

4.2 Using MIME

Mime is a standard for waving your hands through the air, aimlessly, while people stand around yawning.

MIME, however, is a standard for encoding your articles, aimlessly, while all newsreaders die of fear.

MIME may specify what character set the article uses, the encoding of the characters, and it also makes it possible to embed pictures and other naughty stuff in innocent-looking articles.

Gnus pushes MIME articles through `gnus-display-mime-function` to display the MIME parts. This is `gnus-display-mime` by default, which creates a bundle of clickable buttons that can be used to display, save and manipulate the MIME objects.

The following commands are available when you have placed point over a MIME button:

⟨RET⟩ (Article)
Mouse-2 (Article)
> Toggle displaying of the MIME object (`gnus-article-press-button`).

M-⟨RET⟩ (Article)
v (Article) Prompt for a method, and then view the MIME object using this method (`gnus-mime-view-part`).

o (Article) Prompt for a file name, and then save the MIME object (`gnus-mime-save-part`).

c (Article)   Copy the MIME object to a fresh buffer and display this buffer (`gnus-mime-copy-part`).

t (Article)   View the MIME object as if it were a different MIME media type (`gnus-mime-view-part-as-type`).

| (Article)   Output the MIME object to a process (`gnus-mime-pipe-part`).

i (Article)   Insert the contents of the MIME object into the buffer (`gnus-mime-inline-part`) as text/plain. If given a prefix, insert the raw contents without decoding. If given a numerical prefix, you can do semi-manual charset stuff (see `gnus-summary-show-article-charset-alist` in see Section 3.4 [Paging the Article], page 43).

. (Article)   Interactively run an action on the MIME object (`gnus-mime-action-on-part`).

Gnus will display some MIME objects automatically. The way Gnus determines which parts to do this with is described in the Emacs MIME manual.

It might be best to just use the toggling functions from the article buffer to avoid getting nasty surprises. (For instance, you enter the group 'alt.sing-a-long' and, before you know it, MIME has decoded the sound file in the article and some horrible sing-a-long song comes screaming out your speakers, and you can't find the volume button, because there isn't one, and people are starting to look at you, and you try to stop the program, but you can't, and you can't find the program to control the volume, and everybody else in the room suddenly decides to look at you disdainfully, and you'll feel rather stupid.)

Any similarity to real events and people is purely coincidental. Ahem.

Also see see Section 3.17 [MIME Commands], page 78.

4.3 Customizing Articles

A slew of functions for customizing how the articles are to look like exist. You can call these functions interactively, or you can have them called automatically when you select the articles.

To have them called automatically, you should set the corresponding "treatment" variable. For instance, to have headers hidden, you'd set `gnus-treat-hide-headers`. Below is a list of variables that can be set, but first we discuss the values these variables can have.

Note: Some values, while valid, make little sense. Check the list below for sensible values.

1. `nil`: Don't do this treatment.

2. `t`: Do this treatment on all body parts.

3. `head`: Do the treatment on the headers.

4. `last`: Do this treatment on the last part.

5. An integer: Do this treatment on all body parts that have a length less than this number.

6. A list of strings: Do this treatment on all body parts that are in articles that are read in groups that have names that match one of the regexps in the list.

7. A list where the first element is not a string:

   The list is evaluated recursively. The first element of the list is a predicate. The following predicates are recognized: `or`, `and`, `not` and `typep`. Here's an example:

   ```
   (or last
       (typep "text/x-vcard"))
   ```

You may have noticed that the word *part* is used here. This refers to the fact that some messages are MIME multipart articles that may be divided into several parts. Articles that are not multiparts are considered to contain just a single part.

Are the treatments applied to all sorts of multipart parts? Yes, if you want to, but by default, only 'text/plain' parts are given the treatment. This is controlled by the **gnus-article-treat-types** variable, which is a list of regular expressions that are matched to the type of the part. This variable is ignored if the value of the controlling variable is a predicate list, as described above.

The following treatment options are available. The easiest way to customize this is to examine the **gnus-article-treat** customization group. Values in parenthesis are suggested sensible values. Others are possible but those listed are probably sufficient for most people.

`gnus-treat-highlight-signature (t, last)`

`gnus-treat-buttonize (t, integer)`

`gnus-treat-buttonize-head (head)`

`gnus-treat-emphasize (t, head, integer)`

`gnus-treat-fill-article (t, integer)`

`gnus-treat-strip-cr (t, integer)`

`gnus-treat-hide-headers (head)`

`gnus-treat-hide-boring-headers (head)`

`gnus-treat-hide-signature (t, last)`

`gnus-treat-hide-citation (t, integer)`

`gnus-treat-strip-pgp (t, last, integer)`

`gnus-treat-strip-pem (t, last, integer)`

`gnus-treat-highlight-headers (head)`

`gnus-treat-highlight-citation (t, integer)`

`gnus-treat-highlight-signature (t, last, integer)`

`gnus-treat-date-ut (head)`

`gnus-treat-date-local (head)`

`gnus-treat-date-lapsed (head)`

`gnus-treat-date-original (head)`

`gnus-treat-strip-headers-in-body (t, integer)`

`gnus-treat-strip-trailing-blank-lines (t, last, integer)`

`gnus-treat-strip-leading-blank-lines (t, integer)`

`gnus-treat-strip-multiple-blank-lines (t, integer)`

`gnus-treat-overstrike (t, integer)`

```
gnus-treat-display-xface (head)
gnus-treat-display-smileys (t, integer)
gnus-treat-display-picons (head)
gnus-treat-capitalize-sentences (t, integer)
gnus-treat-fill-long-lines (t, integer)
gnus-treat-play-sounds
gnus-treat-translate
```

You can, of course, write your own functions to be called from `gnus-part-display-hook`. The functions are called narrowed to the part, and you can do anything you like, pretty much. There is no information that you have to keep in the buffer—you can change everything.

4.4 Article Keymap

Most of the keystrokes in the summary buffer can also be used in the article buffer. They should behave as if you typed them in the summary buffer, which means that you don't actually have to have a summary buffer displayed while reading. You can do it all from the article buffer.

A few additional keystrokes are available:

⟨SPC⟩   Scroll forwards one page (`gnus-article-next-page`).

⟨DEL⟩   Scroll backwards one page (`gnus-article-prev-page`).

`C-c ^`   If point is in the neighborhood of a `Message-ID` and you press `C-c ^`, Gnus will try to get that article from the server (`gnus-article-refer-article`).

`C-c C-m`   Send a reply to the address near point (`gnus-article-mail`). If given a prefix, include the mail.

`s`   Reconfigure the buffers so that the summary buffer becomes visible (`gnus-article-show-summary`).

`?`   Give a very brief description of the available keystrokes (`gnus-article-describe-briefly`).

`TAB`   Go to the next button, if any (`gnus-article-next-button`). This only makes sense if you have buttonizing turned on.

`M-TAB`   Go to the previous button, if any (`gnus-article-prev-button`).

4.5 Misc Article

`gnus-single-article-buffer`

   If non-`nil`, use the same article buffer for all the groups. (This is the default.) If `nil`, each group will have its own article buffer.

`gnus-article-decode-hook`

   Hook used to decode MIME articles. The default value is (`article-decode-charset article-decode-encoded-words`)

`gnus-article-prepare-hook`
> This hook is called right after the article has been inserted into the article buffer. It is mainly intended for functions that do something depending on the contents; it should probably not be used for changing the contents of the article buffer.

`gnus-article-mode-hook`
> Hook called in article mode buffers.

`gnus-article-mode-syntax-table`
> Syntax table used in article buffers. It is initialized from `text-mode-syntax-table`.

`gnus-article-mode-line-format`
> This variable is a format string along the same lines as `gnus-summary-mode-line-format` (see Section 8.4.2 [Mode Line Formatting], page 191). It accepts the same format specifications as that variable, with two extensions:
>
> 'w'   The *wash status* of the article. This is a short string with one character for each possible article wash operation that may have been performed.
>
> 'm'   The number of MIME parts in the article.

`gnus-break-pages`
> Controls whether *page breaking* is to take place. If this variable is non-`nil`, the articles will be divided into pages whenever a page delimiter appears in the article. If this variable is `nil`, paging will not be done.

`gnus-page-delimiter`
> This is the delimiter mentioned above. By default, it is '`^L`' (formfeed).

5 Composing Messages

All commands for posting and mailing will put you in a message buffer where you can edit the article all you like, before you send the article by pressing `C-c C-c`. See section "Top" in *The Message Manual*. Where the message will be posted/mailed to depends on your setup (see Section 5.2 [Posting Server], page 99).

Also see see Section 3.5.4 [Canceling and Superseding], page 46 for information on how to remove articles you shouldn't have posted.

5.1 Mail

Variables for customizing outgoing mail:

`gnus-uu-digest-headers`

> List of regexps to match headers included in digested messages. The headers will be included in the sequence they are matched.

`gnus-add-to-list`

> If non-`nil`, add a `to-list` group parameter to mail groups that have none when you do a `a`.

5.2 Posting Server

When you press those magical `C-c C-c` keys to ship off your latest (extremely intelligent, of course) article, where does it go?

Thank you for asking. I hate you.

It can be quite complicated. Normally, Gnus will use the same native server. However. If your native server doesn't allow posting, just reading, you probably want to use some other server to post your (extremely intelligent and fabulously interesting) articles. You can then set the `gnus-post-method` to some other method:

```
(setq gnus-post-method '(nnspool ""))
```

Now, if you've done this, and then this server rejects your article, or this server is down, what do you do then? To override this variable you can use a non-zero prefix to the `C-c C-c` command to force using the "current" server for posting.

If you give a zero prefix (i.e., `C-u 0 C-c C-c`) to that command, Gnus will prompt you for what method to use for posting.

You can also set `gnus-post-method` to a list of select methods. If that's the case, Gnus will always prompt you for what method to use for posting.

Finally, if you want to always post using the same select method as you're reading from (which might be convenient if you're reading lots of groups from different private servers), you can set this variable to `current`.

5.3 Mail and Post

Here's a list of variables relevant to both mailing and posting:

`gnus-mailing-list-groups`

> If your news server offers groups that are really mailing lists gatewayed to the NNTP server, you can read those groups without problems, but you can't post/followup to them without some difficulty. One solution is to add a `to-address` to the group parameters (see Section 2.10 [Group Parameters], page 22). An easier thing to do is set the `gnus-mailing-list-groups` to a regexp that matches the groups that really are mailing lists. Then, at least, followups to the mailing lists will work most of the time. Posting to these groups (`a`) is still a pain, though.

You may want to do spell-checking on messages that you send out. Or, if you don't want to spell-check by hand, you could add automatic spell-checking via the `ispell` package:

```
(add-hook 'message-send-hook 'ispell-message)
```

If you want to change the `ispell` dictionary based on what group you're in, you could say something like the following:

```
(add-hook 'gnus-select-group-hook
          (lambda ()
            (cond
             ((string-match "^de\\." gnus-newsgroup-name)
              (ispell-change-dictionary "deutsch"))
             (t
              (ispell-change-dictionary "english")))))
```

Modify to suit your needs.

5.4 Archived Messages

Gnus provides a few different methods for storing the mail and news you send. The default method is to use the *archive virtual server* to store the messages. If you want to disable this completely, the `gnus-message-archive-group` variable should be `nil`, which is the default.

`gnus-message-archive-method` says what virtual server Gnus is to use to store sent messages. The default is:

```
(nnfolder "archive"
          (nnfolder-directory   "~/Mail/archive")
          (nnfolder-active-file "~/Mail/archive/active")
          (nnfolder-get-new-mail nil)
          (nnfolder-inhibit-expiry t))
```

You can, however, use any mail select method (`nnml`, `nnmbox`, etc.). `nnfolder` is a quite likable select method for doing this sort of thing, though. If you don't like the default directory chosen, you could say something like:

```
(setq gnus-message-archive-method
      '(nnfolder "archive"
                 (nnfolder-inhibit-expiry t)
```

```
                    (nnfolder-active-file "~/News/sent-mail/active")
                    (nnfolder-directory "~/News/sent-mail/")))
```

Gnus will insert `Gcc` headers in all outgoing messages that point to one or more group(s) on that server. Which group to use is determined by the `gnus-message-archive-group` variable.

This variable can be used to do the following:

- a string Messages will be saved in that group.

  Note that you can include a select method in the group name, then the message will not be stored in the select method given by `gnus-message-archive-method`, but in the select method specified by the group name, instead. Suppose `gnus-message-archive-method` has the default value shown above. Then setting `gnus-message-archive-group` to `"foo"` means that outgoing messages are stored in 'nnfolder+archive:foo', but if you use the value `"nnml:foo"`, then outgoing messages will be stored in 'nnml:foo'.

- a list of strings Messages will be saved in all those groups.

- an alist of regexps, functions and forms When a key "matches", the result is used.

- `nil` No message archiving will take place. This is the default.

Let's illustrate:

Just saving to a single group called 'MisK':

```
   (setq gnus-message-archive-group "MisK")
```

Saving to two groups, 'MisK' and 'safe':

```
   (setq gnus-message-archive-group '("MisK" "safe"))
```

Save to different groups based on what group you are in:

```
   (setq gnus-message-archive-group
         '(("^alt" "sent-to-alt")
           ("mail" "sent-to-mail")
           (".*" "sent-to-misc")))
```

More complex stuff:

```
   (setq gnus-message-archive-group
         '((if (message-news-p)
               "misc-news"
             "misc-mail")))
```

How about storing all news messages in one file, but storing all mail messages in one file per month:

```
   (setq gnus-message-archive-group
         '((if (message-news-p)
               "misc-news"
             (concat "mail." (format-time-string "%Y-%m")))))
```

(XEmacs 19.13 doesn't have `format-time-string`, so you'll have to use a different value for `gnus-message-archive-group` there.)

Now, when you send a message off, it will be stored in the appropriate group. (If you want to disable storing for just one particular message, you can just remove the `Gcc` header that has been inserted.) The archive group will appear in the group buffer the next time

you start Gnus, or the next time you press `F` in the group buffer. You can enter it and read the articles in it just like you'd read any other group. If the group gets really big and annoying, you can simply rename if (using `G r` in the group buffer) to something nice—'`misc-mail-september-1995`', or whatever. New messages will continue to be stored in the old (now empty) group.

That's the default method of archiving sent messages. Gnus offers a different way for the people who don't like the default method. In that case you should set `gnus-message-archive-group` to `nil`; this will disable archiving.

`gnus-outgoing-message-group`

>All outgoing messages will be put in this group. If you want to store all your outgoing mail and articles in the group '`nnml:archive`', you set this variable to that value. This variable can also be a list of group names.

>If you want to have greater control over what group to put each message in, you can set this variable to a function that checks the current newsgroup name and then returns a suitable group name (or list of names).

>This variable can be used instead of `gnus-message-archive-group`, but the latter is the preferred method.

5.5 Posting Styles

All them variables, they make my head swim.

So what if you want a different `Organization` and signature based on what groups you post to? And you post both from your home machine and your work machine, and you want different `From` lines, and so on?

One way to do stuff like that is to write clever hooks that change the variables you need to have changed. That's a bit boring, so somebody came up with the bright idea of letting the user specify these things in a handy alist. Here's an example of a `gnus-posting-styles` variable:

```
((".*"
  (signature "Peace and happiness")
  (organization "What me?"))
 ("^comp"
  (signature "Death to everybody"))
 ("comp.emacs.i-love-it"
  (organization "Emacs is it")))
```

As you might surmise from this example, this alist consists of several *styles*. Each style will be applicable if the first element "matches", in some form or other. The entire alist will be iterated over, from the beginning towards the end, and each match will be applied, which means that attributes in later styles that match override the same attributes in earlier matching styles. So '`comp.programming.literate`' will have the '`Death to everybody`' signature and the '`What me?`' `Organization` header.

The first element in each style is called the `match`. If it's a string, then Gnus will try to regexp match it against the group name. If it is the symbol `header`, then Gnus will look for header that match the next element in the match, and compare that to the last header in the match. If it's a function symbol, that function will be called with no arguments. If

it's a variable symbol, then the variable will be referenced. If it's a list, then that list will be `evaled`. In any case, if this returns a non-`nil` value, then the style is said to *match*.

Each style may contain a arbitrary amount of *attributes*. Each attribute consists of a (*name . value*) pair. The attribute name can be one of `signature`, `signature-file`, `organization`, `address`, `name` or `body`. The attribute name can also be a string. In that case, this will be used as a header name, and the value will be inserted in the headers of the article; if the value is `nil`, the header name will be removed. If the attribute name is `eval`, the form is evaluated, and the result is thrown away.

The attribute value can be a string (used verbatim), a function with zero arguments (the return value will be used), a variable (its value will be used) or a list (it will be `evaled` and the return value will be used). The functions and sexps are called/`evaled` in the message buffer that is being set up. The headers of the current article are available through the `message-reply-headers` variable.

If you wish to check whether the message you are about to compose is meant to be a news article or a mail message, you can check the values of the `message-news-p` and `message-mail-p` functions.

So here's a new example:

```
(setq gnus-posting-styles
      '((".*"
         (signature-file "~/.signature")
         (name "User Name")
         ("X-Home-Page" (getenv "WWW_HOME"))
         (organization "People's Front Against MWM"))
        ("^rec.humor"
         (signature my-funny-signature-randomizer))
        ((equal (system-name) "gnarly")
         (signature my-quote-randomizer))
        ((message-news-p)
         (signature my-news-signature))
        (header "From\\|To" "larsi.*org"
                (Organization "Somewhere, Inc."))
        ((posting-from-work-p)
         (signature-file "~/.work-signature")
         (address "user@bar.foo")
         (body "You are fired.\n\nSincerely, your boss.")
         (organization "Important Work, Inc"))
        ("nnml:.*"
         (From (save-excursion
                 (set-buffer gnus-article-buffer)
                 (message-fetch-field "to"))))
        ("^nn.+:"
         (signature-file "~/.mail-signature"))))
```

The 'nnml:.*' rule means that you use the `To` address as the `From` address in all your outgoing replies, which might be handy if you fill many roles.

5.6 Drafts

If you are writing a message (mail or news) and suddenly remember that you have a steak in the oven (or some pesto in the food processor, you craaazy vegetarians), you'll probably wish there was a method to save the message you are writing so that you can continue editing it some other day, and send it when you feel its finished.

Well, don't worry about it. Whenever you start composing a message of some sort using the Gnus mail and post commands, the buffer you get will automatically associate to an article in a special *draft* group. If you save the buffer the normal way (`C-x C-s`, for instance), the article will be saved there. (Auto-save files also go to the draft group.)

The draft group is a special group (which is implemented as an `nndraft` group, if you absolutely have to know) called '`nndraft:drafts`'. The variable `nndraft-directory` says where `nndraft` is to store its files. What makes this group special is that you can't tick any articles in it or mark any articles as read—all articles in the group are permanently unread.

If the group doesn't exist, it will be created and you'll be subscribed to it. The only way to make it disappear from the Group buffer is to unsubscribe it.

When you want to continue editing the article, you simply enter the draft group and push `D e` (`gnus-draft-edit-message`) to do that. You will be placed in a buffer where you left off.

Rejected articles will also be put in this draft group (see Section 5.7 [Rejected Articles], page 104).

If you have lots of rejected messages you want to post (or mail) without doing further editing, you can use the `D s` command (`gnus-draft-send-message`). This command understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189). The `D S` command (`gnus-draft-send-all-messages`) will ship off all messages in the buffer.

If you have some messages that you wish not to send, you can use the `D t` (`gnus-draft-toggle-sending`) command to mark the message as unsendable. This is a toggling command.

5.7 Rejected Articles

Sometimes a news server will reject an article. Perhaps the server doesn't like your face. Perhaps it just feels miserable. Perhaps *there be demons.* Perhaps you have included too much cited text. Perhaps the disk is full. Perhaps the server is down.

These situations are, of course, totally beyond the control of Gnus. (Gnus, of course, loves the way you look, always feels great, has angels fluttering around inside of it, doesn't care about how much cited text you include, never runs full and never goes down.) So Gnus saves these articles until some later time when the server feels better.

The rejected articles will automatically be put in a special draft group (see Section 5.6 [Drafts], page 104). When the server comes back up again, you'd then typically enter that group and send all the articles off.

6 Select Methods

A *foreign group* is a group not read by the usual (or default) means. It could be, for instance, a group from a different NNTP server, it could be a virtual group, or it could be your own personal mail group.

A foreign group (or any group, really) is specified by a *name* and a *select method*. To take the latter first, a select method is a list where the first element says what back end to use (e.g. `nntp`, `nnspool`, `nnml`) and the second element is the *server name*. There may be additional elements in the select method, where the value may have special meaning for the back end in question.

One could say that a select method defines a *virtual server*—so we do just that (see Section 6.1 [The Server Buffer], page 105).

The *name* of the group is the name the back end will recognize the group as.

For instance, the group 'soc.motss' on the NNTP server 'some.where.edu' will have the name 'soc.motss' and select method (`nntp "some.where.edu"`). Gnus will call this group 'nntp+some.where.edu:soc.motss', even though the `nntp` back end just knows this group as 'soc.motss'.

The different methods all have their peculiarities, of course.

6.1 The Server Buffer

Traditionally, a *server* is a machine or a piece of software that one connects to, and then requests information from. Gnus does not connect directly to any real servers, but does all transactions through one back end or other. But that's just putting one layer more between the actual media and Gnus, so we might just as well say that each back end represents a virtual server.

For instance, the `nntp` back end may be used to connect to several different actual NNTP servers, or, perhaps, to many different ports on the same actual NNTP server. You tell Gnus which back end to use, and what parameters to set by specifying a *select method*.

These select method specifications can sometimes become quite complicated—say, for instance, that you want to read from the NNTP server 'news.funet.fi' on port number 13, which hangs if queried for NOV headers and has a buggy select. Ahem. Anyway, if you had to specify that for each group that used this server, that would be too much work, so Gnus offers a way of naming select methods, which is what you do in the server buffer.

To enter the server buffer, use the `^` (`gnus-group-enter-server-mode`) command in the group buffer.

`gnus-server-mode-hook` is run when creating the server buffer.

6.1.1 Server Buffer Format

You can change the look of the server buffer lines by changing the `gnus-server-line-format` variable. This is a `format`-like variable, with some simple extensions:

'h'   How the news is fetched—the back end name.

'n'   The name of this server.

'w'   Where the news is to be fetched from—the address.

's'   The opened/closed/denied status of the server.

The mode line can also be customized by using the `gnus-server-mode-line-format` variable (see Section 8.4.2 [Mode Line Formatting], page 191). The following specs are understood:

'S'   Server name.

'M'   Server method.

Also see Section 8.4 [Formatting Variables], page 190.

6.1.2 Server Commands

a   Add a new server (`gnus-server-add-server`).

e   Edit a server (`gnus-server-edit-server`).

⟨SPC⟩   Browse the current server (`gnus-server-read-server`).

q   Return to the group buffer (`gnus-server-exit`).

k   Kill the current server (`gnus-server-kill-server`).

y   Yank the previously killed server (`gnus-server-yank-server`).

c   Copy the current server (`gnus-server-copy-server`).

l   List all servers (`gnus-server-list-servers`).

s   Request that the server scan its sources for new articles (`gnus-server-scan-server`). This is mainly sensible with mail servers.

g   Request that the server regenerate all its data structures (`gnus-server-regenerate-server`). This can be useful if you have a mail back end that has gotten out of sync.

6.1.3 Example Methods

Most select methods are pretty simple and self-explanatory:

```
(nntp "news.funet.fi")
```

Reading directly from the spool is even simpler:

```
(nnspool "")
```

As you can see, the first element in a select method is the name of the back end, and the second is the *address*, or *name*, if you will.

After these two elements, there may be an arbitrary number of (*variable form*) pairs.

To go back to the first example—imagine that you want to read from port 15 on that machine. This is what the select method should look like then:

```
    (nntp "news.funet.fi" (nntp-port-number 15))
```

You should read the documentation to each back end to find out what variables are relevant, but here's an **nnmh** example:

nnmh is a mail back end that reads a spool-like structure. Say you have two structures that you wish to access: One is your private mail spool, and the other is a public one. Here's the possible spec for your private mail:

```
    (nnmh "private" (nnmh-directory "~/private/mail/"))
```

(This server is then called 'private', but you may have guessed that.)

Here's the method for a public spool:

```
    (nnmh "public"
          (nnmh-directory "/usr/information/spool/")
          (nnmh-get-new-mail nil))
```

If you are behind a firewall and only have access to the NNTP server from the firewall machine, you can instruct Gnus to **rlogin** on the firewall machine and telnet from there to the NNTP server. Doing this can be rather fiddly, but your virtual server definition should probably look something like this:

```
    (nntp "firewall"
          (nntp-address "the.firewall.machine")
          (nntp-open-connection-function nntp-open-rlogin)
          (nntp-end-of-line "\n")
          (nntp-rlogin-parameters
           ("telnet" "the.real.nntp.host" "nntp")))
```

If you want to use the wonderful **ssh** program to provide a compressed connection over the modem line, you could create a virtual server that would look something like this:

```
    (nntp "news"
          (nntp-address "copper.uio.no")
          (nntp-rlogin-program "ssh")
          (nntp-open-connection-function nntp-open-rlogin)
          (nntp-end-of-line "\n")
          (nntp-rlogin-parameters
           ("telnet" "news.uio.no" "nntp")))
```

This means that you have to have set up **ssh-agent** correctly to provide automatic authorization, of course. And to get a compressed connection, you have to have the 'Compression' option in the **ssh** 'config' file.

6.1.4 Creating a Virtual Server

If you're saving lots of articles in the cache by using persistent articles, you may want to create a virtual server to read the cache.

First you need to add a new server. The **a** command does that. It would probably be best to use **nnspool** to read the cache. You could also use **nnml** or **nnmh**, though.

Type *a nnspool* ⟨RET⟩ *cache* ⟨RET⟩.

You should now have a brand new **nnspool** virtual server called 'cache'. You now need to edit it to have the right definitions. Type *e* to edit the server. You'll be entered into a buffer that will contain the following:

```
    (nnspool "cache")
```
Change that to:
```
    (nnspool "cache"
             (nnspool-spool-directory "~/News/cache/")
             (nnspool-nov-directory "~/News/cache/")
             (nnspool-active-file "~/News/cache/active"))
```

Type `C-c C-c` to return to the server buffer. If you now press ⟨RET⟩ over this virtual server, you should be entered into a browse buffer, and you should be able to enter any of the groups displayed.

6.1.5 Server Variables

One sticky point when defining variables (both on back ends and in Emacs in general) is that some variables are typically initialized from other variables when the definition of the variables is being loaded. If you change the "base" variable after the variables have been loaded, you won't change the "derived" variables.

This typically affects directory and file variables. For instance, `nnml-directory` is '`~/Mail/`' by default, and all `nnml` directory variables are initialized from that variable, so `nnml-active-file` will be '`~/Mail/active`'. If you define a new virtual `nnml` server, it will *not* suffice to set just `nnml-directory`—you have to explicitly set all the file variables to be what you want them to be. For a complete list of variables for each back end, see each back end's section later in this manual, but here's an example `nnml` definition:

```
    (nnml "public"
          (nnml-directory "~/my-mail/")
          (nnml-active-file "~/my-mail/active")
          (nnml-newsgroups-file "~/my-mail/newsgroups"))
```

6.1.6 Servers and Methods

Wherever you would normally use a select method (e.g. `gnus-secondary-select-method`, in the group select method, when browsing a foreign server) you can use a virtual server name instead. This could potentially save lots of typing. And it's nice all over.

6.1.7 Unavailable Servers

If a server seems to be unreachable, Gnus will mark that server as `denied`. That means that any subsequent attempt to make contact with that server will just be ignored. "It can't be opened," Gnus will tell you, without making the least effort to see whether that is actually the case or not.

That might seem quite naughty, but it does make sense most of the time. Let's say you have 10 groups subscribed to on server '`nephelococcygia.com`'. This server is located somewhere quite far away from you and the machine is quite slow, so it takes 1 minute just to find out that it refuses connection to you today. If Gnus were to attempt to do that 10 times, you'd be quite annoyed, so Gnus won't attempt to do that. Once it has gotten a single "connection refused", it will regard that server as "down".

So, what happens if the machine was only feeling unwell temporarily? How do you test to see whether the machine has come up again?

You jump to the server buffer (see Section 6.1 [The Server Buffer], page 105) and poke it with the following commands:

O   Try to establish connection to the server on the current line (**gnus-server-open-server**).

C   Close the connection (if any) to the server (**gnus-server-close-server**).

D   Mark the current server as unreachable (**gnus-server-deny-server**).

M-o   Open the connections to all servers in the buffer (**gnus-server-open-all-servers**).

M-c   Close the connections to all servers in the buffer (**gnus-server-close-all-servers**).

R   Remove all marks to whether Gnus was denied connection from any servers (**gnus-server-remove-denials**).

6.2 Getting News

A newsreader is normally used for reading news. Gnus currently provides only two methods of getting news—it can read from an NNTP server, or it can read from a local spool.

6.2.1 NNTP

Subscribing to a foreign group from an NNTP server is rather easy. You just specify **nntp** as method and the address of the NNTP server as the, uhm, address.

If the NNTP server is located at a non-standard port, setting the third element of the select method to this port number should allow you to connect to the right port. You'll have to edit the group info for that (see Section 2.9 [Foreign Groups], page 21).

The name of the foreign group can be the same as a native group. In fact, you can subscribe to the same group from as many different servers you feel like. There will be no name collisions.

The following variables can be used to create a virtual **nntp** server:

nntp-server-opened-hook
        is run after a connection has been made. It can be used to send commands to the NNTP server after it has been contacted. By default it sends the command **MODE READER** to the server with the **nntp-send-mode-reader** function. This function should always be present in this hook.

nntp-authinfo-function
        This function will be used to send 'AUTHINFO' to the NNTP server. The default function is **nntp-send-authinfo**, which looks through your '~/.authinfo' (or whatever you've set the **nntp-authinfo-file** variable to) for applicable entries. If none are found, it will prompt you for a login name and a password. The format of the '~/.authinfo' file is (almost) the same as the **ftp** '~/.netrc' file, which is defined in the **ftp** manual page, but here are the salient facts:

    1. The file contains one or more line, each of which define one server.

2. Each line may contain an arbitrary number of token/value pairs.

   The valid tokens include 'machine', 'login', 'password', 'default'. In
   addition Gnus introduces two new tokens, not present in the original
   '.netrc'/ftp syntax, namely 'port' and 'force'. (This is the only way
   the '.authinfo' file format deviates from the '.netrc' file format.) 'port'
   is used to indicate what port on the server the credentials apply to and
   'force' is explained below.

Here's an example file:

```
machine news.uio.no login larsi password geheimnis
machine nntp.ifi.uio.no login larsi force yes
```

The token/value pairs may appear in any order; 'machine' doesn't have to be
first, for instance.

In this example, both login name and password have been supplied for the
former server, while the latter has only the login name listed, and the user will
be prompted for the password. The latter also has the 'force' tag, which means
that the authinfo will be sent to the *nntp* server upon connection; the default
(i.e., when there is not 'force' tag) is to not send authinfo to the *nntp* server
until the *nntp* server asks for it.

You can also add 'default' lines that will apply to all servers that don't have
matching 'machine' lines.

```
default force yes
```

This will force sending 'AUTHINFO' commands to all servers not previously men-
tioned.

Remember to not leave the '~/.authinfo' file world-readable.

nntp-server-action-alist

This is a list of regexps to match on server types and actions to be taken when
matches are made. For instance, if you want Gnus to beep every time you
connect to innd, you could say something like:

```
(setq nntp-server-action-alist
      '(("innd" (ding))))
```

You probably don't want to do that, though.

The default value is

```
'(("nntpd 1\\.5\\.11t"
   (remove-hook 'nntp-server-opened-hook
                'nntp-send-mode-reader)))
```

This ensures that Gnus doesn't send the MODE READER command to nntpd
1.5.11t, since that command chokes that server, I've been told.

nntp-maximum-request

If the NNTP server doesn't support NOV headers, this back end will collect
headers by sending a series of head commands. To speed things up, the back
end sends lots of these commands without waiting for reply, and then reads all
the replies. This is controlled by the nntp-maximum-request variable, and is
400 by default. If your network is buggy, you should set this to 1.

`nntp-connection-timeout`

> If you have lots of foreign `nntp` groups that you connect to regularly, you're sure to have problems with NNTP servers not responding properly, or being too loaded to reply within reasonable time. This is can lead to awkward problems, which can be helped somewhat by setting `nntp-connection-timeout`. This is an integer that says how many seconds the `nntp` back end should wait for a connection before giving up. If it is `nil`, which is the default, no timeouts are done.

`nntp-server-hook`

> This hook is run as the last step when connecting to an NNTP server.

`nntp-open-connection-function`

> This function is used to connect to the remote system. Four pre-made functions are supplied:

> > `nntp-open-network-stream`
> >
> > > This is the default, and simply connects to some port or other on the remote system.

> > `nntp-open-rlogin`
> >
> > > Does an 'rlogin' on the remote system, and then does a 'telnet' to the NNTP server available there.
> > >
> > > `nntp-open-rlogin`-related variables:
> > >
> > > `nntp-rlogin-program`
> > >
> > > > Program used to log in on remote machines. The default is 'rsh', but 'ssh' is a popular alternative.
> > >
> > > `nntp-rlogin-parameters`
> > >
> > > > This list will be used as the parameter list given to `rsh`.
> > >
> > > `nntp-rlogin-user-name`
> > >
> > > > User name on the remote system.

> > `nntp-open-telnet`
> >
> > > Does a 'telnet' to the remote system and then another 'telnet' to get to the NNTP server.
> > >
> > > `nntp-open-telnet`-related variables:
> > >
> > > `nntp-telnet-command`
> > >
> > > > Command used to start `telnet`.
> > >
> > > `nntp-telnet-switches`
> > >
> > > > List of strings to be used as the switches to the `telnet` command.
> > >
> > > `nntp-telnet-user-name`
> > >
> > > > User name for log in on the remote system.
> > >
> > > `nntp-telnet-passwd`
> > >
> > > > Password to use when logging in.

nntp-telnet-parameters
:   A list of strings executed as a command after logging in via `telnet`.

nntp-telnet-shell-prompt
:   Regexp matching the shell prompt on the remote machine. The default is '`bash\\|\$ *\r?$\\|> *\r?`'.

nntp-open-telnet-envuser
:   If non-`nil`, the `telnet` session (client and server both) will support the `ENVIRON` option and not prompt for login name. This works for Solaris `telnet`, for instance.

nntp-open-ssl-stream
:   Opens a connection to a server over a *secure* channel. To use this you must have SSLay installed (`ftp://ftp.psy.uq.oz.au/pub/Crypto/SSL`, and you also need '`ssl.el`' (from the W3 distribution, for instance). You then define a server as follows:

    ```
    ;; Type 'C-c C-c' after you've finished editing.
    ;;
    ;; "snews" is port 563 and is predefined
    ;; in our /etc/services
    ;;
    (nntp "snews.bar.com"
          (nntp-open-connection-function
            nntp-open-ssl-stream)
          (nntp-port-number "snews")
          (nntp-address "snews.bar.com"))
    ```

nntp-end-of-line
:   String to use as end-of-line marker when talking to the NNTP server. This is '`\r\n`' by default, but should be '`\n`' when using `rlogin` to talk to the server.

nntp-rlogin-user-name
:   User name on the remote system when using the `rlogin` connect function.

nntp-address
:   The address of the remote system running the NNTP server.

nntp-port-number
:   Port number to connect to when using the `nntp-open-network-stream` connect function.

nntp-buggy-select
:   Set this to non-`nil` if your select routine is buggy.

nntp-nov-is-evil
:   If the NNTP server does not support NOV, you could set this variable to `t`, but `nntp` usually checks automatically whether NOV can be used.

nntp-xover-commands
:   List of strings used as commands to fetch NOV lines from a server. The default value of this variable is (`"XOVER"` `"XOVERVIEW"`).

`nntp-nov-gap`

> `nntp` normally sends just one big request for NOV lines to the server. The server responds with one huge list of lines. However, if you have read articles 2-5000 in the group, and only want to read article 1 and 5001, that means that `nntp` will fetch 4999 NOV lines that you will not need. This variable says how big a gap between two consecutive articles is allowed to be before the `XOVER` request is split into several request. Note that if your network is fast, setting this variable to a really small number means that fetching will probably be slower. If this variable is `nil`, `nntp` will never split requests. The default is 5.

`nntp-prepare-server-hook`

> A hook run before attempting to connect to an NNTP server.

`nntp-warn-about-losing-connection`

> If this variable is non-`nil`, some noise will be made when a server closes connection.

`nntp-record-commands`

> If non-`nil`, `nntp` will log all commands it sends to the NNTP server (along with a timestamp) in the '`*nntp-log*`' buffer. This is useful if you are debugging a Gnus/NNTP connection that doesn't seem to work.

6.2.2 News Spool

Subscribing to a foreign group from the local spool is extremely easy, and might be useful, for instance, to speed up reading groups that contain very big articles—'`alt.binaries.pictures.furniture`', for instance.

Anyway, you just specify `nnspool` as the method and `""` (or anything else) as the address.

If you have access to a local spool, you should probably use that as the native select method (see Section 1.1 [Finding the News], page 3). It is normally faster than using an `nntp` select method, but might not be. It depends. You just have to try to find out what's best at your site.

`nnspool-inews-program`

> Program used to post an article.

`nnspool-inews-switches`

> Parameters given to the inews program when posting an article.

`nnspool-spool-directory`

> Where `nnspool` looks for the articles. This is normally '`/usr/spool/news/`'.

`nnspool-nov-directory`

> Where `nnspool` will look for NOV files. This is normally '`/usr/spool/news/over.view/`'.

`nnspool-lib-dir`

> Where the news lib dir is ('`/usr/lib/news/`' by default).

`nnspool-active-file`

> The path to the active file.

`nnspool-newsgroups-file`
>   The path to the group descriptions file.

`nnspool-history-file`
>   The path to the news history file.

`nnspool-active-times-file`
>   The path to the active date file.

`nnspool-nov-is-evil`
>   If non-`nil`, `nnspool` won't try to use any NOV files that it finds.

`nnspool-sift-nov-with-sed`
>   If non-`nil`, which is the default, use `sed` to get the relevant portion from the
>   overview file. If nil, `nnspool` will load the entire file into a buffer and process
>   it there.

6.3 Getting Mail

Reading mail with a newsreader—isn't that just plain WeIrD? But of course.

6.3.1 Mail in a Newsreader

If you are used to traditional mail readers, but have decided to switch to reading mail
with Gnus, you may find yourself experiencing something of a culture shock.

Gnus does not behave like traditional mail readers. If you want to make it behave that
way, you can, but it's an uphill battle.

Gnus, by default, handles all its groups using the same approach. This approach is
very newsreaderly—you enter a group, see the new/unread messages, and when you read
the messages, they get marked as read, and you don't see them any more. (Unless you
explicitly ask for them.)

In particular, you do not do anything explicitly to delete messages.

Does this mean that all the messages that have been marked as read are deleted? How
awful!

But, no, it means that old messages are *expired* according to some scheme or other. For
news messages, the expire process is controlled by the news administrator; for mail, the
expire process is controlled by you. The expire process for mail is covered in depth in see
Section 6.3.9 [Expiring Mail], page 128.

What many Gnus users find, after using it a while for both news and mail, is that the
transport mechanism has very little to do with how they want to treat a message.

Many people subscribe to several mailing lists. These are transported via SMTP, and
are therefore mail. But we might go for weeks without answering, or even reading these
messages very carefully. We may not need to save them because if we should need to read
one again, they are archived somewhere else.

Some people have local news groups which have only a handful of readers. These are
transported via NNTP, and are therefore news. But we may need to read and answer a large
fraction of the messages very carefully in order to do our work. And there may not be an

archive, so we may need to save the interesting messages the same way we would personal mail.

The important distinction turns out to be not the transport mechanism, but other factors such as how interested we are in the subject matter, or how easy it is to retrieve the message if we need to read it again.

Gnus provides many options for sorting mail into "groups" which behave like newsgroups, and for treating each group (whether mail or news) differently.

Some users never get comfortable using the Gnus (ahem) paradigm and wish that Gnus should grow up and be a male, er, mail reader. It is possible to whip Gnus into a more mailreaderly being, but, as said before, it's not easy. People who prefer proper mail readers should try VM instead, which is an excellent, and proper, mail reader.

I don't mean to scare anybody off, but I want to make it clear that you may be required to learn a new way of thinking about messages. After you've been subjected to The Gnus Way, you will come to love it. I can guarantee it. (At least the guy who sold me the Emacs Subliminal Brain-Washing Functions that I've put into Gnus did guarantee it. You Will Be Assimilated. You Love Gnus. You Love The Gnus Mail Way. You Do.)

6.3.2 Getting Started Reading Mail

It's quite easy to use Gnus to read your new mail. You just plonk the mail back end of your choice into `gnus-secondary-select-methods`, and things will happen automatically.

For instance, if you want to use `nnml` (which is a "one file per mail" back end), you could put the following in your '`.gnus`' file:

```
(setq gnus-secondary-select-methods
      '((nnml "private")))
```

Now, the next time you start Gnus, this back end will be queried for new articles, and it will move all the messages in your spool file to its directory, which is `~/Mail/` by default. The new group that will be created ('`mail.misc`') will be subscribed, and you can read it like any other group.

You will probably want to split the mail into several groups, though:

```
(setq nnmail-split-methods
      '(("junk" "^From:.*Lars Ingebrigtsen")
        ("crazy" "^Subject:.*die\\|^Organization:.*flabby")
        ("other" "")))
```

This will result in three new `nnml` mail groups being created: '`nnml:junk`', '`nnml:crazy`', and '`nnml:other`'. All the mail that doesn't fit into the first two groups will be placed in the last group.

This should be sufficient for reading mail with Gnus. You might want to give the other sections in this part of the manual a perusal, though. Especially see Section 6.3.13 [Choosing a Mail Back End], page 133 and see Section 6.3.9 [Expiring Mail], page 128.

6.3.3 Splitting Mail

The `nnmail-split-methods` variable says how the incoming mail is to be split into groups.

```
(setq nnmail-split-methods
  '(("mail.junk" "^From:.*Lars Ingebrigtsen")
    ("mail.crazy" "^Subject:.*die\\|^Organization:.*flabby")
    ("mail.other" "")))
```

This variable is a list of lists, where the first element of each of these lists is the name of the mail group (they do not have to be called something beginning with 'mail', by the way), and the second element is a regular expression used on the header of each mail to determine if it belongs in this mail group. The first string may contain '\\1' forms, like the ones used by `replace-match` to insert sub-expressions from the matched text. For instance:

```
("list.\\1" "From:.* \\(.*\\)-list@majordomo.com")
```

The second element can also be a function. In that case, it will be called narrowed to the headers with the first element of the rule as the argument. It should return a non-`nil` value if it thinks that the mail belongs in that group.

The last of these groups should always be a general one, and the regular expression should *always* be '' so that it matches any mails that haven't been matched by any of the other regexps. (These rules are processed from the beginning of the alist toward the end. The first rule to make a match will "win", unless you have crossposting enabled. In that case, all matching rules will "win".)

If you like to tinker with this yourself, you can set this variable to a function of your choice. This function will be called without any arguments in a buffer narrowed to the headers of an incoming mail message. The function should return a list of group names that it thinks should carry this mail message.

Note that the mail back ends are free to maul the poor, innocent, incoming headers all they want to. They all add `Lines` headers; some add `X-Gnus-Group` headers; most rename the Unix mbox 'From ' line to something else.

The mail back ends all support cross-posting. If several regexps match, the mail will be "cross-posted" to all those groups. `nnmail-crosspost` says whether to use this mechanism or not. Note that no articles are crossposted to the general ('') group.

`nnmh` and `nnml` makes crossposts by creating hard links to the crossposted articles. However, not all file systems support hard links. If that's the case for you, set `nnmail-crosspost-link-function` to `copy-file`. (This variable is `add-name-to-file` by default.)

If you wish to see where the previous mail split put the messages, you can use the `M-x nnmail-split-history` command. If you wish to see where re-spooling messages would put the messages, you can use `gnus-summary-respool-trace` and related commands (see Section 3.24 [Mail Group Commands], page 86).

Gnus gives you all the opportunity you could possibly want for shooting yourself in the foot. Let's say you create a group that will contain all the mail you get from your boss. And then you accidentally unsubscribe from the group. Gnus will still put all the mail from your boss in the unsubscribed group, and so, when your boss mails you "Have that report ready by Monday or you're fired!", you'll never see it and, come Tuesday, you'll still believe that you're gainfully employed while you really should be out collecting empty bottles to save up for next month's rent money.

6.3.4 Mail Sources

Mail can be gotten from many different sources—the mail spool, from a POP mail server, from a procmail directory, or from a maildir, for instance.

6.3.4.1 Mail Source Specifiers

You tell Gnus how to fetch mail by setting `mail-sources` (see Section 6.3.4.4 [Fetching Mail], page 122) to a *mail source specifier*.

Here's an example:

```
(pop :server "pop3.mailserver.com" :user "myname")
```

As can be observed, a mail source specifier is a list where the first element is a *mail source type*, followed by an arbitrary number of *keywords*. Keywords that are not explicitly specified are given default values.

The following mail source types are available:

`file`   Get mail from a single file; typically from the mail spool.

Keywords:

`:path`   The path of the file. Defaults to the value of the `MAIL` environment variable or '`/usr/mail/spool/user-name`'.

An example file mail source:

```
(file :path "/usr/spool/mail/user-name")
```

Or using the default path:

```
(file)
```

If the mail spool file is not located on the local machine, it's best to use POP or IMAP or the like to fetch the mail. You can not use ange-ftp file names here—it has no way to lock the mail spool while moving the mail.

If it's impossible to set up a proper server, you can use ssh instead.

```
(setq mail-sources
      '((file :prescript "ssh host bin/getmail >%t")))
```

The '`getmail`' script would look something like the following:

```
#!/bin/sh
#  getmail - move mail from spool to stdout
#  flu@iki.fi

MOVEMAIL=/usr/lib/emacs/20.3/i386-redhat-linux/movemail
TMP=$HOME/Mail/tmp
rm -f $TMP; $MOVEMAIL $MAIL $TMP >/dev/null && cat $TMP
```

Alter this script to fit find the '`movemail`' you want to use.

`directory`

Get mail from several files in a directory. This is typically used when you have procmail split the incoming mail into several files. Setting `nnmail-scan-directory-mail-source-once` to non-nil force Gnus to scan the mail source only once.

Keywords:

:path   The path of the directory where the files are. There is no default
           value.

:suffix   Only files ending with this suffix are used. The default is '.spool'.

:predicate
           Only files that have this predicate return non-nil are returned. The
           default is identity. This is used as an additional filter—only files
           that have the right suffix *and* satisfy this predicate are considered.

:prescript
:postscript
           Script run before/after fetching mail.

An example directory mail source:

```
(directory :path "/home/user-name/procmail-dir/"
           :suffix ".prcml")
```

pop   Get mail from a POP server.

           Keywords:

:server   The name of the POP server. The default is taken from the
           MAILHOST environment variable.

:port   The port number of the POP server. This can be a number (e.g.
           ':port 110') or a string (e.g. ':port "pop3"'). If it is a string,
           it should be a service name as listed in '/etc/services' on Unix
           systems. The default is '"pop3"'. On some systems you might need
           to specify it as '"pop-3"' instead.

:user   The user name to give to the POP server. The default is the login
           name.

:password
           The password to give to the POP server. If not specified, the user
           is prompted.

:program   The program to use to fetch mail from the POP server. This should
           be a format-like string. Here's an example:

```
fetchmail %u@%s -P %p %t
```

           The valid format specifier characters are:

           't'   The name of the file the mail is to be moved to. This
                      must always be included in this string.

           's'   The name of the server.

           'P'   The port number of the server.

           'u'   The user name to use.

           'p'   The password to use.

           The values used for these specs are taken from the values you give
           the corresponding keywords.

:prescript

> A script to be run before fetching the mail. The syntax is the same as the :program keyword. This can also be a function to be run.

:postscript

> A script to be run after fetching the mail. The syntax is the same as the :program keyword. This can also be a function to be run.

:function

> The function to use to fetch mail from the POP server. The function is called with one parameter—the name of the file where the mail should be moved to.

:authentication

> This can be either the symbol password or the symbol apop and says what authentication scheme to use. The default is password.

If the :program and :function keywords aren't specified, pop3-movemail will be used.

Here are some examples. Fetch from the default POP server, using the default user name, and default fetcher:

```
(pop)
```

Fetch from a named server with a named user and password:

```
(pop :server "my.pop.server"
     :user "user-name" :password "secret")
```

Use 'movemail' to move the mail:

```
(pop :program "movemail po:%u %t %p")
```

maildir   Get mail from a maildir. This is a type of mailbox that is supported by at least qmail and postfix, where each file in a special directory contains exactly one mail.

Keywords:

:path   The path of the directory where the mails are stored. The default is taken from the MAILDIR environment variable or '~/Maildir/'.

:subdirs   The subdirectories of the Maildir. The default is '("new" "cur")'.

> You can also get mails from remote hosts (because maildirs don't suffer from locking problems).

Two example maildir mail sources:

```
(maildir :path "/home/user-name/Maildir/"
         :subdirs ("cur" "new"))
```

```
(maildir :path "/user@remotehost.org:~/Maildir/"
         :subdirs ("new"))
```

imap   Get mail from a IMAP server. If you don't want to use IMAP as intended, as a network mail reading protocol (ie with nnimap), for some reason or other, Gnus let you treat it similar to a POP server and fetches articles from a given IMAP mailbox. See Section 6.5.6 [IMAP], page 149, for more information.

Note that for SSL support you may need external programs and libraries, See
Section 6.5.6 [IMAP], page 149.

Keywords:

`:server`   The name of the IMAP server. The default is taken from the
           `MAILHOST` environment variable.

`:port`   The port number of the IMAP server. The default is '`143`', or '`993`'
           for SSL connections.

`:user`   The user name to give to the IMAP server. The default is the login
           name.

`:password`

           The password to give to the IMAP server. If not specified, the user
           is prompted.

`:stream`   What stream to use for connecting to the server, this is one
           of the symbols in `imap-stream-alist`. Right now, this means
           '`kerberos4`', '`ssl`' or the default '`network`'.

`:authentication`

           Which authenticator to use for authenticating to the server, this
           is one of the symbols in `imap-authenticator-alist`. Right now,
           this means '`kerberos4`', '`cram-md5`', '`anonymous`' or the default
           '`login`'.

`:program`   When using the 'shell' :stream, the contents of this variable is
           mapped into the 'imap-shell-program' variable. This should be a
           `format`-like string (or list of strings). Here's an example:

               `ssh %s imapd`

           The valid format specifier characters are:

           '`s`'   The name of the server.

           '`l`'   User name from 'imap-default-user'.

           '`p`'   The port number of the server.

           The values used for these specs are taken from the values you give
           the corresponding keywords.

`:mailbox`   The name of the mailbox to get mail from. The default is '`INBOX`'
           which normally is the mailbox which receive incoming mail.

`:predicate`

           The predicate used to find articles to fetch. The default, '`UNSEEN`
           `UNDELETED`', is probably the best choice for most people, but if you
           sometimes peek in your mailbox with a IMAP client and mark some
           articles as read (or; SEEN) you might want to set this to '`nil`'.
           Then all articles in the mailbox is fetched, no matter what. For a
           complete list of predicates, see RFC 2060 6.4.4.

:fetchflag

> How to flag fetched articles on the server, the default '\Deleted'
> will mark them as deleted, an alternative would be '\Seen' which
> would simply mark them as read. These are the two most likely
> choices, but more flags are defined in RFC 2060 2.3.2.

:dontexpunge

> If non-nil, don't remove all articles marked as deleted in the mailbox
> after finishing the fetch.

An example IMAP mail source:

```
(imap :server "mail.mycorp.com"
      :stream kerberos4
      :fetchflag "\\Seen")
```

webmail   Get mail from a webmail server, such as www.hotmail.com, web-
          mail.netscape.com, www.netaddress.com, www.my-deja.com.

NOTE: Now mail.yahoo.com provides POP3 service, so POP mail source is
suggested.

NOTE: Webmail largely depends cookies. A "one-line-cookie" patch is required
for url "4.0pre.46".

WARNING: Mails may lost. NO WARRANTY.

Keywords:

:subtype   The type of the webmail server. The default is `hotmail`. The
           alternatives are `netscape`, `netaddress`, `my-deja`.

:user   The user name to give to the webmail server. The default is the
           login name.

:password

> The password to give to the webmail server. If not specified, the
> user is prompted.

:dontexpunge

> If non-nil, only fetch unread articles and don't move them to trash
> folder after finishing the fetch.

An example webmail source:

```
(webmail :subtype 'hotmail
         :user "user-name"
         :password "secret")
```

Common Keywords

> Common keywords can be used in any type of mail source.

Keywords:

:plugged   If non-nil, fetch the mail even when Gnus is unplugged. If you use
           directory source to get mail, you can specify it as in this example:

```
(setq mail-sources
      '((directory :path "/home/pavel/.Spool/"
```

```
                                            :suffix ""
                                            :plugged t)))
```
Gnus will then fetch your mail even when you are unplugged. This is useful when you use local mail and news.

6.3.4.2 Function Interface

Some of the above keywords specify a Lisp function to be executed. For each keyword `:foo`, the Lisp variable `foo` is bound to the value of the keyword while the function is executing. For example, consider the following mail-source setting:

```
(setq mail-sources '((pop :user "jrl"
                          :server "pophost" :function fetchfunc)))
```

While the function `fetchfunc` is executing, the symbol `user` is bound to `"jrl"`, and the symbol `server` is bound to `"pophost"`. The symbols `port`, `password`, `program`, `prescript`, `postscript`, `function`, and `authentication` are also bound (to their default values).

See above for a list of keywords for each type of mail source.

6.3.4.3 Mail Source Customization

The following is a list of variables that influence how the mail is fetched. You would normally not need to set or change any of these variables.

`mail-source-crash-box`
> File where mail will be stored while processing it. The default is '`~/.emacs-mail-crash-box`'.

`mail-source-delete-incoming`
> If non-`nil`, delete incoming files after handling them.

`mail-source-directory`
> Directory where files (if any) will be stored. The default is '`~/Mail/`'. At present, the only thing this is used for is to say where the incoming files will be stored if the previous variable is `nil`.

`mail-source-incoming-file-prefix`
> Prefix for file name for storing incoming mail. The default is '`Incoming`', in which case files will end up with names like '`Incoming30630D_`' or '`Incoming298602ZD`'. This is really only relevant if `mail-source-delete-incoming` is `nil`.

`mail-source-default-file-modes`
> All new mail files will get this file mode. The default is 384.

6.3.4.4 Fetching Mail

The way to actually tell Gnus where to get new mail from is to set `mail-sources` to a list of mail source specifiers (see Section 6.3.4.1 [Mail Source Specifiers], page 117).

If this variable (and the obsolescent `nnmail-spool-file`) is `nil`, the mail back ends will never attempt to fetch mail by themselves.

If you want to fetch mail both from your local spool as well as a POP mail server, you'd say something like:

```
(setq mail-sources
      '((file)
        (pop :server "pop3.mail.server"
             :password "secret")))
```

Or, if you don't want to use any of the keyword defaults:

```
(setq mail-sources
      '((file :path "/var/spool/mail/user-name")
        (pop :server "pop3.mail.server"
             :user "user-name"
             :port "pop3"
             :password "secret")))
```

When you use a mail back end, Gnus will slurp all your mail from your inbox and plonk it down in your home directory. Gnus doesn't move any mail if you're not using a mail back end—you have to do a lot of magic invocations first. At the time when you have finished drawing the pentagram, lightened the candles, and sacrificed the goat, you really shouldn't be too surprised when Gnus moves your mail.

6.3.5 Mail Back End Variables

These variables are (for the most part) pertinent to all the various mail back ends.

`nnmail-read-incoming-hook`

> The mail back ends all call this hook after reading new mail. You can use this hook to notify any mail watch programs, if you want to.

`nnmail-split-hook`

> Hook run in the buffer where the mail headers of each message is kept just before the splitting based on these headers is done. The hook is free to modify the buffer contents in any way it sees fit—the buffer is discarded after the splitting has been done, and no changes performed in the buffer will show up in any files. `gnus-article-decode-encoded-words` is one likely function to add to this hook.

`nnmail-pre-get-new-mail-hook`
`nnmail-post-get-new-mail-hook`

> These are two useful hooks executed when treating new incoming mail—`nnmail-pre-get-new-mail-hook` (is called just before starting to handle the new mail) and `nnmail-post-get-new-mail-hook` (is called when the mail handling is done). Here's and example of using these two hooks to change the default file modes the new mail files get:
>
> ```
> (add-hook 'gnus-pre-get-new-mail-hook
>           (lambda () (set-default-file-modes 511)))
>
> (add-hook 'gnus-post-get-new-mail-hook
>           (lambda () (set-default-file-modes 551)))
> ```

`nnmail-use-long-file-names`
>   If non-`nil`, the mail back ends will use long file and directory names. Groups like 'mail.misc' will end up in directories (assuming use of `nnml` back end) or files (assuming use of `nnfolder` back end) like 'mail.misc'. If it is `nil`, the same group will end up in '`mail/misc`'.

`nnmail-delete-file-function`
>   Function called to delete files. It is `delete-file` by default.

`nnmail-cache-accepted-message-ids`
>   If non-`nil`, put the `Message-IDs` of articles imported into the back end (via `Gcc`, for instance) into the mail duplication discovery cache. The default is `nil`.

6.3.6 Fancy Mail Splitting

If the rather simple, standard method for specifying how to split mail doesn't allow you to do what you want, you can set `nnmail-split-methods` to `nnmail-split-fancy`. Then you can play with the `nnmail-split-fancy` variable.

Let's look at an example value of this variable first:

```
;; Messages from the mailer daemon are not crossposted to any of
;; the ordinary groups.  Warnings are put in a separate group
;; from real errors.
(| ("from" mail (| ("subject" "warn.*" "mail.warning")
                   "mail.misc"))
   ;; Non-error messages are crossposted to all relevant
   ;; groups, but we don't crosspost between the group for the
   ;; (ding) list and the group for other (ding) related mail.
   (& (| (any "ding@ifi\\.uio\\.no" "ding.list")
         ("subject" "ding" "ding.misc"))
      ;; Other mailing lists...
      (any "procmail@informatik\\.rwth-aachen\\.de" "procmail.list")
      (any "SmartList@informatik\\.rwth-aachen\\.de" "SmartList.list")
      ;; Both lists below have the same suffix, so prevent
      ;; cross-posting to mkpkg.list of messages posted only to
      ;; the bugs- list, but allow cross-posting when the
      ;; message was really cross-posted.
      (any "bugs-mypackage@somewhere" "mypkg.bugs")
      (any "mypackage@somewhere\" - "bugs-mypackage" "mypkg.list")
      ;; People...
      (any "larsi@ifi\\.uio\\.no" "people.Lars_Magne_Ingebrigtsen"))
   ;; Unmatched mail goes to the catch all group.
   "misc.misc")
```

This variable has the format of a *split*. A split is a (possibly) recursive structure where each split may contain other splits. Here are the five possible split syntaxes:

1. '`group`': If the split is a string, that will be taken as a group name. Normal regexp match expansion will be done. See below for examples.

2. (*field value* [- *restrict* [...]] *split*): If the split is a list, the first element of which is a string, then store the message as specified by *split*, if header *field* (a regexp) contains

value (also a regexp). If *restrict* (yet another regexp) matches some string after *field* and before the end of the matched *value*, the *split* is ignored. If none of the *restrict* clauses match, *split* is processed.

3. (| *split*...): If the split is a list, and the first element is | (vertical bar), then process each *split* until one of them matches. A *split* is said to match if it will cause the mail message to be stored in one or more groups.

4. (& *split*...): If the split is a list, and the first element is &, then process all *split*s in the list.

5. junk: If the split is the symbol junk, then don't save this message. Use with extreme caution.

6. (: *function arg1 arg2* ...): If the split is a list, and the first element is :, then the second element will be called as a function with *args* given as arguments. The function should return a *split*.

   For instance, the following function could be used to split based on the body of the messages:

   ```
   (defun split-on-body ()
     (save-excursion
       (set-buffer " *nnmail incoming*")
       (goto-char (point-min))
       (when (re-search-forward "Some.*string" nil t)
         "string.group")))
   ```

7. (! *func split*): If the split is a list, and the first element is !, then SPLIT will be processed, and FUNC will be called as a function with the result of SPLIT as argument. FUNC should return a split.

8. nil: If the split is nil, it is ignored.

In these splits, *field* must match a complete field name. *value* must match a complete word according to the fundamental mode syntax table. You can use .* in the regexps to match partial field names or words. In other words, all *value*'s are wrapped in '\<' and '\>' pairs.

field and *value* can also be lisp symbols, in that case they are expanded as specified by the variable nnmail-split-abbrev-alist. This is an alist of cons cells, where the car of a cell contains the key, and the cdr contains the associated value.

nnmail-split-fancy-syntax-table is the syntax table in effect when all this splitting is performed.

If you want to have Gnus create groups dynamically based on some information in the headers (i.e., do replace-match-like substitutions in the group names), you can say things like:

```
(any "debian-\\b\\(\\w+\\)@lists.debian.org" "mail.debian.\\1")
```

In this example, messages sent to 'debian-foo@lists.debian.org' will be filed in 'mail.debian.foo'.

If the string contains the element '\&', then the previously matched string will be substituted. Similarly, the elements '\\1' up to '\\9' will be substituted with the text matched by the groupings 1 through 9.

nnmail-split-fancy-with-parent is a function which allows you to split followups into the same groups their parents are in. Sometimes you can't make splitting rules for all your mail. For example, your boss might send you personal mail regarding different projects you are working on, and as you can't tell your boss to put a distinguishing string into the subject line, you have to resort to manually moving the messages into the right group. With this function, you only have to do it once per thread.

To use this feature, you have to set nnmail-treat-duplicates to a non-nil value. And then you can include nnmail-split-fancy-with-parent using the colon feature, like so:

```
(setq nnmail-split-fancy
      '(| (: nnmail-split-fancy-with-parent)
          ;; other splits go here
        ))
```

This feature works as follows: when nnmail-treat-duplicates is non-nil, Gnus records the message id of every message it sees in the file specified by the variable nnmail-message-id-cache-file, together with the group it is in (the group is omitted for non-mail messages). When mail splitting is invoked, the function nnmail-split-fancy-with-parent then looks at the References (and In-Reply-To) header of each message to split and searches the file specified by nnmail-message-id-cache-file for the message ids. When it has found a parent, it returns the corresponding group name. It is recommended that you set nnmail-message-id-cache-length to a somewhat higher number than the default so that the message ids are still in the cache. (A value of 5000 appears to create a file some 300 kBytes in size.) When nnmail-cache-accepted-message-ids is non-nil, Gnus also records the message ids of moved articles, so that the followup messages goes into the new group.

6.3.7 Group Mail Splitting

If you subscribe to dozens of mailing lists but you don't want to maintain mail splitting rules manually, group mail splitting is for you. You just have to set *to-list* and/or *to-address* in group parameters or group customization and set nnmail-split-methods to gnus-group-split. This splitting function will scan all groups for those parameters and split mail accordingly, i.e., messages posted from or to the addresses specified in the parameters *to-list* or *to-address* of a mail group will be stored in that group.

Sometimes, mailing lists have multiple addresses, and you may want mail splitting to recognize them all: just set the *extra-aliases* group parameter to the list of additional addresses and it's done. If you'd rather use a regular expression, set *split-regexp*.

All these parameters in a group will be used to create an nnmail-split-fancy split, in which the *field* is 'any', the *value* is a single regular expression that matches *to-list*, *to-address*, all of *extra-aliases* and all matches of *split-regexp*, and the *split* is the name of the group. *restrict*s are also supported: just set the *split-exclude* parameter to a list of regular expressions.

If you can't get the right split to be generated using all these parameters, or you just need something fancier, you can set the parameter *split-spec* to an nnmail-split-fancy split. In this case, all other aforementioned parameters will be ignored by gnus-group-split. In particular, *split-spec* may be set to nil, in which case the group will be ignored by gnus-group-split.

gnus-group-split will do cross-posting on all groups that match, by defining a single & fancy split containing one split for each group. If a message doesn't match any split, it will be stored in the group named in gnus-group-split-default-catch-all-group, unless some group has *split-spec* set to catch-all, in which case that group is used as the catch-all group. Even though this variable is often used just to name a group, it may also be set to an arbitrarily complex fancy split (after all, a group name is a fancy split), and this may be useful to split mail that doesn't go to any mailing list to personal mail folders. Note that this fancy split is added as the last element of a | split list that also contains a & split with the rules extracted from group parameters.

It's time for an example. Assume the following group parameters have been defined:

```
nnml:mail.bar:
((to-address . "bar@femail.com")
 (split-regexp . ".*@femail\\.com"))
nnml:mail.foo:
((to-list . "foo@nowhere.gov")
 (extra-aliases "foo@localhost" "foo-redist@home")
 (split-exclude "bugs-foo" "rambling-foo")
 (admin-address . "foo-request@nowhere.gov"))
nnml:mail.others:
((split-spec . catch-all))
```

Setting nnmail-split-methods to gnus-group-split will behave as if nnmail-split-fancy had been selected and variable nnmail-split-fancy had been set as follows:

```
(| (& (any "\\(bar@femail\\.com\\|.*@femail\\.com\\)" "mail.bar")
      (any "\\(foo@nowhere\\.gov\\|foo@localhost\\|foo-redist@home\\)"
           - "bugs-foo" - "rambling-foo" "mail.foo"))
   "mail.others")
```

If you'd rather not use group splitting for all your mail groups, you may use it for only some of them, by using nnmail-split-fancy splits like this:

```
(: gnus-mlsplt-fancy GROUPS NO-CROSSPOST CATCH-ALL)
```

groups may be a regular expression or a list of group names whose parameters will be scanned to generate the output split. *no-crosspost* can be used to disable cross-posting; in this case, a single | split will be output. *catch-all* is the fallback fancy split, used like *gnus-group-split-default-catch-all-group*. If *catch-all* is nil, or if *split-regexp* matches the empty string in any selected group, no catch-all split will be issued. Otherwise, if some group has *split-spec* set to catch-all, this group will override the value of the *catch-all* argument.

Unfortunately, scanning all groups and their parameters can be quite slow, especially considering that it has to be done for every message. But don't despair! The function gnus-group-split-setup can be used to enable gnus-group-split in a much more efficient way. It sets nnmail-split-methods to nnmail-split-fancy and sets nnmail-split-fancy to the split produced by gnus-group-split-fancy. Thus, the group parameters are only scanned once, no matter how many messages are split.

However, if you change group parameters, you have to update nnmail-split-fancy manually. You can do it by running gnus-group-split-update. If you'd rather have it updated automatically, just tell gnus-group-split-setup to do it for you. For example, add to your '.gnus':

```
(gnus-group-split-setup AUTO-UPDATE CATCH-ALL)
```

If *auto-update* is non-`nil`, `gnus-group-split-update` will be added to `nnmail-pre-get-new-mail-hook`, so you won't ever have to worry about updating `nnmail-split-fancy` again. If you don't omit *catch-all* (it's optional, equivalent to `nil`), `gnus-group-split-default-catch-all-group` will be set to its value.

Because you may want to change `nnmail-split-fancy` after it is set by `gnus-group-split-update`, this function will run `gnus-group-split-updated-hook` just before finishing.

6.3.8 Incorporating Old Mail

Most people have lots of old mail stored in various file formats. If you have set up Gnus to read mail using one of the spiffy Gnus mail back ends, you'll probably wish to have that old mail incorporated into your mail groups.

Doing so can be quite easy.

To take an example: You're reading mail using `nnml` (see Section 6.3.13.3 [Mail Spool], page 133), and have set `nnmail-split-methods` to a satisfactory value (see Section 6.3.3 [Splitting Mail], page 115). You have an old Unix mbox file filled with important, but old, mail. You want to move it into your `nnml` groups.

Here's how:

1. Go to the group buffer.
2. Type *G f* and give the path to the mbox file when prompted to create an `nndoc` group from the mbox file (see Section 2.9 [Foreign Groups], page 21).
3. Type ⟨SPC⟩ to enter the newly created group.
4. Type *M P b* to process-mark all articles in this group's buffer (see Section 3.6.6 [Setting Process Marks], page 51).
5. Type *B r* to respool all the process-marked articles, and answer 'nnml' when prompted (see Section 3.24 [Mail Group Commands], page 86).

All the mail messages in the mbox file will now also be spread out over all your `nnml` groups. Try entering them and check whether things have gone without a glitch. If things look ok, you may consider deleting the mbox file, but I wouldn't do that unless I was absolutely sure that all the mail has ended up where it should be.

Respooling is also a handy thing to do if you're switching from one mail back end to another. Just respool all the mail in the old mail groups using the new mail back end.

6.3.9 Expiring Mail

Traditional mail readers have a tendency to remove mail articles when you mark them as read, in some way. Gnus takes a fundamentally different approach to mail reading.

Gnus basically considers mail just to be news that has been received in a rather peculiar manner. It does not think that it has the power to actually change the mail, or delete any mail messages. If you enter a mail group, and mark articles as "read", or kill them in some other fashion, the mail articles will still exist on the system. I repeat: Gnus will not delete your old, read mail. Unless you ask it to, of course.

To make Gnus get rid of your unwanted mail, you have to mark the articles as *expirable*. This does not mean that the articles will disappear right away, however. In general, a mail article will be deleted from your system if, 1) it is marked as expirable, AND 2) it is more than one week old. If you do not mark an article as expirable, it will remain on your system until hell freezes over. This bears repeating one more time, with some spurious capitalizations: IF you do NOT mark articles as EXPIRABLE, Gnus will NEVER delete those ARTICLES.

You do not have to mark articles as expirable by hand. Groups that match the regular expression `gnus-auto-expirable-newsgroups` will have all articles that you read marked as expirable automatically. All articles marked as expirable have an 'E' in the first column in the summary buffer.

By default, if you have auto expiry switched on, Gnus will mark all the articles you read as expirable, no matter if they were read or unread before. To avoid having articles marked as read marked as expirable automatically, you can put something like the following in your '`.gnus`' file:

```
(remove-hook 'gnus-mark-article-hook
             'gnus-summary-mark-read-and-unread-as-read)
(add-hook 'gnus-mark-article-hook 'gnus-summary-mark-unread-as-read)
```

Note that making a group auto-expirable doesn't mean that all read articles are expired— only the articles marked as expirable will be expired. Also note that using the *d* command won't make groups expirable—only semi-automatic marking of articles as read will mark the articles as expirable in auto-expirable groups.

Let's say you subscribe to a couple of mailing lists, and you want the articles you have read to disappear after a while:

```
(setq gnus-auto-expirable-newsgroups
      "mail.nonsense-list\\|mail.nice-list")
```

Another way to have auto-expiry happen is to have the element `auto-expire` in the group parameters of the group.

If you use adaptive scoring (see Section 7.6 [Adaptive Scoring], page 176) and auto-expiring, you'll have problems. Auto-expiring and adaptive scoring don't really mix very well.

The `nnmail-expiry-wait` variable supplies the default time an expirable article has to live. Gnus starts counting days from when the message *arrived*, not from when it was sent. The default is seven days.

Gnus also supplies a function that lets you fine-tune how long articles are to live, based on what group they are in. Let's say you want to have one month expiry period in the '`mail.private`' group, a one day expiry period in the '`mail.junk`' group, and a six day expiry period everywhere else:

```
(setq nnmail-expiry-wait-function
      (lambda (group)
       (cond ((string= group "mail.private")
              31)
             ((string= group "mail.junk")
              1)
             ((string= group "important")
```

```
            'never)
        (t
            6))))
```

The group names this function is fed are "unadorned" group names—no 'nnml:' prefixes and the like.

The nnmail-expiry-wait variable and nnmail-expiry-wait-function function can either be a number (not necessarily an integer) or one of the symbols immediate or never.

You can also use the expiry-wait group parameter to selectively change the expiry period (see Section 2.10 [Group Parameters], page 22).

The normal action taken when expiring articles is to delete them. However, in some circumstances it might make more sense to move them to other groups instead of deleting them. The variable nnmail-expiry-target (and the expiry-target group parameter) controls this. The variable supplies a default value for all groups, which can be overridden for specific groups by the group parameter. default value is delete, but this can also be a string (which should be the name of the group the message should be moved to), or a function (which will be called in a buffer narrowed to the message in question, and with the name of the group being moved from as its parameter) which should return a target – either a group name or delete.

Here's an example for specifying a group name:

```
(setq nnmail-expiry-target "nnml:expired")
```

If nnmail-keep-last-article is non-nil, Gnus will never expire the final article in a mail newsgroup. This is to make life easier for procmail users.

By the way: That line up there, about Gnus never expiring non-expirable articles, is a lie. If you put total-expire in the group parameters, articles will not be marked as expirable, but all read articles will be put through the expiry process. Use with extreme caution. Even more dangerous is the gnus-total-expirable-newsgroups variable. All groups that match this regexp will have all read articles put through the expiry process, which means that *all* old mail articles in the groups in question will be deleted after a while. Use with extreme caution, and don't come crying to me when you discover that the regexp you used matched the wrong group and all your important mail has disappeared. Be a *man*! Or a *woman*! Whatever you feel more comfortable with! So there!

Most people make most of their mail groups total-expirable, though.

If gnus-inhibit-user-auto-expire is non-nil, user marking commands will not mark an article as expirable, even if the group has auto-expire turned on.

6.3.10 Washing Mail

Mailers and list servers are notorious for doing all sorts of really, really stupid things with mail. "Hey, RFC 822 doesn't explicitly prohibit us from adding the string wE aRe ElItE!!!!!1!! to the end of all lines passing through our server, so let's do that!!!!1!" Yes, but RFC 822 wasn't designed to be read by morons. Things that were considered to be self-evident were not discussed. So. Here we are.

Case in point: The German version of Microsoft Exchange adds 'AW: ' to the subjects of replies instead of 'Re: '. I could pretend to be shocked and dismayed by this, but I haven't got the energy. It is to laugh.

Gnus provides a plethora of functions for washing articles while displaying them, but it might be nicer to do the filtering before storing the mail to disc. For that purpose, we have three hooks and various functions that can be put in these hooks.

`nnmail-prepare-incoming-hook`

> This hook is called before doing anything with the mail and is meant for grand, sweeping gestures. It is called in a buffer that contains all the new, incoming mail. Functions to be used include:

> `nnheader-ms-strip-cr`

>> Remove trailing carriage returns from each line. This is default on Emacs running on MS machines.

`nnmail-prepare-incoming-header-hook`

> This hook is called narrowed to each header. It can be used when cleaning up the headers. Functions that can be used include:

> `nnmail-remove-leading-whitespace`

>> Clear leading white space that "helpful" listservs have added to the headers to make them look nice. Aaah.

> `nnmail-remove-list-identifiers`

>> Some list servers add an identifier—for example, '`(idm)`'—to the beginning of all `Subject` headers. I'm sure that's nice for people who use stone age mail readers. This function will remove strings that match the `nnmail-list-identifiers` regexp, which can also be a list of regexp. `nnmail-list-identifiers` may not contain `\\(..\\)`.

>> For instance, if you want to remove the '`(idm)`' and the '`nagnagnag`' identifiers:

>>```
(setq nnmail-list-identifiers
      '("(idm)" "nagnagnag"))
```

>> This can also be done non-destructively with `gnus-list-identifiers`, See Section 3.16.3 [Article Hiding], page 71.

> `nnmail-remove-tabs`

>> Translate all tab characters into space characters.

> `nnmail-fix-eudora-headers`

>> Eudora produces broken `References` headers, but OK `In-Reply-To` headers. This function will get rid of the `References` headers.

`nnmail-prepare-incoming-message-hook`

> This hook is called narrowed to each message. Functions to be used include:

> `article-de-quoted-unreadable`

>> Decode Quoted Readable encoding.

## 6.3.11 Duplicates

If you are a member of a couple of mailing lists, you will sometimes receive two copies of the same mail. This can be quite annoying, so `nnmail` checks for and treats any duplicates

it might find. To do this, it keeps a cache of old `Message-IDs`— `nnmail-message-id-cache-file`, which is '`~/.nnmail-cache`' by default. The approximate maximum number of `Message-IDs` stored there is controlled by the `nnmail-message-id-cache-length` variable, which is 1000 by default. (So 1000 `Message-IDs` will be stored.) If all this sounds scary to you, you can set `nnmail-treat-duplicates` to `warn` (which is what it is by default), and `nnmail` won't delete duplicate mails. Instead it will insert a warning into the head of the mail saying that it thinks that this is a duplicate of a different message.

This variable can also be a function. If that's the case, the function will be called from a buffer narrowed to the message in question with the `Message-ID` as a parameter. The function must return either `nil`, `warn`, or `delete`.

You can turn this feature off completely by setting the variable to `nil`.

If you want all the duplicate mails to be put into a special *duplicates* group, you could do that using the normal mail split methods:

```
(setq nnmail-split-fancy
      '(| ;; Messages duplicates go to a separate group.
          ("gnus-warning" "duplication of message" "duplicate")
          ;; Message from daemons, postmaster, and the like to another.
          (any mail "mail.misc")
          ;; Other rules.
          [ ... ] ))
```

Or something like:

```
(setq nnmail-split-methods
      '(("duplicates" "^Gnus-Warning:")
        ;; Other rules.
        [...]))
```

Here's a neat feature: If you know that the recipient reads her mail with Gnus, and that she has `nnmail-treat-duplicates` set to `delete`, you can send her as many insults as you like, just by using a `Message-ID` of a mail that you know that she's already received. Think of all the fun! She'll never see any of it! Whee!

6.3.12 Not Reading Mail

If you start using any of the mail back ends, they have the annoying habit of assuming that you want to read mail with them. This might not be unreasonable, but it might not be what you want.

If you set `mail-sources` and `nnmail-spool-file` to `nil`, none of the back ends will ever attempt to read incoming mail, which should help.

This might be too much, if, for instance, you are reading mail quite happily with `nnml` and just want to peek at some old RMAIL file you have stashed away with `nnbabyl`. All back ends have variables called back-end-`get-new-mail`. If you want to disable the `nnbabyl` mail reading, you edit the virtual server for the group to have a setting where `nnbabyl-get-new-mail` to `nil`.

All the mail back ends will call `nn*-prepare-save-mail-hook` narrowed to the article to be saved before saving it when reading incoming mail.

6.3.13 Choosing a Mail Back End

Gnus will read the mail spool when you activate a mail group. The mail file is first copied to your home directory. What happens after that depends on what format you want to store your mail in.

There are five different mail back ends in the standard Gnus, and more back ends are available separately. The mail back end most people use (because it is the fastest and most flexible) is `nnml` (see Section 6.3.13.3 [Mail Spool], page 133).

6.3.13.1 Unix Mail Box

The *nnmbox* back end will use the standard Un*x mbox file to store mail. `nnmbox` will add extra headers to each mail article to say which group it belongs in.

Virtual server settings:

`nnmbox-mbox-file`
> The name of the mail box in the user's home directory.

`nnmbox-active-file`
> The name of the active file for the mail box.

`nnmbox-get-new-mail`
> If non-`nil`, `nnmbox` will read incoming mail and split it into groups.

6.3.13.2 Rmail Babyl

The *nnbabyl* back end will use a babyl mail box (aka. *rmail mbox*) to store mail. `nnbabyl` will add extra headers to each mail article to say which group it belongs in.

Virtual server settings:

`nnbabyl-mbox-file`
> The name of the rmail mbox file.

`nnbabyl-active-file`
> The name of the active file for the rmail box.

`nnbabyl-get-new-mail`
> If non-`nil`, `nnbabyl` will read incoming mail.

6.3.13.3 Mail Spool

The *nnml* spool mail format isn't compatible with any other known format. It should be used with some caution.

If you use this back end, Gnus will split all incoming mail into files, one file for each mail, and put the articles into the corresponding directories under the directory specified by the `nnml-directory` variable. The default value is '`~/Mail/`'.

You do not have to create any directories beforehand; Gnus will take care of all that.

If you have a strict limit as to how many files you are allowed to store in your account, you should not use this back end. As each mail gets its own file, you might very well occupy thousands of inodes within a few weeks. If this is no problem for you, and it isn't a problem

for you having your friendly systems administrator walking around, madly, shouting "Who is eating all my inodes?! Who? Who!?!", then you should know that this is probably the fastest format to use. You do not have to trudge through a big mbox file just to read your new mail.

`nnml` is probably the slowest back end when it comes to article splitting. It has to create lots of files, and it also generates NOV databases for the incoming mails. This makes it the fastest back end when it comes to reading mail.

Virtual server settings:

`nnml-directory`
> All `nnml` directories will be placed under this directory.

`nnml-active-file`
> The active file for the `nnml` server.

`nnml-newsgroups-file`
> The `nnml` group descriptions file. See Section 10.7.9.2 [Newsgroups File Format], page 253.

`nnml-get-new-mail`
> If non-`nil`, `nnml` will read incoming mail.

`nnml-nov-is-evil`
> If non-`nil`, this back end will ignore any NOV files.

`nnml-nov-file-name`
> The name of the NOV files. The default is '`.overview`'.

`nnml-prepare-save-mail-hook`
> Hook run narrowed to an article before saving.

If your `nnml` groups and NOV files get totally out of whack, you can do a complete update by typing *M-x nnml-generate-nov-databases*. This command will trawl through the entire `nnml` hierarchy, looking at each and every article, so it might take a while to complete. A better interface to this functionality can be found in the server buffer (see Section 6.1.2 [Server Commands], page 106).

6.3.13.4 MH Spool

`nnmh` is just like `nnml`, except that is doesn't generate NOV databases and it doesn't keep an active file. This makes `nnmh` a *much* slower back end than `nnml`, but it also makes it easier to write procmail scripts for.

Virtual server settings:

`nnmh-directory`
> All `nnmh` directories will be located under this directory.

`nnmh-get-new-mail`
> If non-`nil`, `nnmh` will read incoming mail.

`nnmh-be-safe`
> If non-`nil`, `nnmh` will go to ridiculous lengths to make sure that the articles in the folder are actually what Gnus thinks they are. It will check date stamps

and stat everything in sight, so setting this to `t` will mean a serious slow-down. If you never use anything but Gnus to read the `nnmh` articles, you do not have to set this variable to `t`.

6.3.13.5 Mail Folders

`nnfolder` is a back end for storing each mail group in a separate file. Each file is in the standard Un*x mbox format. `nnfolder` will add extra headers to keep track of article numbers and arrival dates.

Virtual server settings:

`nnfolder-directory`

>   All the `nnfolder` mail boxes will be stored under this directory.

`nnfolder-active-file`

>   The name of the active file.

`nnfolder-newsgroups-file`

>   The name of the group descriptions file. See Section 10.7.9.2 [Newsgroups File Format], page 253.

`nnfolder-get-new-mail`

>   If non-`nil`, `nnfolder` will read incoming mail.

`nnfolder-save-buffer-hook`

>   Hook run before saving the folders. Note that Emacs does the normal backup renaming of files even with the `nnfolder` buffers. If you wish to switch this off, you could say something like the following in your '`.emacs`' file:

```
(defun turn-off-backup ()
  (set (make-local-variable 'backup-inhibited) t))

(add-hook 'nnfolder-save-buffer-hook 'turn-off-backup)
```

`nnfolder-delete-mail-hook`

>   Hook run in a buffer narrowed to the message that is to be deleted. This function can be used to copy the message to somewhere else, or to extract some information from it before removing it.

If you have lots of `nnfolder`-like files you'd like to read with `nnfolder`, you can use the `M-x nnfolder-generate-active-file` command to make `nnfolder` aware of all likely files in `nnfolder-directory`. This only works if you use long file names, though.

6.3.13.6 Comparing Mail Back Ends

First, just for terminology, the *back end* is the common word for a low-level access method—a transport, if you will, by which something is acquired. The sense is that one's mail has to come from somewhere, and so selection of a suitable back end is required in order to get that mail within spitting distance of Gnus.

The same concept exists for Usenet itself: Though access to articles is typically done by NNTP these days, once upon a midnight dreary, everyone in the world got at Usenet by running a reader on the machine where the articles lay (the machine which today we call

an NNTP server), and access was by the reader stepping into the articles' directory spool area directly. One can still select between either the `nntp` or `nnspool` back ends, to select between these methods, if one happens actually to live on the server (or can see its spool directly, anyway, via NFS).

The goal in selecting a mail back end is to pick one which simultaneously represents a suitable way of dealing with the original format plus leaving mail in a form that is convenient to use in the future. Here are some high and low points on each:

`nnmbox`

UNIX systems have historically had a single, very common, and well- defined format. All messages arrive in a single *spool file*, and they are delineated by a line whose regular expression matches '`^From_`'. (My notational use of '`_`' is to indicate a space, to make it clear in this instance that this is not the RFC-specified '`From:`' header.) Because Emacs and therefore Gnus emanate historically from the Unix environment, it is simplest if one does not mess a great deal with the original mailbox format, so if one chooses this back end, Gnus' primary activity in getting mail from the real spool area to Gnus' preferred directory is simply to copy it, with no (appreciable) format change in the process. It is the "dumbest" way to move mail into availability in the Gnus environment. This makes it fast to move into place, but slow to parse, when Gnus has to look at what's where.

`nnbabyl`

Once upon a time, there was the DEC-10 and DEC-20, running operating systems called TOPS and related things, and the usual (only?) mail reading environment was a thing called Babyl. I don't know what format was used for mail landing on the system, but Babyl had its own internal format to which mail was converted, primarily involving creating a spool-file-like entity with a scheme for inserting Babyl-specific headers and status bits above the top of each message in the file. RMAIL was Emacs' first mail reader, it was written by Richard Stallman, and Stallman came out of that TOPS/Babyl environment, so he wrote RMAIL to understand the mail files folks already had in existence. Gnus (and VM, for that matter) continue to support this format because it's perceived as having some good qualities in those mailer-specific headers/status bits stuff. RMAIL itself still exists as well, of course, and is still maintained by Stallman.

Both of the above forms leave your mail in a single file on your filesystem, and they must parse that entire file each time you take a look at your mail.

`nnml`

`nnml` is the back end which smells the most as though you were actually operating with an `nnspool`-accessed Usenet system. (In fact, I believe `nnml` actually derived from `nnspool` code, lo these years ago.) One's mail is taken from the original spool file, and is then cut up into individual message files, 1:1. It maintains a Usenet-style active file (analogous to what one finds in an INN- or CNews-based news system in (for instance) '`/var/lib/news/active`', or what is returned via the '`NNTP LIST`' verb) and also creates *overview* files for efficient group entry, as has been defined for NNTP servers for some years now. It

is slower in mail-splitting, due to the creation of lots of files, updates to the `nnml` active file, and additions to overview files on a per-message basis, but it is extremely fast on access because of what amounts to the indexing support provided by the active file and overviews.

`nnml` costs *inodes* in a big way; that is, it soaks up the resource which defines available places in the filesystem to put new files. Sysadmins take a dim view of heavy inode occupation within tight, shared filesystems. But if you live on a personal machine where the filesystem is your own and space is not at a premium, `nnml` wins big.

It is also problematic using this back end if you are living in a FAT16-based Windows world, since much space will be wasted on all these tiny files.

`nnmh`

The Rand MH mail-reading system has been around UNIX systems for a very long time; it operates by splitting one's spool file of messages into individual files, but with little or no indexing support – `nnmh` is considered to be semantically equivalent to "`nnml` without active file or overviews". This is arguably the worst choice, because one gets the slowness of individual file creation married to the slowness of access parsing when learning what's new in one's groups.

`nnfolder`

Basically the effect of `nnfolder` is `nnmbox` (the first method described above) on a per-group basis. That is, `nnmbox` itself puts *all* one's mail in one file; `nnfolder` provides a little bit of optimization to this so that each of one's mail groups has a Unix mail box file. It's faster than `nnmbox` because each group can be parsed separately, and still provides the simple Unix mail box format requiring minimal effort in moving the mail around. In addition, it maintains an "active" file making it much faster for Gnus to figure out how many messages there are in each separate group.

If you have groups that are expected to have a massive amount of messages, `nnfolder` is not the best choice, but if you receive only a moderate amount of mail, `nnfolder` is probably the most friendly mail back end all over.

6.4 Browsing the Web

Web-based discussion forums are getting more and more popular. On many subjects, the web-based forums have become the most important forums, eclipsing the importance of mailing lists and news groups. The reason is easy to understand—they are friendly to new users; you just point and click, and there's the discussion. With mailing lists, you have to go through a cumbersome subscription procedure, and most people don't even know what a news group is.

The problem with this scenario is that web browsers are not very good at being newsreaders. They do not keep track of what articles you've read; they do not allow you to score on subjects you're interested in; they do not allow off-line browsing; they require you to click around and drive you mad in the end.

So—if web browsers suck at reading discussion forums, why not use Gnus to do it instead?

Gnus has been getting a bit of a collection of back ends for providing interfaces to these sources.

All the web sources require Emacs/w3 and the url library to work.

The main caveat with all these web sources is that they probably won't work for a very long time. Gleaning information from the HTML data is guesswork at best, and when the layout is altered, the Gnus back end will fail. If you have reasonably new versions of these back ends, though, you should be ok.

One thing all these Web methods have in common is that the Web sources are often down, unavailable or just plain too slow to be fun. In those cases, it makes a lot of sense to let the Gnus Agent (see Section 6.7 [Gnus Unplugged], page 156) handle downloading articles, and then you can read them at leisure from your local disk. No more World Wide Wait for you.

6.4.1 Web Searches

It's, like, too neat to search the Usenet for articles that match a string, but it, like, totally *sucks*, like, totally, to use one of those, like, Web browsers, and you, like, have to, rilly, like, look at the commercials, so, like, with Gnus you can do *rad*, rilly, searches without having to use a browser.

The `nnweb` back end allows an easy interface to the mighty search engine. You create an `nnweb` group, enter a search pattern, and then enter the group and read the articles like you would any normal group. The `G w` command in the group buffer (see Section 2.9 [Foreign Groups], page 21) will do this in an easy-to-use fashion.

`nnweb` groups don't really lend themselves to being solid groups—they have a very fleeting idea of article numbers. In fact, each time you enter an `nnweb` group (not even changing the search pattern), you are likely to get the articles ordered in a different manner. Not even using duplicate suppression (see Section 3.28 [Duplicate Suppression], page 91) will help, since `nnweb` doesn't even know the `Message-ID` of the articles before reading them using some search engines (DejaNews, for instance). The only possible way to keep track of which articles you've read is by scoring on the `Date` header—mark all articles posted before the last date you read the group as read.

If the search engine changes its output substantially, `nnweb` won't be able to parse it and will fail. One could hardly fault the Web providers if they were to do this—their *raison d'tre* is to make money off of advertisements, not to provide services to the community. Since `nnweb` washes the ads off all the articles, one might think that the providers might be somewhat miffed. We'll see.

You must have the `url` and `w3` package installed to be able to use `nnweb`.

Virtual server variables:

`nnweb-type`

What search engine type is being used. The currently supported types are `dejanews`, `dejanewsold`, `altavista` and `reference`.

`nnweb-search`

The search string to feed to the search engine.

`nnweb-max-hits`

Advisory maximum number of hits per search to display. The default is 100.

`nnweb-type-definition`

Type-to-definition alist. This alist says what `nnweb` should do with the various search engine types. The following elements must be present:

`article`   Function to decode the article and provide something that Gnus understands.

`map`   Function to create an article number to message header and URL alist.

`search`   Function to send the search string to the search engine.

`address`   The address the aforementioned function should send the search string to.

`id`   Format string URL to fetch an article by `Message-ID`.

6.4.2 Slashdot

Slashdot (`http://slashdot.org/`) is a popular news site, with lively discussion following the news articles. `nnslashdot` will let you read this forum in a convenient manner.

The easiest way to read this source is to put something like the following in your '`.gnus.el`' file:

```
(setq gnus-secondary-select-methods
      '((nnslashdot "")))
```

This will make Gnus query the `nnslashdot` back end for new comments and groups. The `F` command will subscribe each new news article as a new Gnus group, and you can read the comments by entering these groups. (Note that the default subscription method is to subscribe new groups as zombies. Other methods are available (see Section 1.6.2 [Subscription Methods], page 6).

If you want to remove an old `nnslashdot` group, the `G` $\overline{\text{DEL}}$ command is the most handy tool (see Section 2.9 [Foreign Groups], page 21).

When following up to `nnslashdot` comments (or posting new comments), some light HTMLizations will be performed. In particular, text quoted with '`>` ' will be quoted with `blockquote` instead, and signatures will have `br` added to the end of each line. Other than that, you can just write HTML directly into the message buffer. Note that Slashdot filters out some HTML forms.

The following variables can be altered to change its behavior:

`nnslashdot-threaded`

Whether `nnslashdot` should display threaded groups or not. The default is `t`. To be able to display threads, `nnslashdot` has to retrieve absolutely all comments in a group upon entry. If a threaded display is not required, `nnslashdot` will only retrieve the comments that are actually wanted by the user. Threading is nicer, but much, much slower than untreaded.

`nnslashdot-login-name`

The login name to use when posting.

`nnslashdot-password`

The password to use when posting.

`nnslashdot-directory`

        Where `nnslashdot` will store its files. The default is '`~/News/slashdot/`'.

`nnslashdot-active-url`

        The URL format string that will be used to fetch the information on news articles and comments. Default: '`http://slashdot.org/search.pl?section=&min=%d`'.

`nnslashdot-comments-url`

        The URL format string that will be used to fetch comments. The default is '`http://slashdot.org/comments.pl?sid=%s&threshold=%d&commentsort=%d&mode=flat&start`'

`nnslashdot-article-url`

        The URL format string that will be used to fetch the news article. The default is '`http://slashdot.org/article.pl?sid=%s&mode=nocomment`'.

`nnslashdot-threshold`

        The score threshold. The default is -1.

`nnslashdot-group-number`

        The number of old groups, in addition to the ten latest, to keep updated. The default is 0.

6.4.3 Ultimate

The Ultimate Bulletin Board (`http://www.ultimatebb.com/`) is probably the most popular Web bulletin board system used. It has a quite regular and nice interface, and it's possible to get the information Gnus needs to keep groups updated.

The easiest way to get started with `nnultimate` is to say something like the following in the group buffer: *B nnultimate* (RET) *http://www.tcj.com/messboard/ubbcgi/* (RET). (Substitute the URL (not including '`Ultimate.cgi`' or the like at the end) for a forum you're interested in; there's quite a list of them on the Ultimate web site.) Then subscribe to the groups you're interested in from the server buffer, and read them from the group buffer.

The following `nnultimate` variables can be altered:

`nnultimate-directory`

        The directory where `nnultimate` stores its files. The default is '`~/News/ultimate/`'.

6.4.4 Web Archive

Some mailing lists only have archives on Web servers, such as `http://www.egroups.com/` and `http://www.mail-archive.com/`. It has a quite regular and nice interface, and it's possible to get the information Gnus needs to keep groups updated.

The easiest way to get started with `nnwarchive` is to say something like the following in the group buffer: *M-x gnus-group-make-warchive-group* (RET) *an_egroup* (RET) *egroups* (RET) *www.egroups.com* (RET) *your@email.address* (RET). (Substitute the AN_EGROUP with the mailing list you subscribed, the YOUR@EMAIL.ADDRESS with your email address.), or to browse the back end by *B nnwarchive* (RET) *mail-archive* (RET).

The following `nnwarchive` variables can be altered:

`nnwarchive-directory`

> The directory where `nnwarchive` stores its files. The default is
> '`~/News/warchive/`'.

`nnwarchive-login`

> The account name on the web server.

`nnwarchive-passwd`

> The password for your account on the web server.

6.4.5 Customizing w3

Gnus uses the url library to fetch web pages and Emacs/w3 to display web pages. Emacs/w3 is documented in its own manual, but there are some things that may be more relevant for Gnus users.

For instance, a common question is how to make Emacs/w3 follow links using the `browse-url` functions (which will call some external web browser like Netscape). Here's one way:

```
(eval-after-load "w3"
  '(progn
    (fset 'w3-fetch-orig (symbol-function 'w3-fetch))
    (defun w3-fetch (&optional url target)
      (interactive (list (w3-read-url-with-default)))
      (if (eq major-mode 'gnus-article-mode)
          (browse-url url)
        (w3-fetch-orig url target)))))
```

Put that in your '`.emacs`' file, and hitting links in w3-rendered HTML in the Gnus article buffers will use `browse-url` to follow the link.

6.5 Other Sources

Gnus can do more than just read news or mail. The methods described below allow Gnus to view directories and files as if they were newsgroups.

6.5.1 Directory Groups

If you have a directory that has lots of articles in separate files in it, you might treat it as a newsgroup. The files have to have numerical names, of course.

This might be an opportune moment to mention `ange-ftp` (and its successor `efs`), that most wonderful of all wonderful Emacs packages. When I wrote `nndir`, I didn't think much about it—a back end to read directories. Big deal.

`ange-ftp` changes that picture dramatically. For instance, if you enter the `ange-ftp` file name '`/ftp.hpc.uh.edu:/pub/emacs/ding-list/`' as the directory name, `ange-ftp` or `efs` will actually allow you to read this directory over at '`sina`' as a newsgroup. Distributed news ahoy!

`nndir` will use NOV files if they are present.

`nndir` is a "read-only" back end—you can't delete or expire articles with this method. You can use `nnmh` or `nnml` for whatever you use `nndir` for, so you could switch to any of those methods if you feel the need to have a non-read-only `nndir`.

6.5.2 Anything Groups

From the `nndir` back end (which reads a single spool-like directory), it's just a hop and a skip to `nneething`, which pretends that any arbitrary directory is a newsgroup. Strange, but true.

When `nneething` is presented with a directory, it will scan this directory and assign article numbers to each file. When you enter such a group, `nneething` must create "headers" that Gnus can use. After all, Gnus is a newsreader, in case you're forgetting. `nneething` does this in a two-step process. First, it snoops each file in question. If the file looks like an article (i.e., the first few lines look like headers), it will use this as the head. If this is just some arbitrary file without a head (e.g. a C source file), `nneething` will cobble up a header out of thin air. It will use file ownership, name and date and do whatever it can with these elements.

All this should happen automatically for you, and you will be presented with something that looks very much like a newsgroup. Totally like a newsgroup, to be precise. If you select an article, it will be displayed in the article buffer, just as usual.

If you select a line that represents a directory, Gnus will pop you into a new summary buffer for this `nneething` group. And so on. You can traverse the entire disk this way, if you feel like, but remember that Gnus is not dired, really, and does not intend to be, either.

There are two overall modes to this action—ephemeral or solid. When doing the ephemeral thing (i.e., `G D` from the group buffer), Gnus will not store information on what files you have read, and what files are new, and so on. If you create a solid `nneething` group the normal way with `G m`, Gnus will store a mapping table between article numbers and file names, and you can treat this group like any other groups. When you activate a solid `nneething` group, you will be told how many unread articles it contains, etc., etc.

Some variables:

`nneething-map-file-directory`

> All the mapping files for solid `nneething` groups will be stored in this directory, which defaults to '`~/.nneething/`'.

`nneething-exclude-files`

> All files that match this regexp will be ignored. Nice to use to exclude auto-save files and the like, which is what it does by default.

`nneething-include-files`

> Regexp saying what files to include in the group. If this variable is non-`nil`, only files matching this regexp will be included.

`nneething-map-file`

> Name of the map files.

6.5.3 Document Groups

`nndoc` is a cute little thing that will let you read a single file as a newsgroup. Several files types are supported:

`babyl`   The babyl (rmail) mail box.

`mbox`   The standard Unix mbox file.

`mmdf`   The MMDF mail box format.

`news`   Several news articles appended into a file.

`rnews`   The rnews batch transport format.

`forward`   Forwarded articles.

`nsmail`   Netscape mail boxes.

`mime-parts`
              MIME multipart messages.

`standard-digest`
              The standard (RFC 1153) digest format.

`slack-digest`
              Non-standard digest format—matches most things, but does it badly.

You can also use the special "file type" `guess`, which means that `nndoc` will try to guess what file type it is looking at. `digest` means that `nndoc` should guess what digest type the file is.

`nndoc` will not try to change the file or insert any extra headers into it—it will simply, like, let you use the file as the basis for a group. And that's it.

If you have some old archived articles that you want to insert into your new & spiffy Gnus mail back end, `nndoc` can probably help you with that. Say you have an old 'RMAIL' file with mail that you now want to split into your new `nnml` groups. You look at that file using `nndoc` (using the `G f` command in the group buffer (see Section 2.9 [Foreign Groups], page 21)), set the process mark on all the articles in the buffer (`M P b`, for instance), and then re-spool (`B r`) using `nnml`. If all goes well, all the mail in the 'RMAIL' file is now also stored in lots of `nnml` directories, and you can delete that pesky 'RMAIL' file. If you have the guts!

Virtual server variables:

`nndoc-article-type`
              This should be one of `mbox`, `babyl`, `digest`, `news`, `rnews`, `mmdf`, `forward`, `rfc934`, `rfc822-forward`, `mime-parts`, `standard-digest`, `slack-digest`, `clari-briefs`, `nsmail` or `guess`.

`nndoc-post-type`
              This variable says whether Gnus is to consider the group a news group or a mail group. There are two valid values: `mail` (the default) and `news`.

6.5.3.1 Document Server Internals

Adding new document types to be recognized by `nndoc` isn't difficult. You just have to whip up a definition of what the document looks like, write a predicate function to recognize that document type, and then hook into `nndoc`.

First, here's an example document type definition:

```
(mmdf
 (article-begin .  "^\^A\^A\^A\^A\n")
 (body-end .  "^\^A\^A\^A\^A\n"))
```

The definition is simply a unique *name* followed by a series of regexp pseudo-variable settings. Below are the possible variables—don't be daunted by the number of variables; most document types can be defined with very few settings:

`first-article`
> If present, `nndoc` will skip past all text until it finds something that match this regexp. All text before this will be totally ignored.

`article-begin`
> This setting has to be present in all document type definitions. It says what the beginning of each article looks like.

`head-begin-function`
> If present, this should be a function that moves point to the head of the article.

`nndoc-head-begin`
> If present, this should be a regexp that matches the head of the article.

`nndoc-head-end`
> This should match the end of the head of the article. It defaults to '`^$`'—the empty line.

`body-begin-function`
> If present, this function should move point to the beginning of the body of the article.

`body-begin`
> This should match the beginning of the body of the article. It defaults to '`^\n`'.

`body-end-function`
> If present, this function should move point to the end of the body of the article.

`body-end` If present, this should match the end of the body of the article.

`file-end` If present, this should match the end of the file. All text after this regexp will be totally ignored.

So, using these variables `nndoc` is able to dissect a document file into a series of articles, each with a head and a body. However, a few more variables are needed since not all document types are all that news-like—variables needed to transform the head or the body into something that's palatable for Gnus:

`prepare-body-function`
> If present, this function will be called when requesting an article. It will be called with point at the start of the body, and is useful if the document has encoded some parts of its contents.

article-transform-function

> If present, this function is called when requesting an article. It's meant to be used for more wide-ranging transformation of both head and body of the article.

generate-head-function

> If present, this function is called to generate a head that Gnus can understand. It is called with the article number as a parameter, and is expected to generate a nice head for the article in question. It is called when requesting the headers of all articles.

Let's look at the most complicated example I can come up with—standard digests:

```
(standard-digest
 (first-article . ,(concat "^" (make-string 70 ?-) "\n\n+"))
 (article-begin . ,(concat "\n\n" (make-string 30 ?-) "\n\n+"))
 (prepare-body-function . nndoc-unquote-dashes)
 (body-end-function . nndoc-digest-body-end)
 (head-end . "^ ?$")
 (body-begin . "^ ?\n")
 (file-end . "^End of .*digest.*[0-9].*\n\\*\\*\\|^End of.*Digest *$")
 (subtype digest guess))
```

We see that all text before a 70-width line of dashes is ignored; all text after a line that starts with that '`^End of`' is also ignored; each article begins with a 30-width line of dashes; the line separating the head from the body may contain a single space; and that the body is run through `nndoc-unquote-dashes` before being delivered.

To hook your own document definition into `nndoc`, use the `nndoc-add-type` function. It takes two parameters—the first is the definition itself and the second (optional) parameter says where in the document type definition alist to put this definition. The alist is traversed sequentially, and `nndoc-TYPE-type-p` is called for a given type `TYPE`. So `nndoc-mmdf-type-p` is called to see whether a document is of `mmdf` type, and so on. These type predicates should return `nil` if the document is not of the correct type; `t` if it is of the correct type; and a number if the document might be of the correct type. A high number means high probability; a low number means low probability with '`0`' being the lowest valid number.

6.5.4 SOUP

In the PC world people often talk about "offline" newsreaders. These are thingies that are combined reader/news transport monstrosities. With built-in modem programs. Yecchh!

Of course, us Unix Weenie types of human beans use things like `uucp` and, like, `nntpd` and set up proper news and mail transport things like Ghod intended. And then we just use normal newsreaders.

However, it can sometimes be convenient to do something that's a bit easier on the brain if you have a very slow modem, and you're not really that interested in doing things properly.

A file format called SOUP has been developed for transporting news and mail from servers to home machines and back again. It can be a bit fiddly.

First some terminology:

server   This is the machine that is connected to the outside world and where you get
              news and/or mail from.

home machine

              This is the machine that you want to do the actual reading and responding on.
              It is typically not connected to the rest of the world in any way.

packet   Something that contains messages and/or commands. There are two kinds of
              packets:

              message packets

                          These are packets made at the server, and typically contain lots
                          of messages for you to read. These are called 'SoupoutX.tgz' by
                          default, where x is a number.

              response packets

                          These are packets made at the home machine, and typically con-
                          tains replies that you've written. These are called 'SoupinX.tgz'
                          by default, where x is a number.

 1. You log in on the server and create a SOUP packet. You can either use a dedicated SOUP
    thingie (like the awk program), or you can use Gnus to create the packet with its SOUP
    commands (O s and/or G s b; and then G s p) (see Section 6.5.4.1 [SOUP Commands],
    page 146).

 2. You transfer the packet home. Rail, boat, car or modem will do fine.

 3. You put the packet in your home directory.

 4. You fire up Gnus on your home machine using the nnsoup back end as the native or
    secondary server.

 5. You read articles and mail and answer and followup to the things you want (see Sec-
    tion 6.5.4.3 [SOUP Replies], page 148).

 6. You do the G s r command to pack these replies into a SOUP packet.

 7. You transfer this packet to the server.

 8. You use Gnus to mail this packet out with the G s s command.

 9. You then repeat until you die.

   So you basically have a bipartite system—you use nnsoup for reading and Gnus for
packing/sending these SOUP packets.

6.5.4.1 SOUP Commands

   These are commands for creating and manipulating SOUP packets.

G s b   Pack all unread articles in the current group (gnus-group-brew-soup). This
            command understands the process/prefix convention.

G s w   Save all SOUP data files (gnus-soup-save-areas).

G s s   Send all replies from the replies packet (gnus-soup-send-replies).

G s p   Pack all files into a SOUP packet (gnus-soup-pack-packet).

G s r   Pack all replies into a replies packet (nnsoup-pack-replies).

O s   This summary-mode command adds the current article to a SOUP packet (`gnus-soup-add-article`). It understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

There are a few variables to customize where Gnus will put all these thingies:

`gnus-soup-directory`
> Directory where Gnus will save intermediate files while composing SOUP packets. The default is '`~/SoupBrew/`'.

`gnus-soup-replies-directory`
> This is what Gnus will use as a temporary directory while sending our reply packets. '`~/SoupBrew/SoupReplies/`' is the default.

`gnus-soup-prefix-file`
> Name of the file where Gnus stores the last used prefix. The default is '`gnus-prefix`'.

`gnus-soup-packer`
> A format string command for packing a SOUP packet. The default is '`tar cf - %s | gzip > $HOME/Soupout%d.tgz`'.

`gnus-soup-unpacker`
> Format string command for unpacking a SOUP packet. The default is '`gunzip -c %s | tar xvf –`'.

`gnus-soup-packet-directory`
> Where Gnus will look for reply packets. The default is '`~/`'.

`gnus-soup-packet-regexp`
> Regular expression matching SOUP reply packets in `gnus-soup-packet-directory`.

6.5.4.2 SOUP Groups

`nnsoup` is the back end for reading SOUP packets. It will read incoming packets, unpack them, and put them in a directory where you can read them at leisure.

These are the variables you can use to customize its behavior:

`nnsoup-tmp-directory`
> When `nnsoup` unpacks a SOUP packet, it does it in this directory. ('`/tmp/`' by default.)

`nnsoup-directory`
> `nnsoup` then moves each message and index file to this directory. The default is '`~/SOUP/`'.

`nnsoup-replies-directory`
> All replies will be stored in this directory before being packed into a reply packet. The default is '`~/SOUP/replies/"`'.

`nnsoup-replies-format-type`
> The SOUP format of the replies packets. The default is '`?n`' (rnews), and I don't think you should touch that variable. I probably shouldn't even have documented it. Drats! Too late!

`nnsoup-replies-index-type`
> The index type of the replies packet. The default is '`?n`', which means "none". Don't fiddle with this one either!

`nnsoup-active-file`
> Where **nnsoup** stores lots of information. This is not an "active file" in the **nntp** sense; it's an Emacs Lisp file. If you lose this file or mess it up in any way, you're dead. The default is '`~/SOUP/active`'.

`nnsoup-packer`
> Format string command for packing a reply SOUP packet. The default is '`tar cf - %s | gzip > $HOME/Soupin%d.tgz`'.

`nnsoup-unpacker`
> Format string command for unpacking incoming SOUP packets. The default is '`gunzip -c %s | tar xvf -`'.

`nnsoup-packet-directory`
> Where **nnsoup** will look for incoming packets. The default is '`~/`'.

`nnsoup-packet-regexp`
> Regular expression matching incoming SOUP packets. The default is '`Soupout`'.

`nnsoup-always-save`
> If non-`nil`, save the replies buffer after each posted message.

6.5.4.3 SOUP Replies

Just using **nnsoup** won't mean that your postings and mailings end up in SOUP reply packets automagically. You have to work a bit more for that to happen.

The **nnsoup-set-variables** command will set the appropriate variables to ensure that all your followups and replies end up in the SOUP system.

In specific, this is what it does:

```
(setq message-send-news-function 'nnsoup-request-post)
(setq message-send-mail-function 'nnsoup-request-mail)
```

And that's it, really. If you only want news to go into the SOUP system you just use the first line. If you only want mail to be SOUPed you use the second.

6.5.5 Mail-To-News Gateways

If your local **nntp** server doesn't allow posting, for some reason or other, you can post using one of the numerous mail-to-news gateways. The **nngateway** back end provides the interface.

Note that you can't read anything from this back end—it can only be used to post with.

Server variables:

`nngateway-address`
> This is the address of the mail-to-news gateway.

`nngateway-header-transformation`

> News headers often have to be transformed in some odd way or other for the mail-to-news gateway to accept it. This variable says what transformation should be called, and defaults to `nngateway-simple-header-transformation`. The function is called narrowed to the headers to be transformed and with one parameter—the gateway address.
>
> This default function just inserts a new `To` header based on the `Newsgroups` header and the gateway address. For instance, an article with this `Newsgroups` header:
>
>   Newsgroups: alt.religion.emacs
>
> will get this `From` header inserted:
>
>   To: alt-religion-emacs@GATEWAY
>
> The following pre-defined functions exist:

`nngateway-simple-header-transformation`

> Creates a `To` header that looks like *newsgroup*`@nngateway-address`.

`nngateway-mail2news-header-transformation`

> Creates a `To` header that looks like `nngateway-address`.
>
> Here's an example:
>
> ```
> (setq gnus-post-method
>       '(nngateway
>         "mail2news@replay.com"
>         (nngateway-header-transformation
>          nngateway-mail2news-header-transformation)))
> ```

So, to use this, simply say something like:

```
(setq gnus-post-method '(nngateway "GATEWAY.ADDRESS"))
```

6.5.6 IMAP

IMAP is a network protocol for reading mail (or news, or...), think of it as a modernized NNTP. Connecting to a IMAP server is much similar to connecting to a news server, you just specify the network address of the server.

IMAP has two properties. First, IMAP can do everything that POP can, it can hence be viewed as POP++. Secondly, IMAP is a mail storage protocol, similar to NNTP being a news storage protocol. (IMAP offers more features than NNTP because news is more or less read-only whereas mail is read-write.)

If you want to use IMAP as POP++, use an imap entry in mail-sources. With this, Gnus will fetch mails from the IMAP server and store them on the local disk. This is not the usage described in this section. See Section 6.3.4 [Mail Sources], page 117.

If you want to use IMAP as a mail storage protocol, use an nnimap entry in gnus-secondary-select-methods. With this, Gnus will manipulate mails stored on the IMAP server. This is the kind of usage explained in this section.

A server configuration in `~/.gnus` with a few IMAP servers might look something like the following. (Note that for SSL, you need external programs and libraries, see below.)

```
(setq gnus-secondary-select-methods
      '((nnimap "simpleserver") ; no special configuration
        ; perhaps a ssh port forwarded server:
        (nnimap "dolk"
                (nnimap-address "localhost")
                (nnimap-server-port 1430))
        ; a UW server running on localhost
        (nnimap "barbar"
                (nnimap-server-port 143)
                (nnimap-address "localhost")
                (nnimap-list-pattern ("INBOX" "mail/*")))
        ; anonymous public cyrus server:
        (nnimap "cyrus.andrew.cmu.edu"
                (nnimap-authenticator anonymous)
                (nnimap-list-pattern "archive.*")
                (nnimap-stream network))
        ; a ssl server on a non-standard port:
        (nnimap "vic20"
                (nnimap-address "vic20.somewhere.com")
                (nnimap-server-port 9930)
                (nnimap-stream ssl))))
```

The following variables can be used to create a virtual **nnimap** server:

nnimap-address

> The address of the remote IMAP server. Defaults to the virtual server name if not specified.

nnimap-server-port

> Port on server to contact. Defaults to port 143, or 993 for SSL.
>
> Note that this should be a integer, example server specification:
>
> ```
> (nnimap "mail.server.com"
>         (nnimap-server-port 4711))
> ```

nnimap-list-pattern

> String or list of strings of mailboxes to limit available groups to. This is used when the server has very many mailboxes and you're only interested in a few – some servers export your home directory via IMAP, you'll probably want to limit the mailboxes to those in '`~/Mail/*`' then.
>
> The string can also be a cons of REFERENCE and the string as above, what REFERENCE is used for is server specific, but on the University of Washington server it's a directory that will be concatenated with the mailbox.
>
> Example server specification:
>
> ```
> (nnimap "mail.server.com"
>         (nnimap-list-pattern ("INBOX" "Mail/*" "alt.sex.*"
>                                 ("~friend/Mail/" . "list/*"))))
> ```

nnimap-stream

> The type of stream used to connect to your server. By default, nnimap will detect and automatically use all of the below, with the exception of SSL. (SSL

is being replaced by STARTTLS, which can be automatically detected, but it's not widely deployed yet).

Example server specification:

```
(nnimap "mail.server.com"
        (nnimap-stream ssl))
```

Please note that the value of `nnimap-stream` is a symbol!

- *gssapi:* Connect with GSSAPI (usually Kerberos 5). Requires the `imtest` program.

- *kerberos4:* Connect with Kerberos 4. Requires the `imtest` program.

- *starttls:* Connect via the STARTTLS extension (similar to SSL). Requires the library 'starttls.el' and program `starttls`.

- *ssl:* Connect through SSL. Requires OpenSSL (the program `openssl`) or SSLeay (`s_client`) as well as the external library 'ssl.el'.

- *shell:* Use a shell command to start an IMAP connection.

- *network:* Plain, TCP/IP network connection.

The `imtest` program is shipped with Cyrus IMAPD. Nnimap supports both `imtest` version 1.5.x and version 1.6.x. The variable `imap-kerberos4-program` contains parameters to pass to the `imtest` program.

For SSL connections, the OpenSSL program is available from `http://www.openssl.org/`. OpenSSL was formerly known as SSLeay, and nnimap supports it too. However, the most recent versions of SSLeay, 0.9.x, are known to have serious bugs making it useless. Earlier versions, especially 0.8.x, of SSLeay are known to work. The variable `imap-ssl-program` contains parameters to pass to OpenSSL/SSLeay. You also need 'ssl.el' (from the W3 distribution, for instance).

For IMAP connections using the `shell` stream, the variable `imap-shell-program` specifies what program to call.

`nnimap-authenticator`

The authenticator used to connect to the server. By default, nnimap will use the most secure authenticator your server supports.

Example server specification:

```
(nnimap "mail.server.com"
        (nnimap-authenticator anonymous))
```

Please note that the value of `nnimap-authenticator` is a symbol!

- *gssapi:* GSSAPI (usually Kerberos 5) authentication. Requires the external program `imtest`.

- *kerberos4:* Kerberos authentication. Requires the external program `imtest`.

- *digest-md5:* Encrypted username/password via DIGEST-MD5. Requires external library `digest-md5.el`.

- *cram-md5:* Encrypted username/password via CRAM-MD5.

- *login:* Plain-text username/password via LOGIN.

- *anonymous:* Login as 'anonymous', supplying your email address as password.

`nnimap-expunge-on-close`

Unlike Parmenides, the IMAP designers decided that things that don't exist actually do exist. More specifically, IMAP has the concept of marking articles `Deleted` which doesn't actually delete them, and this (marking them `Deleted`, that is) is what nnimap does when you delete a article in Gnus (with `G` (DEL) or similar).

Since the articles aren't really removed when we mark them with the `Deleted` flag we'll need a way to actually delete them. Feel like running in circles yet?

Traditionally, nnimap has removed all articles marked as `Deleted` when closing a mailbox but this is now configurable by this server variable.

The possible options are:

`always`   The default behavior, delete all articles marked as "Deleted" when closing a mailbox.

`never`   Never actually delete articles. Currently there is no way of showing the articles marked for deletion in nnimap, but other IMAP clients may allow you to do this. If you ever want to run the EXPUNGE command manually, See Section 6.5.6.3 [Expunging mailboxes], page 154.

`ask`   When closing mailboxes, nnimap will ask if you wish to expunge deleted articles or not.

`nnimap-authinfo-file`

A file containing credentials used to log in on servers. The format is (almost) the same as the `ftp` '`~/.netrc`' file. See 'nntp-authinfo-file' for exact syntax.

A file containing credentials used to log in on servers. The format is (almost) the same as the `ftp` '`~/.netrc`' file. See the variable `nntp-authinfo-file` for exact syntax; also see See Section 6.2.1 [NNTP], page 109.

6.5.6.1 Splitting in IMAP

Splitting is something Gnus users have loved and used for years, and now the rest of the world is catching up. Yeah, dream on; not many IMAP servers have server side splitting and those that have splitting seem to use some non-standard protocol. This means that IMAP support for Gnus has to do its own splitting.

And it does.

Here are the variables of interest:

`nnimap-split-crosspost`

If non-nil, do crossposting if several split methods match the mail. If nil, the first match in `nnimap-split-rule` found will be used.

Nnmail equivalent: `nnmail-crosspost`.

`nnimap-split-inbox`

A string or a list of strings that gives the name(s) of IMAP mailboxes to split from. Defaults to `nil`, which means that splitting is disabled!

```
(setq nnimap-split-inbox
      '("INBOX" ("~/friend/Mail" . "lists/*") "lists.imap"))
```

No nnmail equivalent.

nnimap-split-rule

New mail found in `nnimap-split-inbox` will be split according to this variable.

This variable contains a list of lists, where the first element in the sublist gives the name of the IMAP mailbox to move articles matching the regexp in the second element in the sublist. Got that? Neither did I, we need examples.

```
(setq nnimap-split-rule
      '(("INBOX.nnimap"
          "^Sender: owner-nnimap@vic20.globalcom.se")
        ("INBOX.junk"    "^Subject:.*MAKE MONEY")
        ("INBOX.private" "")))
```

This will put all articles from the nnimap mailing list into mailbox INBOX.nnimap, all articles containing MAKE MONEY in the Subject: line into INBOX.spam and everything else in INBOX.private.

The first string may contain '*digit*' forms, like the ones used by replace-match to insert sub-expressions from the matched text. For instance:

```
("INBOX.lists.\\1"    "^Sender: owner-\\([a-z-]+\\)@")
```

The second element can also be a function. In that case, it will be called with the first element of the rule as the argument, in a buffer containing the headers of the article. It should return a non-nil value if it thinks that the mail belongs in that group.

Nnmail users might recollect that the last regexp had to be empty to match all articles (like in the example above). This is not required in nnimap. Articles not matching any of the regexps will not be moved out of your inbox. (This might affect performance if you keep lots of unread articles in your inbox, since the splitting code would go over them every time you fetch new mail.)

These rules are processed from the beginning of the alist toward the end. The first rule to make a match will 'win', unless you have crossposting enabled. In that case, all matching rules will 'win'.

This variable can also have a function as its value, the function will be called with the headers narrowed and should return a group to where it thinks the article should be split. See `nnimap-split-fancy`.

The splitting code tries to create mailboxes if it needs too.

To allow for different split rules on different virtual servers, and even different split rules in different inboxes on the same server, the syntax of this variable has been extended along the lines of:

```
(setq nnimap-split-rule
      '(("my1server"    (".*"    (("ding"    "ding@gnus.org")
                                   ("junk"    "From:.*Simon"))))
        ("my2server"    ("INBOX" nnimap-split-fancy))
        ("my[34]server" (".*"    (("private" "To:.*Simon")
                                   ("junk"    my-junk-func)))))
```

The virtual server name is in fact a regexp, so that the same rules may apply to several servers. In the example, the servers `my3server` and `my4server` both use the same rules. Similarly, the inbox string is also a regexp. The actual splitting rules are as before, either a function, or a list with group/regexp or group/function elements.

Nnmail equivalent: `nnmail-split-methods`.

`nnimap-split-predicate`

Mail matching this predicate in `nnimap-split-inbox` will be split; it is a string and the default is 'UNSEEN UNDELETED'.

This might be useful if you use another IMAP client to read mail in your inbox but would like Gnus to split all articles in the inbox regardless of readedness. Then you might change this to 'UNDELETED'.

`nnimap-split-fancy`

It's possible to set `nnimap-split-rule` to `nnmail-split-fancy` if you want to use fancy splitting. See Section 6.3.6 [Fancy Mail Splitting], page 124.

However, to be able to have different fancy split rules for nnmail and nnimap back ends you can set `nnimap-split-rule` to `nnimap-split-fancy` and define the nnimap specific fancy split rule in `nnimap-split-fancy`.

Example:

```
(setq nnimap-split-rule 'nnimap-split-fancy
      nnimap-split-fancy ...)
```

Nnmail equivalent: `nnmail-split-fancy`.

6.5.6.2 Editing IMAP ACLs

ACL stands for Access Control List. ACLs are used in IMAP for limiting (or enabling) other users access to your mail boxes. Not all IMAP servers support this, this function will give an error if it doesn't.

To edit a ACL for a mailbox, type `G l` (`gnus-group-edit-nnimap-acl`) and you'll be presented with a ACL editing window with detailed instructions.

Some possible uses:

- Giving "anyone" the "lrs" rights (lookup, read, keep seen/unseen flags) on your mailing list mailboxes enables other users on the same server to follow the list without subscribing to it.
- At least with the Cyrus server, you are required to give the user "anyone" posting ("p") capabilities to have "plussing" work (that is, mail sent to user+mailbox@domain ending up in the IMAP mailbox INBOX.mailbox).

6.5.6.3 Expunging mailboxes

If you're using the `never` setting of `nnimap-expunge-close`, you may want the option of expunging all deleted articles in a mailbox manually. This is exactly what `G x` does.

Currently there is no way of showing deleted articles, you can just delete them.

6.6 Combined Groups

Gnus allows combining a mixture of all the other group types into bigger groups.

6.6.1 Virtual Groups

An *nnvirtual group* is really nothing more than a collection of other groups.

For instance, if you are tired of reading many small groups, you can put them all in one big group, and then grow tired of reading one big, unwieldy group. The joys of computing!

You specify `nnvirtual` as the method. The address should be a regexp to match component groups.

All marks in the virtual group will stick to the articles in the component groups. So if you tick an article in a virtual group, the article will also be ticked in the component group from whence it came. (And vice versa—marks from the component groups will also be shown in the virtual group.)

Here's an example `nnvirtual` method that collects all Andrea Dworkin newsgroups into one, big, happy newsgroup:

```
(nnvirtual "^alt\\.fan\\.andrea-dworkin$\\|^rec\\.dworkin.*")
```

The component groups can be native or foreign; everything should work smoothly, but if your computer explodes, it was probably my fault.

Collecting the same group from several servers might actually be a good idea if users have set the Distribution header to limit distribution. If you would like to read 'soc.motss' both from a server in Japan and a server in Norway, you could use the following as the group regexp:

```
"^nntp\\+server\\.jp:soc\\.motss$\\|^nntp\\+server\\.no:soc\\.motss$"
```

(Remember, though, that if you're creating the group with `G m`, you shouldn't double the backslashes, and you should leave off the quote characters at the beginning and the end of the string.)

This should work kinda smoothly—all articles from both groups should end up in this one, and there should be no duplicates. Threading (and the rest) will still work as usual, but there might be problems with the sequence of articles. Sorting on date might be an option here (see Section 2.3 [Selecting a Group], page 16).

One limitation, however—all groups included in a virtual group have to be alive (i.e., subscribed or unsubscribed). Killed or zombie groups can't be component groups for `nnvirtual` groups.

If the `nnvirtual-always-rescan` is non-`nil`, `nnvirtual` will always scan groups for unread articles when entering a virtual group. If this variable is `nil` (which is the default) and you read articles in a component group after the virtual group has been activated, the read articles from the component group will show up when you enter the virtual group. You'll also see this effect if you have two virtual groups that have a component group in common. If that's the case, you should set this variable to `t`. Or you can just tap `M-g` on the virtual group every time before you enter it—it'll have much the same effect.

`nnvirtual` can have both mail and news groups as component groups. When responding to articles in `nnvirtual` groups, `nnvirtual` has to ask the back end of the component group

the article comes from whether it is a news or mail back end. However, when you do a `^`, there is typically no sure way for the component back end to know this, and in that case `nnvirtual` tells Gnus that the article came from a not-news back end. (Just to be on the safe side.)

`C-c C-t` in the message buffer will insert the `Newsgroups` line from the article you respond to in these cases.

6.6.2 Kibozed Groups

Kibozing is defined by OED as "grepping through (parts of) the news feed". `nnkiboze` is a back end that will do this for you. Oh joy! Now you can grind any NNTP server down to a halt with useless requests! Oh happiness!

To create a kibozed group, use the `G k` command in the group buffer.

The address field of the `nnkiboze` method is, as with `nnvirtual`, a regexp to match groups to be "included" in the `nnkiboze` group. That's where most similarities between `nnkiboze` and `nnvirtual` end.

In addition to this regexp detailing component groups, an `nnkiboze` group must have a score file to say what articles are to be included in the group (see Chapter 7 [Scoring], page 167).

You must run *M-x nnkiboze-generate-groups* after creating the `nnkiboze` groups you want to have. This command will take time. Lots of time. Oodles and oodles of time. Gnus has to fetch the headers from all the articles in all the component groups and run them through the scoring process to determine if there are any articles in the groups that are to be part of the `nnkiboze` groups.

Please limit the number of component groups by using restrictive regexps. Otherwise your sysadmin may become annoyed with you, and the NNTP site may throw you off and never let you back in again. Stranger things have happened.

`nnkiboze` component groups do not have to be alive—they can be dead, and they can be foreign. No restrictions.

The generation of an `nnkiboze` group means writing two files in `nnkiboze-directory`, which is '`~/News/`' by default. One contains the NOV header lines for all the articles in the group, and the other is an additional '`.newsrc`' file to store information on what groups have been searched through to find component articles.

Articles marked as read in the `nnkiboze` group will have their NOV lines removed from the NOV file.

6.7 Gnus Unplugged

In olden times (ca. February '88), people used to run their newsreaders on big machines with permanent connections to the net. News transport was dealt with by news servers, and all the newsreaders had to do was to read news. Believe it or not.

Nowadays most people read news and mail at home, and use some sort of modem to connect to the net. To avoid running up huge phone bills, it would be nice to have a way to slurp down all the news and mail, hang up the phone, read for several hours, and then upload any responses you have to make. And then you repeat the procedure.

Of course, you can use news servers for doing this as well. I've used `inn` together with `slurp`, `pop` and `sendmail` for some years, but doing that's a bore. Moving the news server functionality up to the newsreader makes sense if you're the only person reading news on a machine.

Using Gnus as an "offline" newsreader is quite simple.

- First, set up Gnus as you would do if you were running it on a machine that has full connection to the net. Go ahead. I'll still be waiting here.

- Then, put the following magical incantation at the end of your '`.gnus.el`' file:

  ```
  (gnus-agentize)
  ```

That's it. Gnus is now an "offline" newsreader.

Of course, to use it as such, you have to learn a few new commands.

6.7.1 Agent Basics

First, let's get some terminology out of the way.

The Gnus Agent is said to be *unplugged* when you have severed the connection to the net (and notified the Agent that this is the case). When the connection to the net is up again (and Gnus knows this), the Agent is *plugged*.

The *local* machine is the one you're running on, and which isn't connected to the net continuously.

Downloading means fetching things from the net to your local machine. *Uploading* is doing the opposite.

Let's take a typical Gnus session using the Agent.

- You start Gnus with `gnus-unplugged`. This brings up the Gnus Agent in a disconnected state. You can read all the news that you have already fetched while in this mode.

- You then decide to see whether any new news has arrived. You connect your machine to the net (using PPP or whatever), and then hit *J j* to make Gnus become *plugged* and use *g* to check for new mail as usual. To check for new mail in unplugged mode, see (see Section 6.3.4.1 [Mail Source Specifiers], page 117).

- You can then read the new news immediately, or you can download the news onto your local machine. If you want to do the latter, you press *g* to check if there are any new news and then *J s* to fetch all the eligible articles in all the groups. (To let Gnus know which articles you want to download, see Section 6.7.2 [Agent Categories], page 158.)

- After fetching the articles, you press *J j* to make Gnus become unplugged again, and you shut down the PPP thing (or whatever). And then you read the news offline.

- And then you go to step 2.

Here are some things you should do the first time (or so) that you use the Agent.

- Decide which servers should be covered by the Agent. If you have a mail back end, it would probably be nonsensical to have it covered by the Agent. Go to the server buffer (`^` in the group buffer) and press *J a* the server (or servers) that you wish to have covered by the Agent (see Section 6.7.3.3 [Server Agent Commands], page 163). This will typically be only the primary select method, which is listed on the bottom in the buffer.

- Decide on download policy. See Section 6.7.2 [Agent Categories], page 158.
- Uhm... that's it.

6.7.2 Agent Categories

One of the main reasons to integrate the news transport layer into the newsreader is to allow greater control over what articles to download. There's not much point in downloading huge amounts of articles, just to find out that you're not interested in reading any of them. It's better to be somewhat more conservative in choosing what to download, and then mark the articles for downloading manually if it should turn out that you're interested in the articles anyway.

The main way to control what is to be downloaded is to create a *category* and then assign some (or all) groups to this category. Groups that do not belong in any other category belong to the `default` category. Gnus has its own buffer for creating and managing categories.

6.7.2.1 Category Syntax

A category consists of two things.

1. A predicate which (generally) gives a rough outline of which articles are eligible for downloading; and

2. a score rule which (generally) gives you a finer granularity when deciding what articles to download. (Note that this *download score* is not necessarily related to normal scores.)

A predicate in its simplest form can be a single predicate such as `true` or `false`. These two will download every available article or nothing respectively. In the case of these two special predicates an additional score rule is superfluous.

Predicates of `high` or `low` download articles in respect of their scores in relationship to `gnus-agent-high-score` and `gnus-agent-low-score` as described below.

To gain even finer control of what is to be regarded eligible for download a predicate can consist of a number of predicates with logical operators sprinkled in between.

Perhaps some examples are in order.

Here's a simple predicate. (It's the default predicate, in fact, used for all groups that don't belong to any other category.)

```
short
```

Quite simple, eh? This predicate is true if and only if the article is short (for some value of "short").

Here's a more complex predicate:

```
(or high
    (and
     (not low)
     (not long)))
```

This means that an article should be downloaded if it has a high score, or if the score is not low and the article is not long. You get the drift.

The available logical operators are `or`, `and` and `not`. (If you prefer, you can use the more "C"-ish operators '`|`', `&` and `!` instead.)

The following predicates are pre-defined, but if none of these fit what you want to do, you can write your own.

`short`   True iff the article is shorter than `gnus-agent-short-article` lines; default 100.

`long`   True iff the article is longer than `gnus-agent-long-article` lines; default 200.

`low`   True iff the article has a download score less than `gnus-agent-low-score`; default 0.

`high`   True iff the article has a download score greater than `gnus-agent-high-score`; default 0.

`spam`   True iff the Gnus Agent guesses that the article is spam. The heuristics may change over time, but at present it just computes a checksum and sees whether articles match.

`true`   Always true.

`false`   Always false.

If you want to create your own predicate function, here's what you have to know: The functions are called with no parameters, but the `gnus-headers` and `gnus-score` dynamic variables are bound to useful values.

For example, you could decide that you don't want to download articles that were posted more than a certain number of days ago (e.g. posted more than `gnus-agent-expire-days` ago) you might write a function something along the lines of the following:

```
(defun my-article-old-p ()
  "Say whether an article is old."
  (< (time-to-days (date-to-time (mail-header-date gnus-headers)))
     (- (time-to-days (current-time)) gnus-agent-expire-days)))
```

with the predicate then defined as:

```
(not my-article-old-p)
```

or you could append your predicate to the predefined `gnus-category-predicate-alist` in your '`~/.gnus.el`' or wherever. (Note: this would have to be at a point *after* `gnus-agent` has been loaded via (`gnus-agentize`))

```
(setq  gnus-category-predicate-alist
  (append gnus-category-predicate-alist
 '((old . my-article-old-p))))
```

and simply specify your predicate as:

```
(not old)
```

If/when using something like the above, be aware that there are many misconfigured systems/mailers out there and so an article's date is not always a reliable indication of when it was posted. Hell, some people just don't give a damn.

The above predicates apply to *all* the groups which belong to the category. However, if you wish to have a specific predicate for an individual group within a category, or you're

just too lazy to set up a new category, you can enter a group's individual predicate in it's group parameters like so:

```
(agent-predicate . short)
```

This is the group parameter equivalent of the agent category default. Note that when specifying a single word predicate like this, the `agent-predicate` specification must be in dotted pair notation.

The equivalent of the longer example from above would be:

```
(agent-predicate or high (and (not low) (not long)))
```

The outer parenthesis required in the category specification are not entered here as, not being in dotted pair notation, the value of the predicate is assumed to be a list.

Now, the syntax of the download score is the same as the syntax of normal score files, except that all elements that require actually seeing the article itself are verboten. This means that only the following headers can be scored on: `Subject`, `From`, `Date`, `Message-ID`, `References`, `Chars`, `Lines`, and `Xref`.

As with predicates, the specification of the `download score rule` to use in respect of a group can be in either the category definition if it's to be applicable to all groups in therein, or a group's parameters if it's to be specific to that group.

In both of these places the `download score rule` can take one of three forms:

1. Score rule

   This has the same syntax as a normal gnus score file except only a subset of scoring keywords are available as mentioned above.

   example:

   - Category specification

     ```
     (("from"
             ("Lars Ingebrigtsen" 1000000 nil s))
     ("lines"
             (500 -100 nil <)))
     ```

   - Group Parameter specification

     ```
     (agent-score ("from"
                             ("Lars Ingebrigtsen" 1000000 nil s))
                     ("lines"
                             (500 -100 nil <)))
     ```

   Again, note the omission of the outermost parenthesis here.

2. Agent score file

   These score files must *only* contain the permitted scoring keywords stated above.

   example:

   - Category specification

     ```
     ("~/News/agent.SCORE")
     ```

     or perhaps

     ```
     ("~/News/agent.SCORE" "~/News/agent.group.SCORE")
     ```

   - Group Parameter specification

```
(agent-score "~/News/agent.SCORE")
```

Additional score files can be specified as above. Need I say anything about parenthesis?

3. Use `normal` score files

   If you don't want to maintain two sets of scoring rules for a group, and your desired `downloading` criteria for a group are the same as your `reading` criteria then you can tell the agent to refer to your `normal` score files when deciding what to download.

   These directives in either the category definition or a group's parameters will cause the agent to read in all the applicable score files for a group, *filtering out* those sections that do not relate to one of the permitted subset of scoring keywords.

   - Category Specification
     ```
     file
     ```
   - Group Parameter specification
     ```
     (agent-score . file)
     ```

6.7.2.2 The Category Buffer

You'd normally do all category maintenance from the category buffer. When you enter it for the first time (with the *J c* command from the group buffer), you'll only see the `default` category.

The following commands are available in this buffer:

q   Return to the group buffer (`gnus-category-exit`).

k   Kill the current category (`gnus-category-kill`).

c   Copy the current category (`gnus-category-copy`).

a   Add a new category (`gnus-category-add`).

p   Edit the predicate of the current category (`gnus-category-edit-predicate`).

g   Edit the list of groups belonging to the current category (`gnus-category-edit-groups`).

s   Edit the download score rule of the current category (`gnus-category-edit-score`).

l   List all the categories (`gnus-category-list`).

6.7.2.3 Category Variables

`gnus-category-mode-hook`
         Hook run in category buffers.

`gnus-category-line-format`
         Format of the lines in the category buffer (see Section 8.4 [Formatting Variables], page 190). Valid elements are:

         'c'   The name of the category.

'g'   The number of groups in the category.

`gnus-category-mode-line-format`
           Format of the category mode line (see Section 8.4.2 [Mode Line Formatting],
           page 191).

`gnus-agent-short-article`
           Articles that have fewer lines than this are short. Default 100.

`gnus-agent-long-article`
           Articles that have more lines than this are long. Default 200.

`gnus-agent-low-score`
           Articles that have a score lower than this have a low score. Default 0.

`gnus-agent-high-score`
           Articles that have a score higher than this have a high score. Default 0.

6.7.3 Agent Commands

All the Gnus Agent commands are on the *J* submap. The *J j* (`gnus-agent-toggle-plugged` command works in all modes, and toggles the plugged/unplugged state of the Gnus Agent.

You can run a complete batch fetch from the command line with the following incantation:

```
$ emacs -batch -l ~/.gnus.el -f gnus-agent-batch-fetch
```

6.7.3.1 Group Agent Commands

J u   Fetch all eligible articles in the current group (`gnus-agent-fetch-groups`).

J c   Enter the Agent category buffer (`gnus-enter-category-buffer`).

J s   Fetch all eligible articles in all groups (`gnus-agent-fetch-session`).

J S   Send all sendable messages in the draft group (`gnus-group-send-drafts`). See Section 5.6 [Drafts], page 104.

J a   Add the current group to an Agent category (`gnus-agent-add-group`). This command understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

J r   Remove the current group from its category, if any (`gnus-agent-remove-group`). This command understands the process/prefix convention (see Section 8.1 [Process/Prefix], page 189).

J Y   Synchronize flags changed while unplugged with remote server, if any.

6.7.3.2 Summary Agent Commands

J #   Mark the article for downloading (`gnus-agent-mark-article`).

J M-#   Remove the downloading mark from the article (`gnus-agent-unmark-article`).

@   Toggle whether to download the article (`gnus-agent-toggle-mark`).

J c   Mark all undownloaded articles as read (`gnus-agent-catchup`).

6.7.3.3 Server Agent Commands

J a   Add the current server to the list of servers covered by the Gnus Agent (`gnus-agent-add-server`).

J r   Remove the current server from the list of servers covered by the Gnus Agent (`gnus-agent-remove-server`).

6.7.4 Agent Expiry

`nnagent` doesn't handle expiry. Instead, there's a special `gnus-agent-expire` command that will expire all read articles that are older than `gnus-agent-expire-days` days. It can be run whenever you feel that you're running out of space. It's not particularly fast or efficient, and it's not a particularly good idea to interrupt it (with *C-g* or anything else) once you've started it.

if `gnus-agent-expire-all` is non-`nil`, this command will expire all articles—unread, read, ticked and dormant. If `nil` (which is the default), only read articles are eligible for expiry, and unread, ticked and dormant articles will be kept indefinitely.

6.7.5 Agent and IMAP

The Agent work with any Gnus back end, including nnimap. However, since there are some conceptual differences between NNTP and IMAP, this section (should) provide you with some information to make Gnus Agent work smoother as a IMAP Disconnected Mode client.

The first thing to keep in mind is that all flags (read, ticked, etc) are kept on the IMAP server, rather than in '`.newsrc`' as is the case for nntp. Thus Gnus need to remember flag changes when disconnected, and synchronize these flags when you plug back in.

Gnus keep track of flag changes when reading nnimap groups under the Agent by default. When you plug back in, by default Gnus will check if you have any changed any flags and ask if you wish to synchronize these with the server. This behavior is customizable with `gnus-agent-synchronize-flags`.

If `gnus-agent-synchronize-flags` is `nil`, the Agent will never automatically synchronize flags. If it is `ask`, the default, the Agent will check if you made any changes and if so ask if you wish to synchronize these when you re-connect. If it has any other value, all flags will be synchronized automatically.

If you do not wish to automatically synchronize flags when you re-connect, this can be done manually with the `gnus-agent-synchronize-flags` command that is bound to *J Y* in the group buffer by default.

Some things are currently not implemented in the Agent that you'd might expect from a disconnected IMAP client, including:

- Copying/moving articles into nnimap groups when unplugged.
- Creating/deleting nnimap groups when unplugged.

Technical note: the synchronization algorithm does not work by "pushing" all local flags to the server, but rather incrementally update the server view of flags by changing only those flags that were changed by the user. Thus, if you set one flag on a article, quit the group and re-select the group and remove the flag; the flag will be set and removed from the server when you "synchronize". The queued flag operations can be found in the per-server `flags` file in the Agent directory. It's emptied when you synchronize flags.

6.7.6 Outgoing Messages

When Gnus is unplugged, all outgoing messages (both mail and news) are stored in the draft groups (see Section 5.6 [Drafts], page 104). You can view them there after posting, and edit them at will.

When Gnus is plugged again, you can send the messages either from the draft group with the special commands available there, or you can use the `J S` command in the group buffer to send all the sendable messages in the draft group.

6.7.7 Agent Variables

`gnus-agent-directory`
>   Where the Gnus Agent will store its files. The default is '`~/News/agent/`'.

`gnus-agent-handle-level`
>   Groups on levels (see Section 2.6 [Group Levels], page 19) higher than this variable will be ignored by the Agent. The default is `gnus-level-subscribed`, which means that only subscribed group will be considered by the Agent by default.

`gnus-agent-plugged-hook`
>   Hook run when connecting to the network.

`gnus-agent-unplugged-hook`
>   Hook run when disconnecting from the network.

6.7.8 Example Setup

If you don't want to read this manual, and you have a fairly standard setup, you may be able to use something like the following as your '`.gnus.el`' file to get started.

```
;;; Define how Gnus is to fetch news.  We do this over NNTP
;;; from your ISP's server.
(setq gnus-select-method '(nntp "news.your-isp.com"))

;;; Define how Gnus is to read your mail.  We read mail from
;;; your ISP's POP server.
(setq mail-sources '((pop :server "pop.your-isp.com")))

;;; Say how Gnus is to store the mail.  We use nnml groups.
(setq gnus-secondary-select-methods '((nnml "")))

;;; Make Gnus into an offline newsreader.
```

```
(gnus-agentize)
```

That should be it, basically. Put that in your '`~/.gnus.el`' file, edit to suit your needs, start up PPP (or whatever), and type *M-x gnus*.

If this is the first time you've run Gnus, you will be subscribed automatically to a few default newsgroups. You'll probably want to subscribe to more groups, and to do that, you have to query the NNTP server for a complete list of groups with the *A A* command. This usually takes quite a while, but you only have to do it once.

After reading and parsing a while, you'll be presented with a list of groups. Subscribe to the ones you want to read with the *u* command. *l* to make all the killed groups disappear after you've subscribe to all the groups you want to read. (*A k* will bring back all the killed groups.)

You can now read the groups at once, or you can download the articles with the *J s* command. And then read the rest of this manual to find out which of the other gazillion things you want to customize.

6.7.9 Batching Agents

Having the Gnus Agent fetch articles (and post whatever messages you've written) is quite easy once you've gotten things set up properly. The following shell script will do everything that is necessary:

```
#!/bin/sh
emacs -batch -l ~/.emacs -f gnus-agent-batch >/dev/null
```

6.7.10 Agent Caveats

The Gnus Agent doesn't seem to work like most other offline newsreaders. Here are some common questions that some imaginary people may ask:

If I read an article while plugged, do they get entered into the
         Agent?

         No.

If I read an article while plugged, and the article already exists
         in the Agent, will it get downloaded once more?

         Yes.

In short, when Gnus is unplugged, it only looks into the locally stored articles; when it's plugged, it only talks to your ISP.

7 Scoring

Other people use *kill files*, but we here at Gnus Towers like scoring better than killing, so we'd rather switch than fight. They do something completely different as well, so sit up straight and pay attention!

All articles have a default score (`gnus-summary-default-score`), which is 0 by default. This score may be raised or lowered either interactively or by score files. Articles that have a score lower than `gnus-summary-mark-below` are marked as read.

Gnus will read any *score files* that apply to the current group before generating the summary buffer.

There are several commands in the summary buffer that insert score entries based on the current article. You can, for instance, ask Gnus to lower or increase the score of all articles with a certain subject.

There are two sorts of scoring entries: Permanent and temporary. Temporary score entries are self-expiring entries. Any entries that are temporary and have not been used for, say, a week, will be removed silently to help keep the sizes of the score files down.

7.1 Summary Score Commands

The score commands that alter score entries do not actually modify real score files. That would be too inefficient. Gnus maintains a cache of previously loaded score files, one of which is considered the *current score file alist*. The score commands simply insert entries into this list, and upon group exit, this list is saved.

The current score file is by default the group's local score file, even if no such score file actually exists. To insert score commands into some other score file (e.g. 'all.SCORE'), you must first make this score file the current one.

General score commands that don't actually change the score file:

V s   Set the score of the current article (`gnus-summary-set-score`).

V S   Display the score of the current article (`gnus-summary-current-score`).

V t   Display all score rules that have been used on the current article (`gnus-score-find-trace`).

V R   Run the current summary through the scoring process (`gnus-summary-rescore`). This might be useful if you're playing around with your score files behind Gnus' back and want to see the effect you're having.

V c   Make a different score file the current (`gnus-score-change-score-file`).

V e   Edit the current score file (`gnus-score-edit-current-scores`). You will be popped into a `gnus-score-mode` buffer (see Section 7.5 [Score File Editing], page 176).

V f   Edit a score file and make this score file the current one (`gnus-score-edit-file`).

V F   Flush the score cache (`gnus-score-flush-cache`). This is useful after editing score files.

V C   Customize a score file in a visually pleasing manner (`gnus-score-customize`).

   The rest of these commands modify the local score file.

V m   Prompt for a score, and mark all articles with a score below this as read (`gnus-score-set-mark-below`).

V x   Prompt for a score, and add a score rule to the current score file to expunge all articles below this score (`gnus-score-set-expunge-below`).

   The keystrokes for actually making score entries follow a very regular pattern, so there's no need to list all the commands. (Hundreds of them.)

1. The first key is either *I* (upper case i) for increasing the score or *L* for lowering the score.

2. The second key says what header you want to score on. The following keys are available:

   a   Score on the author name.

   s   Score on the subject line.

   x   Score on the `Xref` line—i.e., the cross-posting line.

   r   Score on the `References` line.

   d   Score on the date.

   l   Score on the number of lines.

   i   Score on the `Message-ID` header.

   f   Score on followups—this matches the author name, and adds scores to the followups to this author.

   b   Score on the body.

   h   Score on the head.

   t   Score on thread.

3. The third key is the match type. Which match types are valid depends on what headers you are scoring on.

   strings

                   e   Exact matching.

                   s   Substring matching.

                   f   Fuzzy matching (see Section 8.17 [Fuzzy Matching], page 208).

                   r   Regexp matching

   date

                   b   Before date.

                   a   After date.

                   n   This date.

   number

<	Less than number.
=	Equal to number.
>	Greater than number.

4. The fourth and final key says whether this is a temporary (i.e., expiring) score entry, or a permanent (i.e., non-expiring) score entry, or whether it is to be done immediately, without adding to the score file.

t	Temporary score entry.
p	Permanent score entry.
i	Immediately scoring.

So, let's say you want to increase the score on the current author with exact matching permanently: *I a e p*. If you want to lower the score based on the subject line, using substring matching, and make a temporary score entry: *L s s t*. Pretty easy.

To make things a bit more complicated, there are shortcuts. If you use a capital letter on either the second or third keys, Gnus will use defaults for the remaining one or two keystrokes. The defaults are "substring" and "temporary". So *I A* is the same as *I a s t*, and *I a R* is the same as *I a r t*.

These functions take both the numerical prefix and the symbolic prefix (see Section 8.3 [Symbolic Prefixes], page 190). A numerical prefix says how much to lower (or increase) the score of the article. A symbolic prefix of `a` says to use the 'all.SCORE' file for the command instead of the current score file.

The `gnus-score-mimic-keymap` says whether these commands will pretend they are keymaps or not.

7.2 Group Score Commands

There aren't many of these as yet, I'm afraid.

W f   Gnus maintains a cache of score alists to avoid having to reload them all the time. This command will flush the cache (`gnus-score-flush-cache`).

You can do scoring from the command line by saying something like:

```
$ emacs -batch -l ~/.emacs -l ~/.gnus.el -f gnus-batch-score
```

7.3 Score Variables

`gnus-use-scoring`
        If `nil`, Gnus will not check for score files, and will not, in general, do any score-related work. This is `t` by default.

`gnus-kill-killed`
        If this variable is `nil`, Gnus will never apply score files to articles that have already been through the kill process. While this may save you lots of time, it also means that if you apply a kill file to a group, and then change the kill file and want to run it over you group again to kill more articles, it won't work. You have to set this variable to `t` to do that. (It is `t` by default.)

`gnus-kill-files-directory`
>   All kill and score files will be stored in this directory, which is initialized from the `SAVEDIR` environment variable by default. This is '`~/News/`' by default.

`gnus-score-file-suffix`
>   Suffix to add to the group name to arrive at the score file name ('`SCORE`' by default.)

`gnus-score-uncacheable-files`
>   All score files are normally cached to avoid excessive re-loading of score files. However, if this might make your Emacs grow big and bloated, so this regexp can be used to weed out score files unlikely to be needed again. It would be a bad idea to deny caching of '`all.SCORE`', while it might be a good idea to not cache '`comp.infosystems.www.authoring.misc.ADAPT`'. In fact, this variable is '`ADAPT$`' by default, so no adaptive score files will be cached.

`gnus-save-score`
>   If you have really complicated score files, and do lots of batch scoring, then you might set this variable to `t`. This will make Gnus save the scores into the '`.newsrc.eld`' file.
>
>   If you do not set this to `t`, then manual scores (like those set with `V s` (`gnus-summary-set-score`)) will not be preserved across group visits.

`gnus-score-interactive-default-score`
>   Score used by all the interactive raise/lower commands to raise/lower score with. Default is 1000, which may seem excessive, but this is to ensure that the adaptive scoring scheme gets enough room to play with. We don't want the small changes from the adaptive scoring to overwrite manually entered data.

`gnus-summary-default-score`
>   Default score of an article, which is 0 by default.

`gnus-summary-expunge-below`
>   Don't display the summary lines of articles that have scores lower than this variable. This is `nil` by default, which means that no articles will be hidden. This variable is local to the summary buffers, and has to be set from `gnus-summary-mode-hook`.

`gnus-score-over-mark`
>   Mark (in the third column) used for articles with a score over the default. Default is '`+`'.

`gnus-score-below-mark`
>   Mark (in the third column) used for articles with a score below the default. Default is '`-`'.

`gnus-score-find-score-files-function`
>   Function used to find score files for the current group. This function is called with the name of the group as the argument.
>
>   Predefined functions available are:
>
>   `gnus-score-find-single`
>   >   Only apply the group's own score file.

gnus-score-find-bnews

> Apply all score files that match, using bnews syntax. This is the default. If the current group is 'gnu.emacs.gnus', for instance, 'all.emacs.all.SCORE', 'not.alt.all.SCORE' and 'gnu.all.SCORE' would all apply. In short, the instances of 'all' in the score file names are translated into '.*', and then a regexp match is done.
>
> This means that if you have some score entries that you want to apply to all groups, then you put those entries in the 'all.SCORE' file.
>
> The score files are applied in a semi-random order, although Gnus will try to apply the more general score files before the more specific score files. It does this by looking at the number of elements in the score file names—discarding the 'all' elements.

gnus-score-find-hierarchical

> Apply all score files from all the parent groups. This means that you can't have score files like 'all.SCORE', but you can have 'SCORE', 'comp.SCORE' and 'comp.emacs.SCORE' for each server.

This variable can also be a list of functions. In that case, all these functions will be called with the group name as argument, and all the returned lists of score files will be applied. These functions can also return lists of score alists directly. In that case, the functions that return these non-file score alists should probably be placed before the "real" score file functions, to ensure that the last score file returned is the local score file. Phu.

For example, to do hierarchical scoring but use a non-server-specific overall score file, you could use the value

```
(list (lambda (group) ("all.SCORE"))
      'gnus-score-find-hierarchical)
```

gnus-score-expiry-days

> This variable says how many days should pass before an unused score file entry is expired. If this variable is `nil`, no score file entries are expired. It's 7 by default.

gnus-update-score-entry-dates

> If this variable is non-`nil`, matching score entries will have their dates updated. (This is how Gnus controls expiry—all non-matching entries will become too old while matching entries will stay fresh and young.) However, if you set this variable to `nil`, even matching entries will grow old and will have to face that oh-so grim reaper.

gnus-score-after-write-file-function

> Function called with the name of the score file just written.

gnus-score-thread-simplify

> If this variable is non-`nil`, article subjects will be simplified for subject scoring purposes in the same manner as with threading—according to the current value

of gnus-simplify-subject-functions. If the scoring entry uses `substring` or `exact` matching, the match will also be simplified in this manner.

7.4 Score File Format

A score file is an `emacs-lisp` file that normally contains just a single form. Casual users are not expected to edit these files; everything can be changed from the summary buffer.

Anyway, if you'd like to dig into it yourself, here's an example:

```
(("from"
  ("Lars Ingebrigtsen" -10000)
  ("Per Abrahamsen")
  ("larsi\\|lmi" -50000 nil R))
 ("subject"
  ("Ding is Badd" nil 728373))
 ("xref"
  ("alt.politics" -1000 728372 s))
 ("lines"
  (2 -100 nil <))
 (mark 0)
 (expunge -1000)
 (mark-and-expunge -10)
 (read-only nil)
 (orphan -10)
 (adapt t)
 (files "/hom/larsi/News/gnu.SCORE")
 (exclude-files "all.SCORE")
 (local (gnus-newsgroup-auto-expire t)
        (gnus-summary-make-false-root empty))
 (eval (ding)))
```

This example demonstrates most score file elements. For a different approach, see see Section 7.15 [Advanced Scoring], page 185.

Even though this looks much like lisp code, nothing here is actually `eval`ed. The lisp reader is used to read this form, though, so it has to be valid syntactically, if not semantically.

Six keys are supported by this alist:

STRING   If the key is a string, it is the name of the header to perform the match on. Scoring can only be performed on these eight headers: `From`, `Subject`, `References`, `Message-ID`, `Xref`, `Lines`, `Chars` and `Date`. In addition to these headers, there are three strings to tell Gnus to fetch the entire article and do the match on larger parts of the article: `Body` will perform the match on the body of the article, `Head` will perform the match on the head of the article, and `All` will perform the match on the entire article. Note that using any of these last three keys will slow down group entry *considerably*. The final "header" you can score on is `Followup`. These score entries will result in new score entries being added for all follow-ups to articles that matches these score entries.

Following this key is a arbitrary number of score entries, where each score entry has one to four elements.

1. The first element is the *match element*. On most headers this will be a string, but on the Lines and Chars headers, this must be an integer.

2. If the second element is present, it should be a number—the *score element*. This number should be an integer in the neginf to posinf interval. This number is added to the score of the article if the match is successful. If this element is not present, the `gnus-score-interactive-default-score` number will be used instead. This is 1000 by default.

3. If the third element is present, it should be a number—the *date element*. This date says when the last time this score entry matched, which provides a mechanism for expiring the score entries. It this element is not present, the score entry is permanent. The date is represented by the number of days since December 31, 1 BCE.

4. If the fourth element is present, it should be a symbol—the *type element*. This element specifies what function should be used to see whether this score entry matches the article. What match types that can be used depends on what header you wish to perform the match on.

   From, Subject, References, Xref, Message-ID
   > For most header types, there are the `r` and `R` (regexp), as well as `s` and `S` (substring) types, and `e` and `E` (exact match), and `w` (word match) types. If this element is not present, Gnus will assume that substring matching should be used. `R`, `S`, and `E` differ from the others in that the matches will be done in a case-sensitive manner. All these one-letter types are really just abbreviations for the `regexp`, `string`, `exact`, and `word` types, which you can use instead, if you feel like.

   Lines, Chars
   > These two headers use different match types: `<`, `>`, `=`, `>=` and `<=`.
   >
   > These predicates are true if
   >
   > > `(PREDICATE HEADER MATCH)`
   >
   > evaluates to non-`nil`. For instance, the advanced match (`"lines" 4 <`) (see Section 7.15 [Advanced Scoring], page 185) will result in the following form:
   >
   > > `(< header-value 4)`
   >
   > Or to put it another way: When using `<` on `Lines` with 4 as the match, we get the score added if the article has less than 4 lines. (It's easy to get confused and think it's the other way around. But it's not. I think.)
   >
   > When matching on `Lines`, be careful because some back ends (like `nndir`) do not generate `Lines` header, so every article ends up being marked as having 0 lines. This can lead to strange results if you happen to lower score of the articles with few lines.

Date   For the Date header we have three kinda silly match types:
              `before`, `at` and `after`. I can't really imagine this ever being
              useful, but, like, it would feel kinda silly not to provide this
              function. Just in case. You never know. Better safe than sorry.
              Once burnt, twice shy. Don't judge a book by its cover. Never
              not have sex on a first date. (I have been told that at least
              one person, and I quote, "found this function indispensable",
              however.)

              A more useful match type is `regexp`. With it, you
              can match the date string using a regular expression.
              The date is normalized to ISO8601 compact format
              first—*YYYYMMDD*T*HHMMSS*. If you want to match all
              articles that have been posted on April 1st in every year,
              you could use '`....0401.........`' as a match string, for
              instance. (Note that the date is kept in its original time zone,
              so this will match articles that were posted when it was April
              1st where the article was posted from. Time zones are such
              wholesome fun for the whole family, eh?)

Head, Body, All
              These three match keys use the same match types as the `From`
              (etc) header uses.

Followup   This match key is somewhat special, in that it will match the
              `From` header, and affect the score of not only the matching
              articles, but also all followups to the matching articles. This
              allows you e.g. increase the score of followups to your own
              articles, or decrease the score of followups to the articles of
              some known trouble-maker. Uses the same match types as the
              `From` header uses. (Using this match key will lead to creation
              of '`ADAPT`' files.)

Thread   This match key works along the same lines as the `Followup`
              match key. If you say that you want to score on a (sub-
)thread started by an article with a `Message-ID` *x*, then you
              add a '`thread`' match. This will add a new '`thread`' match for
              each article that has *x* in its `References` header. (These new
              '`thread`' matches will use the `Message-ID`s of these matching
              articles.) This will ensure that you can raise/lower the score
              of an entire thread, even though some articles in the thread
              may not have complete `References` headers. Note that using
              this may lead to undeterministic scores of the articles in the
              thread. (Using this match key will lead to creation of '`ADAPT`'
              files.)

mark   The value of this entry should be a number. Any articles with a score lower
              than this number will be marked as read.

expunge   The value of this entry should be a number. Any articles with a score lower
              than this number will be removed from the summary buffer.

mark-and-expunge

> The value of this entry should be a number. Any articles with a score lower than this number will be marked as read and removed from the summary buffer.

thread-mark-and-expunge

> The value of this entry should be a number. All articles that belong to a thread that has a total score below this number will be marked as read and removed from the summary buffer. `gnus-thread-score-function` says how to compute the total score for a thread.

files   The value of this entry should be any number of file names. These files are assumed to be score files as well, and will be loaded the same way this one was.

exclude-files

> The clue of this entry should be any number of files. These files will not be loaded, even though they would normally be so, for some reason or other.

eval   The value of this entry will be `eval`el. This element will be ignored when handling global score files.

read-only

> Read-only score files will not be updated or saved. Global score files should feature this atom (see Section 7.11 [Global Score Files], page 181). (Note: *Global* here really means *global*; not your personal apply-to-all-groups score files.)

orphan   The value of this entry should be a number. Articles that do not have parents will get this number added to their scores. Imagine you follow some high-volume newsgroup, like 'comp.lang.c'. Most likely you will only follow a few of the threads, also want to see any new threads.

> You can do this with the following two score file entries:

>> ```
>> (orphan -500)
>> (mark-and-expunge -100)
>> ```

> When you enter the group the first time, you will only see the new threads. You then raise the score of the threads that you find interesting (with *I T* or *I S*), and ignore (*C y*) the rest. Next time you enter the group, you will see new articles in the interesting threads, plus any new threads.

> I.e.—the orphan score atom is for high-volume groups where there exist a few interesting threads which can't be found automatically by ordinary scoring rules.

adapt   This entry controls the adaptive scoring. If it is `t`, the default adaptive scoring rules will be used. If it is `ignore`, no adaptive scoring will be performed on this group. If it is a list, this list will be used as the adaptive scoring rules. If it isn't present, or is something other than `t` or `ignore`, the default adaptive scoring rules will be used. If you want to use adaptive scoring on most groups, you'd set `gnus-use-adaptive-scoring` to `t`, and insert an `(adapt ignore)` in the groups where you do not want adaptive scoring. If you only want adaptive scoring in a few groups, you'd set `gnus-use-adaptive-scoring` to `nil`, and insert `(adapt t)` in the score files of the groups where you want it.

adapt-file
>   All adaptive score entries will go to the file named by this entry. It will also
>   be applied when entering the group. This atom might be handy if you want to
>   adapt on several groups at once, using the same adaptive file for a number of
>   groups.

local   The value of this entry should be a list of (VAR VALUE) pairs. Each *var* will be
>   made buffer-local to the current summary buffer, and set to the value specified.
>   This is a convenient, if somewhat strange, way of setting variables in some
>   groups if you don't like hooks much. Note that the *value* won't be evaluated.

7.5 Score File Editing

You normally enter all scoring commands from the summary buffer, but you might feel
the urge to edit them by hand as well, so we've supplied you with a mode for that.

It's simply a slightly customized `emacs-lisp` mode, with these additional commands:

C-c C-c   Save the changes you have made and return to the summary buffer (`gnus-score-edit-done`).

C-c C-d   Insert the current date in numerical format (`gnus-score-edit-insert-date`).
>   This is really the day number, if you were wondering.

C-c C-p   The adaptive score files are saved in an unformatted fashion. If you intend
>   to read one of these files, you want to *pretty print* it first. This command
>   (`gnus-score-pretty-print`) does that for you.

Type *M-x gnus-score-mode* to use this mode.

`gnus-score-menu-hook` is run in score mode buffers.

In the summary buffer you can use commands like *V f* and *V e* to begin editing score
files.

7.6 Adaptive Scoring

If all this scoring is getting you down, Gnus has a way of making it all happen
automatically—as if by magic. Or rather, as if by artificial stupidity, to be precise.

When you read an article, or mark an article as read, or kill an article, you leave marks
behind. On exit from the group, Gnus can sniff these marks and add score elements depend-
ing on what marks it finds. You turn on this ability by setting `gnus-use-adaptive-scoring`
to `t` or `(line)`. If you want score adaptively on separate words appearing in the subjects,
you should set this variable to `(word)`. If you want to use both adaptive methods, set this
variable to `(word line)`.

To give you complete control over the scoring process, you can customize the `gnus-default-adaptive-score-alist` variable. For instance, it might look something like this:

```
(setq gnus-default-adaptive-score-alist
  '((gnus-unread-mark)
    (gnus-ticked-mark (from 4))
    (gnus-dormant-mark (from 5))
```

```
(gnus-del-mark (from -4) (subject -1))
(gnus-read-mark (from 4) (subject 2))
(gnus-expirable-mark (from -1) (subject -1))
(gnus-killed-mark (from -1) (subject -3))
(gnus-kill-file-mark)
(gnus-ancient-mark)
(gnus-low-score-mark)
(gnus-catchup-mark (from -1) (subject -1))))
```

As you see, each element in this alist has a mark as a key (either a variable name or a "real" mark—a character). Following this key is a arbitrary number of header/score pairs. If there are no header/score pairs following the key, no adaptive scoring will be done on articles that have that key as the article mark. For instance, articles with `gnus-unread-mark` in the example above will not get adaptive score entries.

Each article can have only one mark, so just a single of these rules will be applied to each article.

To take `gnus-del-mark` as an example—this alist says that all articles that have that mark (i.e., are marked with 'D') will have a score entry added to lower based on the `From` header by -4, and lowered by `Subject` by -1. Change this to fit your prejudices.

If you have marked 10 articles with the same subject with `gnus-del-mark`, the rule for that mark will be applied ten times. That means that that subject will get a score of ten times -1, which should be, unless I'm much mistaken, -10.

If you have auto-expirable (mail) groups (see Section 6.3.9 [Expiring Mail], page 128), all the read articles will be marked with the 'E' mark. This'll probably make adaptive scoring slightly impossible, so auto-expiring and adaptive scoring doesn't really mix very well.

The headers you can score on are `from`, `subject`, `message-id`, `references`, `xref`, `lines`, `chars` and `date`. In addition, you can score on `followup`, which will create an adaptive score entry that matches on the `References` header using the `Message-ID` of the current article, thereby matching the following thread.

You can also score on `thread`, which will try to score all articles that appear in a thread. `thread` matches uses a `Message-ID` to match on the `References` header of the article. If the match is made, the `Message-ID` of the article is added to the `thread` rule. (Think about it. I'd recommend two aspirins afterwards.)

If you use this scheme, you should set the score file atom `mark` to something small—like -300, perhaps, to avoid having small random changes result in articles getting marked as read.

After using adaptive scoring for a week or so, Gnus should start to become properly trained and enhance the authors you like best, and kill the authors you like least, without you having to say so explicitly.

You can control what groups the adaptive scoring is to be performed on by using the score files (see Section 7.4 [Score File Format], page 172). This will also let you use different rules in different groups.

The adaptive score entries will be put into a file where the name is the group name with `gnus-adaptive-file-suffix` appended. The default is 'ADAPT'.

When doing adaptive scoring, substring or fuzzy matching would probably give you the best results in most cases. However, if the header one matches is short, the possibility for

false positives is great, so if the length of the match is less than `gnus-score-exact-adapt-limit`, exact matching will be used. If this variable is `nil`, exact matching will always be used to avoid this problem.

As mentioned above, you can adapt either on individual words or entire headers. If you adapt on words, the `gnus-default-adaptive-word-score-alist` variable says what score each instance of a word should add given a mark.

```
(setq gnus-default-adaptive-word-score-alist
      '((,gnus-read-mark . 30)
        (,gnus-catchup-mark . -10)
        (,gnus-killed-mark . -20)
        (,gnus-del-mark . -15)))
```

This is the default value. If you have adaption on words enabled, every word that appears in subjects of articles marked with `gnus-read-mark` will result in a score rule that increase the score with 30 points.

Words that appear in the `gnus-default-ignored-adaptive-words` list will be ignored. If you wish to add more words to be ignored, use the `gnus-ignored-adaptive-words` list instead.

When the scoring is done, `gnus-adaptive-word-syntax-table` is the syntax table in effect. It is similar to the standard syntax table, but it considers numbers to be non-word-constituent characters.

If `gnus-adaptive-word-minimum` is set to a number, the adaptive word scoring process will never bring down the score of an article to below this number. The default is `nil`.

If `gnus-adaptive-word-no-group-words` is set to `t`, gnus won't adaptively word score any of the words in the group name. Useful for groups like '`comp.editors.emacs`', where most of the subject lines contain the word '`emacs`'.

After using this scheme for a while, it might be nice to write a `gnus-psychoanalyze-user` command to go through the rules and see what words you like and what words you don't like. Or perhaps not.

Note that the adaptive word scoring thing is highly experimental and is likely to change in the future. Initial impressions seem to indicate that it's totally useless as it stands. Some more work (involving more rigorous statistical methods) will have to be done to make this useful.

7.7 Home Score File

The score file where new score file entries will go is called the *home score file*. This is normally (and by default) the score file for the group itself. For instance, the home score file for '`gnu.emacs.gnus`' is '`gnu.emacs.gnus.SCORE`'.

However, this may not be what you want. It is often convenient to share a common home score file among many groups—all '`emacs`' groups could perhaps use the same home score file.

The variable that controls this is `gnus-home-score-file`. It can be:

1. A string. Then this file will be used as the home score file for all groups.
2. A function. The result of this function will be used as the home score file. The function will be called with the name of the group as the parameter.

3. A list. The elements in this list can be:

    1. (*regexp file-name*). If the *regexp* matches the group name, the *file-name* will be used as the home score file.

    2. A function. If the function returns non-nil, the result will be used as the home score file.

    3. A string. Use the string as the home score file.

    The list will be traversed from the beginning towards the end looking for matches.

So, if you want to use just a single score file, you could say:

```
(setq gnus-home-score-file
      "my-total-score-file.SCORE")
```

If you want to use 'gnu.SCORE' for all 'gnu' groups and 'rec.SCORE' for all 'rec' groups (and so on), you can say:

```
(setq gnus-home-score-file
      'gnus-hierarchial-home-score-file)
```

This is a ready-made function provided for your convenience. Other functions include

`gnus-current-home-score-file`
> Return the "current" regular score file. This will make scoring commands add entry to the "innermost" matching score file.

If you want to have one score file for the 'emacs' groups and another for the 'comp' groups, while letting all other groups use their own home score files:

```
(setq gnus-home-score-file
      ;; All groups that match the regexp "\\.emacs"
      '(("\\.emacs" "emacs.SCORE")
        ;; All the comp groups in one score file
        ("^comp" "comp.SCORE")))
```

`gnus-home-adapt-file` works exactly the same way as `gnus-home-score-file`, but says what the home adaptive score file is instead. All new adaptive file entries will go into the file specified by this variable, and the same syntax is allowed.

In addition to using `gnus-home-score-file` and `gnus-home-adapt-file`, you can also use group parameters (see Section 2.10 [Group Parameters], page 22) and topic parameters (see Section 2.16.5 [Topic Parameters], page 32) to achieve much the same. Group and topic parameters take precedence over this variable.

7.8 Followups To Yourself

Gnus offers two commands for picking out the `Message-ID` header in the current buffer. Gnus will then add a score rule that scores using this `Message-ID` on the `References` header of other articles. This will, in effect, increase the score of all articles that respond to the article in the current buffer. Quite useful if you want to easily note when people answer what you've said.

`gnus-score-followup-article`
> This will add a score to articles that directly follow up your own article.

`gnus-score-followup-thread`

> This will add a score to all articles that appear in a thread "below" your own article.

These two functions are both primarily meant to be used in hooks like `message-sent-hook`, like this:

```
(add-hook 'message-sent-hook 'gnus-score-followup-thread)
```

If you look closely at your own `Message-ID`, you'll notice that the first two or three characters are always the same. Here's two of mine:

```
<x6u3u47icf.fsf@eyesore.no>
<x6sp9o7ibw.fsf@eyesore.no>
```

So "my" ident on this machine is 'x6'. This can be exploited—the following rule will raise the score on all followups to myself:

```
("references"
 ("<x6[0-9a-z]+\\.fsf\\(_-_\\)?@.*eyesore\\.no>"
  1000 nil r))
```

Whether it's the first two or first three characters that are "yours" is system-dependent.

7.9 Scoring Tips

Crossposts

> If you want to lower the score of crossposts, the line to match on is the `Xref` header.
>
> ```
> ("xref" (" talk.politics.misc:" -1000))
> ```

Multiple crossposts

> If you want to lower the score of articles that have been crossposted to more than, say, 3 groups:
>
> ```
> ("xref"
>   ("[^:\n]+:[0-9]+ +[^:\n]+:[0-9]+ +[^:\n]+:[0-9]+"
>    -1000 nil r))
> ```

Matching on the body

> This is generally not a very good idea—it takes a very long time. Gnus actually has to fetch each individual article from the server. But you might want to anyway, I guess. Even though there are three match keys (`Head`, `Body` and `All`), you should choose one and stick with it in each score file. If you use any two, each article will be fetched *twice*. If you want to match a bit on the `Head` and a bit on the `Body`, just use `All` for all the matches.

Marking as read

> You will probably want to mark articles that have scores below a certain number as read. This is most easily achieved by putting the following in your 'all.SCORE' file:
>
> ```
> ((mark -100))
> ```
>
> You may also consider doing something similar with `expunge`.

Negated character classes

>   If you say stuff like `[^abcd]*`, you may get unexpected results. That will match
>   newlines, which might lead to, well, The Unknown. Say `[^abcd\n]*` instead.

7.10 Reverse Scoring

If you want to keep just articles that have '`Sex with Emacs`' in the subject header, and
expunge all other articles, you could put something like this in your score file:

```
(("subject"
  ("Sex with Emacs" 2))
 (mark 1)
 (expunge 1))
```

So, you raise all articles that match '`Sex with Emacs`' and mark the rest as read, and
expunge them to boot.

7.11 Global Score Files

Sure, other newsreaders have "global kill files". These are usually nothing more than a
single kill file that applies to all groups, stored in the user's home directory. Bah! Puny,
weak newsreaders!

What I'm talking about here are Global Score Files. Score files from all over the world,
from users everywhere, uniting all nations in one big, happy score file union! Ange-score!
New and untested!

All you have to do to use other people's score files is to set the `gnus-global-score-files` variable. One entry for each score file, or each score file directory. Gnus will decide
by itself what score files are applicable to which group.

To use the score file '`/ftp@ftp.gnus.org:/pub/larsi/ding/score/soc.motss.SCORE`'
and all score files in the '`/ftp@ftp.some-where:/pub/score`' directory, say this:

```
(setq gnus-global-score-files
      '("/ftp@ftp.gnus.org:/pub/larsi/ding/score/soc.motss.SCORE"
        "/ftp@ftp.some-where:/pub/score/"))
```

Simple, eh? Directory names must end with a '`/`'. These directories are typically scanned
only once during each Gnus session. If you feel the need to manually re-scan the remote
directories, you can use the `gnus-score-search-global-directories` command.

Note that, at present, using this option will slow down group entry somewhat. (That
is—a lot.)

If you want to start maintaining score files for other people to use, just put your score file
up for anonymous ftp and announce it to the world. Become a retro-moderator! Participate
in the retro-moderator wars sure to ensue, where retro-moderators battle it out for the
sympathy of the people, luring them to use their score files on false premises! Yay! The net
is saved!

Here are some tips for the would-be retro-moderator, off the top of my head:

- Articles heavily crossposted are probably junk.
- To lower a single inappropriate article, lower by `Message-ID`.

- Particularly brilliant authors can be raised on a permanent basis.
- Authors that repeatedly post off-charter for the group can safely be lowered out of existence.
- Set the `mark` and `expunge` atoms to obliterate the nastiest articles completely.
- Use expiring score entries to keep the size of the file down. You should probably have a long expiry period, though, as some sites keep old articles for a long time.

... I wonder whether other newsreaders will support global score files in the future. *Snicker.* Yup, any day now, newsreaders like Blue Wave, xrn and 1stReader are bound to implement scoring. Should we start holding our breath yet?

7.12 Kill Files

Gnus still supports those pesky old kill files. In fact, the kill file entries can now be expiring, which is something I wrote before Daniel Quinlan thought of doing score files, so I've left the code in there.

In short, kill processing is a lot slower (and I do mean *a lot*) than score processing, so it might be a good idea to rewrite your kill files into score files.

Anyway, a kill file is a normal `emacs-lisp` file. You can put any forms into this file, which means that you can use kill files as some sort of primitive hook function to be run on group entry, even though that isn't a very good idea.

Normal kill files look like this:

```
(gnus-kill "From" "Lars Ingebrigtsen")
(gnus-kill "Subject" "ding")
(gnus-expunge "X")
```

This will mark every article written by me as read, and remove the marked articles from the summary buffer. Very useful, you'll agree.

Other programs use a totally different kill file syntax. If Gnus encounters what looks like a `rn` kill file, it will take a stab at interpreting it.

Two summary functions for editing a GNUS kill file:

M-k   Edit this group's kill file (`gnus-summary-edit-local-kill`).

M-K   Edit the general kill file (`gnus-summary-edit-global-kill`).

Two group mode functions for editing the kill files:

M-k   Edit this group's kill file (`gnus-group-edit-local-kill`).

M-K   Edit the general kill file (`gnus-group-edit-global-kill`).

Kill file variables:

`gnus-kill-file-name`

A kill file for the group 'soc.motss' is normally called 'soc.motss.KILL'. The suffix appended to the group name to get this file name is detailed by the `gnus-kill-file-name` variable. The "global" kill file (not in the score file sense of "global", of course) is just called 'KILL'.

`gnus-kill-save-kill-file`

> If this variable is non-`nil`, Gnus will save the kill file after processing, which is necessary if you use expiring kills.

`gnus-apply-kill-hook`

> A hook called to apply kill files to a group. It is (`gnus-apply-kill-file`) by default. If you want to ignore the kill file if you have a score file for the same group, you can set this hook to (`gnus-apply-kill-file-unless-scored`). If you don't want kill files to be processed, you should set this variable to `nil`.

`gnus-kill-file-mode-hook`

> A hook called in kill-file mode buffers.

7.13 Converting Kill Files

If you have loads of old kill files, you may want to convert them into score files. If they are "regular", you can use the '`gnus-kill-to-score.el`' package; if not, you'll have to do it by hand.

The kill to score conversion package isn't included in Gnus by default. You can fetch it from `http://www.stud.ifi.uio.no/~larsi/ding-various/gnus-kill-to-score.el`.

If your old kill files are very complex—if they contain more non-`gnus-kill` forms than not, you'll have to convert them by hand. Or just let them be as they are. Gnus will still use them as before.

7.14 GroupLens

GroupLens is a collaborative filtering system that helps you work together with other people to find the quality news articles out of the huge volume of news articles generated every day.

To accomplish this the GroupLens system combines your opinions about articles you have already read with the opinions of others who have done likewise and gives you a personalized prediction for each unread news article. Think of GroupLens as a matchmaker. GroupLens watches how you rate articles, and finds other people that rate articles the same way. Once it has found some people you agree with it tells you, in the form of a prediction, what they thought of the article. You can use this prediction to help you decide whether or not you want to read the article.

7.14.1 Using GroupLens

To use GroupLens you must register a pseudonym with your local Better Bit Bureau (BBB). `http://www.cs.umn.edu/Research/GroupLens/bbb.html` is the only better bit in town at the moment.

Once you have registered you'll need to set a couple of variables.

`gnus-use-grouplens`

> Setting this variable to a non-`nil` value will make Gnus hook into all the relevant GroupLens functions.

`grouplens-pseudonym`
>   This variable should be set to the pseudonym you got when registering with the Better Bit Bureau.

`grouplens-newsgroups`
>   A list of groups that you want to get GroupLens predictions for.

That's the minimum of what you need to get up and running with GroupLens. Once you've registered, GroupLens will start giving you scores for articles based on the average of what other people think. But, to get the real benefit of GroupLens you need to start rating articles yourself. Then the scores GroupLens gives you will be personalized for you, based on how the people you usually agree with have already rated.

7.14.2 Rating Articles

In GroupLens, an article is rated on a scale from 1 to 5, inclusive. Where 1 means something like this article is a waste of bandwidth and 5 means that the article was really good. The basic question to ask yourself is, "on a scale from 1 to 5 would I like to see more articles like this one?"

There are four ways to enter a rating for an article in GroupLens.

`r`   This function will prompt you for a rating on a scale of one to five.

`k`   This function will prompt you for a rating, and rate all the articles in the thread. This is really useful for some of those long running giant threads in rec.humor.

The next two commands, `n` and `,` take a numerical prefix to be the score of the article you're reading.

`1-5 n`   Rate the article and go to the next unread article.

`1-5 ,`   Rate the article and go to the next unread article with the highest score.

If you want to give the current article a score of 4 and then go to the next article, just type `4 n`.

7.14.3 Displaying Predictions

GroupLens makes a prediction for you about how much you will like a news article. The predictions from GroupLens are on a scale from 1 to 5, where 1 is the worst and 5 is the best. You can use the predictions from GroupLens in one of three ways controlled by the variable `gnus-grouplens-override-scoring`.

There are three ways to display predictions in grouplens. You may choose to have the GroupLens scores contribute to, or override the regular gnus scoring mechanism. override is the default; however, some people prefer to see the Gnus scores plus the grouplens scores. To get the separate scoring behavior you need to set `gnus-grouplens-override-scoring` to `'separate`. To have the GroupLens predictions combined with the grouplens scores set it to `'override` and to combine the scores set `gnus-grouplens-override-scoring` to `'combine`. When you use the combine option you will also want to set the values for `grouplens-prediction-offset` and `grouplens-score-scale-factor`.

In either case, GroupLens gives you a few choices for how you would like to see your predictions displayed. The display of predictions is controlled by the `grouplens-prediction-display` variable.

The following are valid values for that variable.

`prediction-spot`
> The higher the prediction, the further to the right an '`*`' is displayed.

`confidence-interval`
> A numeric confidence interval.

`prediction-bar`
> The higher the prediction, the longer the bar.

`confidence-bar`
> Numerical confidence.

`confidence-spot`
> The spot gets bigger with more confidence.

`prediction-num`
> Plain-old numeric value.

`confidence-plus-minus`
> Prediction +/- confidence.

7.14.4 GroupLens Variables

`gnus-summary-grouplens-line-format`
> The summary line format used in GroupLens-enhanced summary buffers. It accepts the same specs as the normal summary line format (see Section 3.1.1 [Summary Buffer Lines], page 37). The default is '`%U%R%z%l%I%(%[%4L: %-20,20n%]%) %s\n`'.

`grouplens-bbb-host`
> Host running the bbbd server. '`grouplens.cs.umn.edu`' is the default.

`grouplens-bbb-port`
> Port of the host running the bbbd server. The default is 9000.

`grouplens-score-offset`
> Offset the prediction by this value. In other words, subtract the prediction value by this number to arrive at the effective score. The default is 0.

`grouplens-score-scale-factor`
> This variable allows the user to magnify the effect of GroupLens scores. The scale factor is applied after the offset. The default is 1.

7.15 Advanced Scoring

Scoring on Subjects and From headers is nice enough, but what if you're really interested in what a person has to say only when she's talking about a particular subject? Or what if

you really don't want to read what person A has to say when she's following up to person B, but want to read what she says when she's following up to person C?

By using advanced scoring rules you may create arbitrarily complex scoring patterns.

7.15.1 Advanced Scoring Syntax

Ordinary scoring rules have a string as the first element in the rule. Advanced scoring rules have a list as the first element. The second element is the score to be applied if the first element evaluated to a non-`nil` value.

These lists may consist of three logical operators, one redirection operator, and various match operators.

Logical operators:

`&`
`and`　　　　This logical operator will evaluate each of its arguments until it finds one that evaluates to `false`, and then it'll stop. If all arguments evaluate to `true` values, then this operator will return `true`.

`|`
`or`　　　　This logical operator will evaluate each of its arguments until it finds one that evaluates to `true`. If no arguments are `true`, then this operator will return `false`.

`!`
`not`
　　　　　　This logical operator only takes a single argument. It returns the logical negation of the value of its argument.

There is an *indirection operator* that will make its arguments apply to the ancestors of the current article being scored. For instance, `1-` will make score rules apply to the parent of the current article. `2-` will make score rules apply to the grandparent of the current article. Alternatively, you can write `^^`, where the number of `^`s (carets) says how far back into the ancestry you want to go.

Finally, we have the match operators. These are the ones that do the real work. Match operators are header name strings followed by a match and a match type. A typical match operator looks like '`("from" "Lars Ingebrigtsen" s)`'. The header names are the same as when using simple scoring, and the match types are also the same.

7.15.2 Advanced Scoring Examples

Let's say you want to increase the score of articles written by Lars when he's talking about Gnus:

```
((&
  ("from" "Lars Ingebrigtsen")
  ("subject" "Gnus"))
 1000)
```

Quite simple, huh?

When he writes long articles, he sometimes has something nice to say:

```
((&
  ("from" "Lars Ingebrigtsen")
  (|
   ("subject" "Gnus")
   ("lines" 100 >)))
 1000)
```

However, when he responds to things written by Reig Eigil Logge, you really don't want to read what he's written:

```
((&
  ("from" "Lars Ingebrigtsen")
  (1- ("from" "Reig Eigir Logge")))
 -100000)
```

Everybody that follows up Redmondo when he writes about disappearing socks should have their scores raised, but only when they talk about white socks. However, when Lars talks about socks, it's usually not very interesting:

```
((&
  (1-
   (&
    ("from" "redmondo@.*no" r)
    ("body" "disappearing.*socks" t)))
  (! ("from" "Lars Ingebrigtsen"))
  ("body" "white.*socks"))
 1000)
```

The possibilities are endless.

7.15.3 Advanced Scoring Tips

The & and | logical operators do short-circuit logic. That is, they stop processing their arguments when it's clear what the result of the operation will be. For instance, if one of the arguments of an & evaluates to `false`, there's no point in evaluating the rest of the arguments. This means that you should put slow matches ('`body`', '`header`') last and quick matches ('`from`', '`subject`') first.

The indirection arguments (`1-` and so on) will make their arguments work on previous generations of the thread. If you say something like:

```
...
(1-
 (1-
  ("from" "lars")))
...
```

Then that means "score on the from header of the grandparent of the current article". An indirection is quite fast, but it's better to say:

```
(1-
 (&
  ("from" "Lars")
  ("subject" "Gnus")))
```

than it is to say:

```
(&
 (1- ("from" "Lars"))
 (1- ("subject" "Gnus")))
```

7.16 Score Decays

You may find that your scores have a tendency to grow without bounds, especially if you're using adaptive scoring. If scores get too big, they lose all meaning—they simply max out and it's difficult to use them in any sensible way.

Gnus provides a mechanism for decaying scores to help with this problem. When score files are loaded and `gnus-decay-scores` is non-`nil`, Gnus will run the score files through the decaying mechanism thereby lowering the scores of all non-permanent score rules. The decay itself if performed by the `gnus-decay-score-function` function, which is `gnus-decay-score` by default. Here's the definition of that function:

```
(defun gnus-decay-score (score)
  "Decay SCORE.
This is done according to 'gnus-score-decay-constant'
and 'gnus-score-decay-scale'."
  (floor
   (- score
      (* (if (< score 0) 1 -1)
         (min (abs score)
              (max gnus-score-decay-constant
                   (* (abs score)
                      gnus-score-decay-scale)))))))
```

`gnus-score-decay-constant` is 3 by default and `gnus-score-decay-scale` is 0.05. This should cause the following:

1. Scores between -3 and 3 will be set to 0 when this function is called.

2. Scores with magnitudes between 3 and 60 will be shrunk by 3.

3. Scores with magnitudes greater than 60 will be shrunk by 5% of the score.

If you don't like this decay function, write your own. It is called with the score to be decayed as its only parameter, and it should return the new score, which should be an integer.

Gnus will try to decay scores once a day. If you haven't run Gnus for four days, Gnus will decay the scores four times, for instance.

8 Various

8.1 Process/Prefix

Many functions, among them functions for moving, decoding and saving articles, use what is known as the *Process/Prefix convention*.

This is a method for figuring out what articles the user wants the command to be performed on.

It goes like this:

If the numeric prefix is N, perform the operation on the next N articles, starting with the current one. If the numeric prefix is negative, perform the operation on the previous N articles, starting with the current one.

If `transient-mark-mode` in non-`nil` and the region is active, all articles in the region will be worked upon.

If there is no numeric prefix, but some articles are marked with the process mark, perform the operation on the articles marked with the process mark.

If there is neither a numeric prefix nor any articles marked with the process mark, just perform the operation on the current article.

Quite simple, really, but it needs to be made clear so that surprises are avoided.

Commands that react to the process mark will push the current list of process marked articles onto a stack and will then clear all process marked articles. You can restore the previous configuration with the `M P y` command (see Section 3.6.6 [Setting Process Marks], page 51).

One thing that seems to shock & horrify lots of people is that, for instance, `3 d` does exactly the same as `d d d`. Since each `d` (which marks the current article as read) by default goes to the next unread article after marking, this means that `3 d` will mark the next three unread articles as read, no matter what the summary buffer looks like. Set `gnus-summary-goto-unread` to `nil` for a more straightforward action.

Many commands do not use the process/prefix convention. All commands that do explicitly say so in this manual. To apply the process/prefix convention to commands that do not use it, you can use the `M-&` command. For instance, to mark all the articles in the group as expirable, you could say `M P b M-& E`.

8.2 Interactive

`gnus-novice-user`

> If this variable is non-`nil`, you are either a newcomer to the World of Usenet, or you are very cautious, which is a nice thing to be, really. You will be given questions of the type "Are you sure you want to do this?" before doing anything dangerous. This is `t` by default.

`gnus-expert-user`

> If this variable is non-`nil`, you will seldom be asked any questions by Gnus. It will simply assume you know what you're doing, no matter how strange.

`gnus-interactive-catchup`
>   Require confirmation before catching up a group if non-`nil`. It is `t` by default.

`gnus-interactive-exit`
>   Require confirmation before exiting Gnus. This variable is `t` by default.

8.3 Symbolic Prefixes

Quite a lot of Emacs commands react to the (numeric) prefix. For instance, `C-u 4 C-f` moves point four characters forward, and `C-u 9 0 0 I s s p` adds a permanent `Subject` substring score rule of 900 to the current article.

This is all nice and well, but what if you want to give a command some additional information? Well, what most commands do is interpret the "raw" prefix in some special way. `C-u 0 C-x C-s` means that one doesn't want a backup file to be created when saving the current buffer, for instance. But what if you want to save without making a backup file, and you want Emacs to flash lights and play a nice tune at the same time? You can't, and you're probably perfectly happy that way.

I'm not, so I've added a second prefix—the *symbolic prefix*. The prefix key is `M-i` (`gnus-symbolic-argument`), and the next character typed in is the value. You can stack as many `M-i` prefixes as you want. `M-i a M-C-u` means "feed the `M-C-u` command the symbolic prefix `a`". `M-i a M-i b M-C-u` means "feed the `M-C-u` command the symbolic prefixes `a` and `b`". You get the drift.

Typing in symbolic prefixes to commands that don't accept them doesn't hurt, but it doesn't do any good either. Currently not many Gnus functions make use of the symbolic prefix.

If you're interested in how Gnus implements this, see Section 10.7.7 [Extended Interactive], page 251.

8.4 Formatting Variables

Throughout this manual you've probably noticed lots of variables called things like `gnus-group-line-format` and `gnus-summary-mode-line-format`. These control how Gnus is to output lines in the various buffers. There's quite a lot of them. Fortunately, they all use the same syntax, so there's not that much to be annoyed by.

Here's an example format spec (from the group buffer): '`%M%S%5y: %(%g%)\n`'. We see that it is indeed extremely ugly, and that there are lots of percentages everywhere.

Currently Gnus uses the following formatting variables: `gnus-group-line-format`, `gnus-summary-line-format`, `gnus-server-line-format`, `gnus-topic-line-format`, `gnus-group-mode-line-format`, `gnus-summary-mode-line-format`, `gnus-article-mode-line-format`, `gnus-server-mode-line-format`, and `gnus-summary-pick-line-format`.

All these format variables can also be arbitrary elisp forms. In that case, they will be `eval`ed to insert the required lines.

Gnus includes a command to help you while creating your own format specs. `M-x gnus-update-format` will `eval` the current form, update the spec in question and pop you to a buffer where you can examine the resulting lisp code to be run to generate the line.

8.4.1 Formatting Basics

Each '%' element will be replaced by some string or other when the buffer in question is generated. '%5y' means "insert the 'y' spec, and pad with spaces to get a 5-character field".

As with normal C and Emacs Lisp formatting strings, the numerical modifier between the '%' and the formatting type character will *pad* the output so that it is always at least that long. '%5y' will make the field always (at least) five characters wide by padding with spaces to the left. If you say '%-5y', it will pad to the right instead.

You may also wish to limit the length of the field to protect against particularly wide values. For that you can say '%4,6y', which means that the field will never be more than 6 characters wide and never less than 4 characters wide.

8.4.2 Mode Line Formatting

Mode line formatting variables (e.g., `gnus-summary-mode-line-format`) follow the same rules as other, buffer line oriented formatting variables (see Section 8.4.1 [Formatting Basics], page 191) with the following two differences:

1. There must be no newline ('\n') at the end.

2. The special '%%b' spec can be used to display the buffer name. Well, it's no spec at all, really—'%%' is just a way to quote '%' to allow it to pass through the formatting machinery unmangled, so that Emacs receives '%b', which is something the Emacs mode line display interprets to mean "show the buffer name". For a full list of mode line specs Emacs understands, see the documentation of the `mode-line-format` variable.

8.4.3 Advanced Formatting

It is frequently useful to post-process the fields in some way. Padding, limiting, cutting off parts and suppressing certain values can be achieved by using *tilde modifiers*. A typical tilde spec might look like '%~(cut 3)~(ignore "0")y'.

These are the valid modifiers:

pad
pad-left   Pad the field to the left with spaces until it reaches the required length.

pad-right
            Pad the field to the right with spaces until it reaches the required length.

max
max-left   Cut off characters from the left until it reaches the specified length.

max-right
            Cut off characters from the right until it reaches the specified length.

cut
cut-left   Cut off the specified number of characters from the left.

cut-right
            Cut off the specified number of characters from the right.

ignore   Return an empty string if the field is equal to the specified value.

`form`   Use the specified form as the field value when the '`@`' spec is used.

Let's take an example. The '`%o`' spec in the summary mode lines will return a date in compact ISO8601 format—'`19960809T230410`'. This is quite a mouthful, so we want to shave off the century number and the time, leaving us with a six-character date. That would be '`%~(cut-left 2)~(max-right 6)~(pad 6)o`'. (Cutting is done before maxing, and we need the padding to ensure that the date is never less than 6 characters to make it look nice in columns.)

Ignoring is done first; then cutting; then maxing; and then as the very last operation, padding.

If you use lots of these advanced thingies, you'll find that Gnus gets quite slow. This can be helped enormously by running `M-x gnus-compile` when you are satisfied with the look of your lines. See Section 8.7 [Compilation], page 197.

8.4.4 User-Defined Specs

All the specs allow for inserting user defined specifiers—'`u`'. The next character in the format string should be a letter. Gnus will call the function `gnus-user-format-function-`'`X`', where '`X`' is the letter following '`%u`'. The function will be passed a single parameter—what the parameter means depends on what buffer it's being called from. The function should return a string, which will be inserted into the buffer just like information from any other specifier. This function may also be called with dummy values, so it should protect against that.

You can also use tilde modifiers (see Section 8.4.3 [Advanced Formatting], page 191 to achieve much the same without defining new functions. Here's an example: '`%~(form (count-lines (point-min) (point)))@`'. The form given here will be evaluated to yield the current line number, and then inserted.

8.4.5 Formatting Fonts

There are specs for highlighting, and these are shared by all the format variables. Text inside the '`%(`' and '`%)`' specifiers will get the special `mouse-face` property set, which means that it will be highlighted (with `gnus-mouse-face`) when you put the mouse pointer over it.

Text inside the '`%{`' and '`%}`' specifiers will have their normal faces set using `gnus-face-0`, which is `bold` by default. If you say '`%1{`', you'll get `gnus-face-1` instead, and so on. Create as many faces as you wish. The same goes for the `mouse-face` specs—you can say '`%3(hello%)`' to have '`hello`' mouse-highlighted with `gnus-mouse-face-3`.

Text inside the '`%<`' and '`%>`' specifiers will get the special `balloon-help` property set to `gnus-balloon-face-0`. If you say '`%1<`', you'll get `gnus-balloon-face-1` and so on. The `gnus-balloon-face-*` variables should be either strings or symbols naming functions that return a string. Under `balloon-help-mode`, when the mouse passes over text with this property set, a balloon window will appear and display the string. Please refer to the doc string of `balloon-help-mode` for more information on this.

Here's an alternative recipe for the group buffer:

```
;; Create three face types.
(setq gnus-face-1 'bold)
(setq gnus-face-3 'italic)

;; We want the article count to be in
;; a bold and green face.  So we create
;; a new face called 'my-green-bold'.
(copy-face 'bold 'my-green-bold)
;; Set the color.
(set-face-foreground 'my-green-bold "ForestGreen")
(setq gnus-face-2 'my-green-bold)

;; Set the new & fancy format.
(setq gnus-group-line-format
      "%M%S%3{%5y%}%2[:%] %(%1{%g%}%)\n")
```

I'm sure you'll be able to use this scheme to create totally unreadable and extremely vulgar displays. Have fun!

Note that the '%(' specs (and friends) do not make any sense on the mode-line variables.

8.5 Windows Configuration

No, there's nothing here about X, so be quiet.

If `gnus-use-full-window` non-`nil`, Gnus will delete all other windows and occupy the entire Emacs screen by itself. It is `t` by default.

Setting this variable to `nil` kinda works, but there are glitches. Use at your own peril.

`gnus-buffer-configuration` describes how much space each Gnus buffer should be given. Here's an excerpt of this variable:

```
((group (vertical 1.0 (group 1.0 point)
                      (if gnus-carpal (group-carpal 4))))
 (article (vertical 1.0 (summary 0.25 point)
                        (article 1.0))))
```

This is an alist. The *key* is a symbol that names some action or other. For instance, when displaying the group buffer, the window configuration function will use `group` as the key. A full list of possible names is listed below.

The *value* (i.e., the *split*) says how much space each buffer should occupy. To take the `article` split as an example -

```
(article (vertical 1.0 (summary 0.25 point)
                       (article 1.0)))
```

This *split* says that the summary buffer should occupy 25% of upper half of the screen, and that it is placed over the article buffer. As you may have noticed, 100% + 25% is actually 125% (yup, I saw y'all reaching for that calculator there). However, the special number `1.0` is used to signal that this buffer should soak up all the rest of the space available after the rest of the buffers have taken whatever they need. There should be only one buffer with the `1.0` size spec per split.

Point will be put in the buffer that has the optional third element `point`. In a `frame` split, the last subsplit having a leaf split where the tag `frame-focus` is a member (i.e. is

the third or fourth element in the list, depending on whether the `point` tag is present) gets focus.

Here's a more complicated example:

```
(article (vertical 1.0 (group 4)
                       (summary 0.25 point)
                       (if gnus-carpal (summary-carpal 4))
                       (article 1.0)))
```

If the size spec is an integer instead of a floating point number, then that number will be used to say how many lines a buffer should occupy, not a percentage.

If the *split* looks like something that can be `eval`ed (to be precise—if the `car` of the split is a function or a subr), this split will be `eval`ed. If the result is non-`nil`, it will be used as a split. This means that there will be three buffers if `gnus-carpal` is `nil`, and four buffers if `gnus-carpal` is non-`nil`.

Not complicated enough for you? Well, try this on for size:

```
(article (horizontal 1.0
             (vertical 0.5
                 (group 1.0)
                 (gnus-carpal 4))
             (vertical 1.0
                 (summary 0.25 point)
                 (summary-carpal 4)
                 (article 1.0))))
```

Whoops. Two buffers with the mystery 100% tag. And what's that `horizontal` thingie?

If the first element in one of the split is `horizontal`, Gnus will split the window horizontally, giving you two windows side-by-side. Inside each of these strips you may carry on all you like in the normal fashion. The number following `horizontal` says what percentage of the screen is to be given to this strip.

For each split, there *must* be one element that has the 100% tag. The splitting is never accurate, and this buffer will eat any leftover lines from the splits.

To be slightly more formal, here's a definition of what a valid split may look like:

```
split       = frame | horizontal | vertical | buffer | form
frame       = "(frame " size *split ")"
horizontal  = "(horizontal " size *split ")"
vertical    = "(vertical " size *split ")"
buffer      = "(" buf-name " " size *[ "point" ] *[ "frame-focus"] ")"
size        = number | frame-params
buf-name    = group | article | summary ...
```

The limitations are that the `frame` split can only appear as the top-level split. *form* should be an Emacs Lisp form that should return a valid split. We see that each split is fully recursive, and may contain any number of `vertical` and `horizontal` splits.

Finding the right sizes can be a bit complicated. No window may be less than `gnus-window-min-height` (default 1) characters high, and all windows must be at least `gnus-window-min-width` (default 1) characters wide. Gnus will try to enforce this before applying the splits. If you want to use the normal Emacs window width/height limit, you can just set these two variables to `nil`.

If you're not familiar with Emacs terminology, `horizontal` and `vertical` splits may work the opposite way of what you'd expect. Windows inside a `horizontal` split are shown side-by-side, and windows within a `vertical` split are shown above each other.

If you want to experiment with window placement, a good tip is to call `gnus-configure-frame` directly with a split. This is the function that does all the real work when splitting buffers. Below is a pretty nonsensical configuration with 5 windows; two for the group buffer and three for the article buffer. (I said it was nonsensical.) If you `eval` the statement below, you can get an idea of how that would look straight away, without going through the normal Gnus channels. Play with it until you're satisfied, and then use `gnus-add-configuration` to add your new creation to the buffer configuration list.

```
(gnus-configure-frame
 '(horizontal 1.0
    (vertical 10
      (group 1.0)
      (article 0.3 point))
    (vertical 1.0
      (article 1.0)
      (horizontal 4
        (group 1.0)
        (article 10)))))
```

You might want to have several frames as well. No prob—just use the `frame` split:

```
(gnus-configure-frame
 '(frame 1.0
         (vertical 1.0
                   (summary 0.25 point frame-focus)
                   (article 1.0))
         (vertical ((height . 5) (width . 15)
                    (user-position . t)
                    (left . -1) (top . 1))
                   (picon 1.0))))
```

This split will result in the familiar summary/article window configuration in the first (or "main") frame, while a small additional frame will be created where picons will be shown. As you can see, instead of the normal `1.0` top-level spec, each additional split should have a frame parameter alist as the size spec. See section "Frame Parameters" in *The GNU Emacs Lisp Reference Manual*. Under XEmacs, a frame property list will be accepted, too—for instance, (`height 5 width 15 left -1 top 1`) is such a plist. The list of all possible keys for `gnus-buffer-configuration` can be found in its default value.

Note that the `message` key is used for both `gnus-group-mail` and `gnus-summary-mail-other-window`. If it is desirable to distinguish between the two, something like this might be used:

```
(message (horizontal 1.0
                     (vertical 1.0 (message 1.0 point))
                     (vertical 0.24
                               (if (buffer-live-p gnus-summary-buffer)
                                   '(summary 0.5))
                               (group 1.0)))))
```

One common desire for a multiple frame split is to have a separate frame for composing mail and news while leaving the original frame intact. To accomplish that, something like the following can be done:

```
(message
  (frame 1.0
          (if (not (buffer-live-p gnus-summary-buffer))
              (car (cdr (assoc 'group gnus-buffer-configuration)))
            (car (cdr (assoc 'summary gnus-buffer-configuration))))
          (vertical ((user-position . t) (top . 1) (left . 1)
                      (name . "Message"))
                      (message 1.0 point))))
```

Since the `gnus-buffer-configuration` variable is so long and complicated, there's a function you can use to ease changing the config of a single setting: `gnus-add-configuration`. If, for instance, you want to change the `article` setting, you could say:

```
(gnus-add-configuration
 '(article (vertical 1.0
                  (group 4)
                  (summary .25 point)
                  (article 1.0))))
```

You'd typically stick these `gnus-add-configuration` calls in your '`.gnus.el`' file or in some startup hook—they should be run after Gnus has been loaded.

If all windows mentioned in the configuration are already visible, Gnus won't change the window configuration. If you always want to force the "right" window configuration, you can set `gnus-always-force-window-configuration` to non-`nil`.

If you're using tree displays (see Section 3.23 [Tree Display], page 84), and the tree window is displayed vertically next to another window, you may also want to fiddle with `gnus-tree-minimize-window` to avoid having the windows resized.

8.5.1 Example Window Configurations

- Narrow left hand side occupied by group buffer. Right hand side split between summary buffer (top one-sixth) and article buffer (bottom).

```
(gnus-add-configuration
 '(article
   (horizontal 1.0
                (vertical 25 (group 1.0))
                (vertical 1.0
                          (summary 0.16 point)
                          (article 1.0)))))

(gnus-add-configuration
 '(summary
   (horizontal 1.0
                (vertical 25 (group 1.0))
                (vertical 1.0 (summary 1.0 point)))))
```

8.6 Faces and Fonts

Fiddling with fonts and faces used to be very difficult, but these days it is very simple. You simply say *M-x customize-face*, pick out the face you want to alter, and alter it via the standard Customize interface.

8.7 Compilation

Remember all those line format specification variables? `gnus-summary-line-format`, `gnus-group-line-format`, and so on. Now, Gnus will of course heed whatever these variables are, but, unfortunately, changing them will mean a quite significant slow-down. (The default values of these variables have byte-compiled functions associated with them, while the user-generated versions do not, of course.)

To help with this, you can run *M-x gnus-compile* after you've fiddled around with the variables and feel that you're (kind of) satisfied. This will result in the new specs being byte-compiled, and you'll get top speed again. Gnus will save these compiled specs in the '`.newsrc.eld`' file. (User-defined functions aren't compiled by this function, though—you should compile them yourself by sticking them into the `.gnus.el` file and byte-compiling that file.)

8.8 Mode Lines

`gnus-updated-mode-lines` says what buffers should keep their mode lines updated. It is a list of symbols. Supported symbols include `group`, `article`, `summary`, `server`, `browse`, and `tree`. If the corresponding symbol is present, Gnus will keep that mode line updated with information that may be pertinent. If this variable is `nil`, screen refresh may be quicker.

By default, Gnus displays information on the current article in the mode lines of the summary and article buffers. The information Gnus wishes to display (e.g. the subject of the article) is often longer than the mode lines, and therefore have to be cut off at some point. The `gnus-mode-non-string-length` variable says how long the other elements on the line is (i.e., the non-info part). If you put additional elements on the mode line (e.g. a clock), you should modify this variable:

```
(add-hook 'display-time-hook
          (lambda () (setq gnus-mode-non-string-length
                           (+ 21
                              (if line-number-mode 5 0)
                              (if column-number-mode 4 0)
                              (length display-time-string)))))
```

If this variable is `nil` (which is the default), the mode line strings won't be chopped off, and they won't be padded either. Note that the default is unlikely to be desirable, as even the percentage complete in the buffer may be crowded off the mode line; the user should configure this variable appropriately for her configuration.

8.9 Highlighting and Menus

The `gnus-visual` variable controls most of the Gnus-prettifying aspects. If `nil`, Gnus won't attempt to create menus or use fancy colors or fonts. This will also inhibit loading the 'gnus-vis.el' file.

This variable can be a list of visual properties that are enabled. The following elements are valid, and are all included by default:

`group-highlight`
> Do highlights in the group buffer.

`summary-highlight`
> Do highlights in the summary buffer.

`article-highlight`
> Do highlights in the article buffer.

`highlight`
> Turn on highlighting in all buffers.

`group-menu`
> Create menus in the group buffer.

`summary-menu`
> Create menus in the summary buffers.

`article-menu`
> Create menus in the article buffer.

`browse-menu`
> Create menus in the browse buffer.

`server-menu`
> Create menus in the server buffer.

`score-menu`
> Create menus in the score buffers.

`menu`   Create menus in all buffers.

So if you only want highlighting in the article buffer and menus in all buffers, you could say something like:

    (setq gnus-visual '(article-highlight menu))

If you want highlighting only and no menus whatsoever, you'd say:

    (setq gnus-visual '(highlight))

If `gnus-visual` is `t`, highlighting and menus will be used in all Gnus buffers.

Other general variables that influence the look of all buffers include:

`gnus-mouse-face`
> This is the face (i.e., font) used for mouse highlighting in Gnus. No mouse highlights will be done if `gnus-visual` is `nil`.

There are hooks associated with the creation of all the different menus:

```
gnus-article-menu-hook
```
    Hook called after creating the article mode menu.

```
gnus-group-menu-hook
```
    Hook called after creating the group mode menu.

```
gnus-summary-menu-hook
```
    Hook called after creating the summary mode menu.

```
gnus-server-menu-hook
```
    Hook called after creating the server mode menu.

```
gnus-browse-menu-hook
```
    Hook called after creating the browse mode menu.

```
gnus-score-menu-hook
```
    Hook called after creating the score mode menu.

8.10 Buttons

Those new-fangled *mouse* contraptions are very popular with the young, hep kids who don't want to learn the proper way to do things these days. Why, I remember way back in the summer of '89, when I was using Emacs on a Tops 20 system. Three hundred users on one single machine, and every user was running Simula compilers. Bah!

Right.

Well, you can make Gnus display buffers full of buttons you can click to do anything by setting `gnus-carpal` to `t`. Pretty simple, really. Tell the chiropractor I sent you.

```
gnus-carpal-mode-hook
```
    Hook run in all carpal mode buffers.

```
gnus-carpal-button-face
```
    Face used on buttons.

```
gnus-carpal-header-face
```
    Face used on carpal buffer headers.

```
gnus-carpal-group-buffer-buttons
```
    Buttons in the group buffer.

```
gnus-carpal-summary-buffer-buttons
```
    Buttons in the summary buffer.

```
gnus-carpal-server-buffer-buttons
```
    Buttons in the server buffer.

```
gnus-carpal-browse-buffer-buttons
```
    Buttons in the browse buffer.

All the `buttons` variables are lists. The elements in these list are either cons cells where the `car` contains a text to be displayed and the `cdr` contains a function symbol, or a simple string.

8.11 Daemons

Gnus, being larger than any program ever written (allegedly), does lots of strange stuff that you may wish to have done while you're not present. For instance, you may want it to check for new mail once in a while. Or you may want it to close down all connections to all servers when you leave Emacs idle. And stuff like that.

Gnus will let you do stuff like that by defining various *handlers*. Each handler consists of three elements: A *function*, a *time*, and an *idle* parameter.

Here's an example of a handler that closes connections when Emacs has been idle for thirty minutes:

```
(gnus-demon-close-connections nil 30)
```

Here's a handler that scans for PGP headers every hour when Emacs is idle:

```
(gnus-demon-scan-pgp 60 t)
```

This *time* parameter and than *idle* parameter work together in a strange, but wonderful fashion. Basically, if *idle* is `nil`, then the function will be called every *time* minutes.

If *idle* is `t`, then the function will be called after *time* minutes only if Emacs is idle. So if Emacs is never idle, the function will never be called. But once Emacs goes idle, the function will be called every *time* minutes.

If *idle* is a number and *time* is a number, the function will be called every *time* minutes only when Emacs has been idle for *idle* minutes.

If *idle* is a number and *time* is `nil`, the function will be called once every time Emacs has been idle for *idle* minutes.

And if *time* is a string, it should look like '`07:31`', and the function will then be called once every day somewhere near that time. Modified by the *idle* parameter, of course.

(When I say "minute" here, I really mean `gnus-demon-timestep` seconds. This is 60 by default. If you change that variable, all the timings in the handlers will be affected.)

So, if you want to add a handler, you could put something like this in your '`.gnus`' file:

```
(gnus-demon-add-handler 'gnus-demon-close-connections 30 t)
```

Some ready-made functions to do this have been created: `gnus-demon-add-nocem`, `gnus-demon-add-disconnection`, `gnus-demon-add-nntp-close-connection`, `gnus-demon-add-scan-timestamps`, `gnus-demon-add-rescan`, and `gnus-demon-add-scanmail`. Just put those functions in your '`.gnus`' if you want those abilities.

If you add handlers to `gnus-demon-handlers` directly, you should run `gnus-demon-init` to make the changes take hold. To cancel all daemons, you can use the `gnus-demon-cancel` function.

Note that adding daemons can be pretty naughty if you over do it. Adding functions that scan all news and mail from all servers every two seconds is a sure-fire way of getting booted off any respectable system. So behave.

8.12 NoCeM

Spamming is posting the same article lots and lots of times. Spamming is bad. Spamming is evil.

Spamming is usually canceled within a day or so by various anti-spamming agencies. These agencies usually also send out *NoCeM* messages. NoCeM is pronounced "no see-'em", and means what the name implies—these are messages that make the offending articles, like, go away.

What use are these NoCeM messages if the articles are canceled anyway? Some sites do not honor cancel messages and some sites just honor cancels from a select few people. Then you may wish to make use of the NoCeM messages, which are distributed in the 'alt.nocem.misc' newsgroup.

Gnus can read and parse the messages in this group automatically, and this will make spam disappear.

There are some variables to customize, of course:

gnus-use-nocem

>   Set this variable to `t` to set the ball rolling. It is `nil` by default.

gnus-nocem-groups

>   Gnus will look for NoCeM messages in the groups in this list. The default is `("news.lists.filters" "news.admin.net-abuse.bulletins" "alt.nocem.misc" "news.admin.net-abuse.announce")`.

gnus-nocem-issuers

>   There are many people issuing NoCeM messages. This list says what people you want to listen to. The default is `("Automoose-1" "clewis@ferret.ocunix.on.ca" "cosmo.roadkill" "SpamHippo" "hweede@snafu.de")`; fine, upstanding citizens all of them.

>   Known despammers that you can put in this list are listed at `http://www.xs4all.nl/~rosalind/nocemreg/nocemreg.html`.

>   You do not have to heed NoCeM messages from all these people—just the ones you want to listen to. You also don't have to accept all NoCeM messages from the people you like. Each NoCeM message has a *type* header that gives the message a (more or less, usually less) rigorous definition. Common types are 'spam', 'spew', 'mmf', 'binary', and 'troll'. To specify this, you have to use (*issuer conditions* ...) elements in the list. Each condition is either a string (which is a regexp that matches types you want to use) or a list on the form (not *string*), where *string* is a regexp that matches types you don't want to use.

>   For instance, if you want all NoCeM messages from Chris Lewis except his 'troll' messages, you'd say:

>   `("clewis@ferret.ocunix.on.ca" ".*" (not "troll"))`

>   On the other hand, if you just want nothing but his 'spam' and 'spew' messages, you'd say:

>   `("clewis@ferret.ocunix.on.ca" (not ".*") "spew" "spam")`

>   The specs are applied left-to-right.

gnus-nocem-verifyer

>   This should be a function for verifying that the NoCeM issuer is who she says she is. The default is `mc-verify`, which is a Mailcrypt function. If this is too

slow and you don't care for verification (which may be dangerous), you can set this variable to `nil`.

If you want signed NoCeM messages to be verified and unsigned messages not to be verified (but used anyway), you could do something like:

```
(setq gnus-nocem-verifyer 'my-gnus-mc-verify)

(defun my-gnus-mc-verify ()
  (not (eq 'forged
           (ignore-errors
             (if (mc-verify)
                 t
               'forged)))))
```

This might be dangerous, though.

`gnus-nocem-directory`

> This is where Gnus will store its NoCeM cache files. The default is '`~/News/NoCeM/`'.

`gnus-nocem-expiry-wait`

> The number of days before removing old NoCeM entries from the cache. The default is 15. If you make it shorter Gnus will be faster, but you might then see old spam.

`gnus-nocem-check-from`

> Non-`nil` means check for valid issuers in message bodies. Otherwise don't bother fetching articles unless their author matches a valid issuer; that is much faster if you are selective about the issuers.

`gnus-nocem-check-article-limit`

> If non-`nil`, the maximum number of articles to check in any NoCeM group. NoCeM groups can be huge and very slow to process.

Using NoCeM could potentially be a memory hog. If you have many living (i. e., subscribed or unsubscribed groups), your Emacs process will grow big. If this is a problem, you should kill off all (or most) of your unsubscribed groups (see Section 2.4 [Subscription Commands], page 18).

8.13 Undo

It is very useful to be able to undo actions one has done. In normal Emacs buffers, it's easy enough—you just push the `undo` button. In Gnus buffers, however, it isn't that simple.

The things Gnus displays in its buffer is of no value whatsoever to Gnus—it's all just data designed to look nice to the user. Killing a group in the group buffer with `C-k` makes the line disappear, but that's just a side-effect of the real action—the removal of the group in question from the internal Gnus structures. Undoing something like that can't be done by the normal Emacs `undo` function.

Gnus tries to remedy this somewhat by keeping track of what the user does and coming up with actions that would reverse the actions the user takes. When the user then presses the `undo` key, Gnus will run the code to reverse the previous action, or the previous actions.

However, not all actions are easily reversible, so Gnus currently offers a few key functions to be undoable. These include killing groups, yanking groups, and changing the list of read articles of groups. That's it, really. More functions may be added in the future, but each added function means an increase in data to be stored, so Gnus will never be totally undoable.

The undoability is provided by the `gnus-undo-mode` minor mode. It is used if `gnus-use-undo` is non-`nil`, which is the default. The `M-C-_` key performs the `gnus-undo` command, which should feel kinda like the normal Emacs `undo` command.

8.14 Moderation

If you are a moderator, you can use the '`gnus-mdrtn.el`' package. It is not included in the standard Gnus package. Write a mail to '`larsi@gnus.org`' and state what group you moderate, and you'll get a copy.

The moderation package is implemented as a minor mode for summary buffers. Put

```
(add-hook 'gnus-summary-mode-hook 'gnus-moderate)
```

in your '`.gnus.el`' file.

If you are the moderator of '`rec.zoofle`', this is how it's supposed to work:

1. You split your incoming mail by matching on '`Newsgroups:.*rec.zoofle`', which will put all the to-be-posted articles in some mail group—for instance, '`nnml:rec.zoofle`'.

2. You enter that group once in a while and post articles using the `e` (edit-and-post) or `s` (just send unedited) commands.

3. If, while reading the '`rec.zoofle`' newsgroup, you happen upon some articles that weren't approved by you, you can cancel them with the `c` command.

To use moderation mode in these two groups, say:

```
(setq gnus-moderated-list
      "^nnml:rec.zoofle$\\|^rec.zoofle$")
```

8.15 Emacs Enhancements

Starting with version 21, Emacs is able to display pictures and stuff, so Gnus has taken advantage of that.

Gnus-specific tool bars will be used if Tool Bar mode is on. Currently the group, summary and message buffers have tool bars defined.

MIME image types may be displayed internally if Emacs was built with appropriate support (see variable `image-types`). 'X-Face' headers may be rendered as images internally if you have appropriate support programs (see [X-Face], page 75). You can play sounds internally if Emacs was built with suitable audio support; otherwise Gnus will attempt to play sounds externally.

A simplified version of the XEmacs Smiley support for *emoticons* (see Section 8.16.2 [Smileys], page 207) is available on graphical displays under the control of `gnus-treat-display-smileys`. Text 'smiley' faces—':-)', ':-/', ':-(' and the like—are mapped to pictures which are displayed instead. The mapping is controlled by a list of regexps

`smiley-regexp-alist` mapping matched text to image file names. It contains matches for 'smiley', 'wry' and 'frowny' by default.

There is currently no Emacs support for 'Picons' (see Section 8.16.1 [Picons], page 204), but there is no reason why it couldn't be added.

8.16 XEmacs Enhancements

XEmacs is able to display pictures and stuff, so Gnus has taken advantage of that.

8.16.1 Picons

So. . . You want to slow down your news reader even more! This is a good way to do so. Its also a great way to impress people staring over your shoulder as you read news.

8.16.1.1 Picon Basics

What are Picons? To quote directly from the Picons Web site:

> *Picons* is short for "personal icons". They're small, constrained images used to represent users and domains on the net, organized into databases so that the appropriate image for a given e-mail address can be found. Besides users and domains, there are picon databases for Usenet newsgroups and weather forecasts. The picons are in either monochrome `XBM` format or color `XPM` and `GIF` formats.

If you have a permanent connection to the Internet you can use Steve Kinzler's Picons Search engine by setting `gnus-picons-piconsearch-url` to the string `http://www.cs.indiana.edu/picons/search.html`.

Otherwise you need a local copy of his database. For instructions on obtaining and installing the picons databases, point your Web browser at `http://www.cs.indiana.edu/picons/ftp/index.html`. Gnus expects picons to be installed into a location pointed to by `gnus-picons-database`.

8.16.1.2 Picon Requirements

To have Gnus display Picons for you, you must be running XEmacs 19.13 or greater since all other versions of Emacs aren't yet able to display images.

Additionally, you must have `x` support compiled into XEmacs. To display color picons which are much nicer than the black & white one, you also need one of `xpm` or `gif` compiled into XEmacs.

If you want to display faces from `X-Face` headers, you should have the `xface` support compiled into XEmacs. Otherwise you must have the `netpbm` utilities installed, or munge the `gnus-picons-convert-x-face` variable to use something else.

8.16.1.3 Easy Picons

To enable displaying picons, simply put the following line in your '`~/.gnus`' file and start Gnus.

```
(setq gnus-use-picons t)
(setq gnus-treat-display-picons t)
```

and make sure `gnus-picons-database` points to the directory containing the Picons databases.

Alternatively if you want to use the web piconsearch engine add this:

```
(setq gnus-picons-piconsearch-url
      "http://www.cs.indiana.edu:800/piconsearch")
```

8.16.1.4 Hard Picons

Gnus can display picons for you as you enter and leave groups and articles. It knows how to interact with three sections of the picons database. Namely, it can display the picons newsgroup pictures, author's face picture(s), and the authors domain. To enable this feature, you need to select where to get the picons from, and where to display them.

`gnus-picons-database`
> The location of the picons database. Should point to a directory containing the 'news', 'domains', 'users' (and so on) subdirectories. This is only useful if `gnus-picons-piconsearch-url` is `nil`. Defaults to '`/usr/local/faces/`'.

`gnus-picons-piconsearch-url`
> The URL for the web picons search engine. The only currently known engine is `http://www.cs.indiana.edu:800/piconsearch`. To workaround network delays, icons will be fetched in the background. If this is `nil` 'the default), then picons are fetched from local database indicated by `gnus-picons-database`.

`gnus-picons-display-where`
> Where the picon images should be displayed. It is `picons` by default (which by default maps to the buffer '`*Picons*`'). Other valid places could be `article`, `summary`, or '`*scratch*`' for all I care. Just make sure that you've made the buffer visible using the standard Gnus window configuration routines—see Section 8.5 [Windows Configuration], page 193.

`gnus-picons-group-excluded-groups`
> Groups that are matched by this regexp won't have their group icons displayed.

Note: If you set `gnus-use-picons` to `t`, it will set up your window configuration for you to include the `picons` buffer.

Now that you've made those decision, you need to add the following functions to the appropriate hooks so these pictures will get displayed at the right time.

`gnus-article-display-picons`
> Looks up and displays the picons for the author and the author's domain in the `gnus-picons-display-where` buffer.

`gnus-picons-article-display-x-face`
> Decodes and displays the X-Face header if present.

8.16.1.5 Picon Useless Configuration

The following variables offer further control over how things are done, where things are located, and other useless stuff you really don't need to worry about.

gnus-picons-news-directories
>   List of subdirectories to search in `gnus-picons-database` for newsgroups faces. (`"news"`) is the default.

gnus-picons-user-directories
>   List of subdirectories to search in `gnus-picons-database` for user faces. (`"local"` `"users"` `"usenix"` `"misc"`) is the default.

gnus-picons-domain-directories
>   List of subdirectories to search in `gnus-picons-database` for domain name faces. Defaults to (`"domains"`). Some people may want to add '`"unknown"`' to this list.

gnus-picons-convert-x-face
>   If you don't have `xface` support builtin XEmacs, this is the command to use to convert the `X-Face` header to an X bitmap (`xbm`). Defaults to (`format "{ echo '/* Width=48, Height=48 */'; uncompface; } | icontopbm | pbmtoxbm > %s"` `gnus-picons-x-face-file-name`)

gnus-picons-x-face-file-name
>   Names a temporary file to store the `X-Face` bitmap in. Defaults to (`format "/tmp/picon-xface.%s.xbm"` (`user-login-name`)).

gnus-picons-has-modeline-p
>   If you have set `gnus-picons-display-where` to `picons`, your XEmacs frame will become really cluttered. To alleviate this a bit you can set `gnus-picons-has-modeline-p` to `nil`; this will remove the mode line from the Picons buffer. This is only useful if `gnus-picons-display-where` is `picons`.

gnus-picons-refresh-before-display
>   If non-nil, display the article buffer before computing the picons. Defaults to `nil`.

gnus-picons-display-as-address
>   If `t` display textual email addresses along with pictures. Defaults to `t`.

gnus-picons-file-suffixes
>   Ordered list of suffixes on picon file names to try. Defaults to (`"xpm"` `"gif"` `"xbm"`) minus those not builtin your XEmacs.

gnus-picons-setup-hook
>   Hook run in the picon buffer, if that is displayed.

gnus-picons-display-article-move-p
>   Whether to move point to first empty line when displaying picons. This has only an effect if 'gnus-picons-display-where' has value 'article'.
>
>   If `nil`, display the picons in the `From` and `Newsgroups` lines. This is the default.

`gnus-picons-clear-cache-on-shutdown`
> Whether to clear the picons cache when exiting gnus. Gnus caches every picons it finds while it is running. This saves some time in the search process but eats some memory. If this variable is set to `nil`, Gnus will never clear the cache itself; you will have to manually call `gnus-picons-clear-cache` to clear it. Otherwise the cache will be cleared every time you exit Gnus. Defaults to `t`.

8.16.2 Smileys

Smiley is a package separate from Gnus, but since Gnus is currently the only package that uses Smiley, it is documented here.

In short—to use Smiley in Gnus, put the following in your '`.gnus.el`' file:

```
(setq gnus-treat-display-smileys t)
```

Smiley maps text smiley faces—'`:-)`', '`:-=`', '`:-(`' and the like—to pictures and displays those instead of the text smiley faces. The conversion is controlled by a list of regexps that matches text and maps that to file names.

Smiley supplies two example conversion alists by default: `smiley-deformed-regexp-alist` (which matches '`:)`', '`:(`' and so on), and `smiley-nosey-regexp-alist` (which matches '`:-)`', '`:-(`' and so on).

The alist used is specified by the `smiley-regexp-alist` variable, which defaults to the value of `smiley-deformed-regexp-alist`.

The first item in each element is the regexp to be matched; the second element is the regexp match group that is to be replaced by the picture; and the third element is the name of the file to be displayed.

The following variables customize where Smiley will look for these files, as well as the color to be used and stuff:

`smiley-data-directory`
> Where Smiley will look for smiley faces files.

`smiley-flesh-color`
> Skin color. The default is '`yellow`', which is really racist.

`smiley-features-color`
> Color of the features of the face. The default is '`black`'.

`smiley-tongue-color`
> Color of the tongue. The default is '`red`'.

`smiley-circle-color`
> Color of the circle around the face. The default is '`black`'.

`smiley-mouse-face`
> Face used for mouse highlighting over the smiley face.

8.16.3 Toolbar

`gnus-use-toolbar`
> If `nil`, don't display toolbars. If non-`nil`, it should be one of `default-toolbar`, `top-toolbar`, `bottom-toolbar`, `right-toolbar`, or `left-toolbar`.

gnus-group-toolbar
            The toolbar in the group buffer.

gnus-summary-toolbar
            The toolbar in the summary buffer.

gnus-summary-mail-toolbar
            The toolbar in the summary buffer of mail groups.

8.16.4 Various XEmacs Variables

gnus-xmas-glyph-directory
            This is where Gnus will look for pictures. Gnus will normally auto-detect
            this directory, but you may set it manually if you have an unusual directory
            structure.

gnus-xmas-logo-color-alist
            This is an alist where the key is a type symbol and the values are the foreground
            and background color of the splash page glyph.

gnus-xmas-logo-color-style
            This is the key used to look up the color in the alist described above. Valid
            values include `flame`, `pine`, `moss`, `irish`, `sky`, `tin`, `velvet`, `grape`, `labia`,
            `berry`, `neutral`, and `september`.

gnus-xmas-modeline-glyph
            A glyph displayed in all Gnus mode lines. It is a tiny gnu head by default.

8.17 Fuzzy Matching

Gnus provides *fuzzy matching* of `Subject` lines when doing things like scoring, thread
gathering and thread comparison.

As opposed to regular expression matching, fuzzy matching is very fuzzy. It's so fuzzy
that there's not even a definition of what *fuzziness* means, and the implementation has
changed over time.

Basically, it tries to remove all noise from lines before comparing. '`Re: `', parenthetical
remarks, white space, and so on, are filtered out of the strings before comparing the re-
sults. This often leads to adequate results—even when faced with strings generated by text
manglers masquerading as newsreaders.

8.18 Thwarting Email Spam

In these last days of the Usenet, commercial vultures are hanging about and grepping
through news like crazy to find email addresses they can foist off their scams and products
to. As a reaction to this, many people have started putting nonsense addresses into their
`From` lines. I think this is counterproductive—it makes it difficult for people to send you
legitimate mail in response to things you write, as well as making it difficult to see who wrote
what. This rewriting may perhaps be a bigger menace than the unsolicited commercial email
itself in the end.

The biggest problem I have with email spam is that it comes in under false pretenses. I press *g* and Gnus merrily informs me that I have 10 new emails. I say "Golly gee! Happy is me!" and select the mail group, only to find two pyramid schemes, seven advertisements ("New! Miracle tonic for growing full, lustrous hair on your toes!") and one mail asking me to repent and find some god.

This is annoying.

The way to deal with this is having Gnus split out all spam into a 'spam' mail group (see Section 6.3.3 [Splitting Mail], page 115).

First, pick one (1) valid mail address that you can be reached at, and put it in your `From` header of all your news articles. (I've chosen '`larsi@trym.ifi.uio.no`', but for many addresses on the form '`larsi+usenet@ifi.uio.no`' will be a better choice. Ask your sysadmin whether your sendmail installation accepts keywords in the local part of the mail address.)

```
(setq message-default-news-headers
      "From: Lars Magne Ingebrigtsen <larsi@trym.ifi.uio.no>\n")
```

Then put the following split rule in `nnmail-split-fancy` (see Section 6.3.6 [Fancy Mail Splitting], page 124):

```
(
 ...
 (to "larsi@trym.ifi.uio.no"
     (| ("subject" "re:.*" "misc")
        ("references" ".*@.*" "misc")
        "spam"))
 ...
)
```

This says that all mail to this address is suspect, but if it has a `Subject` that starts with a '`Re:`' or has a `References` header, it's probably ok. All the rest goes to the 'spam' group. (This idea probably comes from Tim Pierce.)

In addition, many mail spammers talk directly to your `smtp` server and do not include your email address explicitly in the `To` header. Why they do this is unknown—perhaps it's to thwart this thwarting scheme? In any case, this is trivial to deal with—you just put anything not addressed to you in the 'spam' group by ending your fancy split rule in this way:

```
(
 ...
 (to "larsi" "misc")
 "spam")
```

In my experience, this will sort virtually everything into the right group. You still have to check the 'spam' group from time to time to check for legitimate mail, though. If you feel like being a good net citizen, you can even send off complaints to the proper authorities on each unsolicited commercial email—at your leisure.

If you are also a lazy net citizen, you will probably prefer complaining automatically with the '`gnus-junk.el`' package, available as free software at `http://stud2.tuwien.ac.at/~e9426626/gnus-junk.html`. Since most e-mail spam is sent automatically, this may reconcile the cosmic balance somewhat.

This works for me. It allows people an easy way to contact me (they can just press `r` in the usual way), and I'm not bothered at all with spam. It's a win-win situation. Forging `From` headers to point to non-existent domains is yucky, in my opinion.

8.19 Various Various

`gnus-home-directory`

> All Gnus path variables will be initialized from this variable, which defaults to '`~/`'.

`gnus-directory`

> Most Gnus storage path variables will be initialized from this variable, which defaults to the '`SAVEDIR`' environment variable, or '`~/News/`' if that variable isn't set.

> Note that Gnus is mostly loaded when the '`.gnus.el`' file is read. This means that other directory variables that are initialized from this variable won't be set properly if you set this variable in '`.gnus.el`'. Set this variable in '`.emacs`' instead.

`gnus-default-directory`

> Not related to the above variable at all—this variable says what the default directory of all Gnus buffers should be. If you issue commands like `C-x C-f`, the prompt you'll get starts in the current buffer's default directory. If this variable is `nil` (which is the default), the default directory will be the default directory of the buffer you were in when you started Gnus.

`gnus-verbose`

> This variable is an integer between zero and ten. The higher the value, the more messages will be displayed. If this variable is zero, Gnus will never flash any messages, if it is seven (which is the default), most important messages will be shown, and if it is ten, Gnus won't ever shut up, but will flash so many messages it will make your head swim.

`gnus-verbose-backends`

> This variable works the same way as `gnus-verbose`, but it applies to the Gnus back ends instead of Gnus proper.

`nnheader-max-head-length`

> When the back ends read straight heads of articles, they all try to read as little as possible. This variable (default 4096) specifies the absolute max length the back ends will try to read before giving up on finding a separator line between the head and the body. If this variable is `nil`, there is no upper read bound. If it is `t`, the back ends won't try to read the articles piece by piece, but read the entire articles. This makes sense with some versions of `ange-ftp` or `efs`.

`nnheader-head-chop-length`

> This variable (default 2048) says how big a piece of each article to read when doing the operation described above.

`nnheader-file-name-translation-alist`

> This is an alist that says how to translate characters in file names. For instance, if ':' is invalid as a file character in file names on your system (you OS/2 user you), you could say something like:
>
> ```
> (setq nnheader-file-name-translation-alist
>       '((?: . ?_)))
> ```
>
> In fact, this is the default value for this variable on OS/2 and MS Windows (phooey) systems.

`gnus-hidden-properties`

> This is a list of properties to use to hide "invisible" text. It is `(invisible t intangible t)` by default on most systems, which makes invisible text invisible and intangible.

`gnus-parse-headers-hook`

> A hook called before parsing headers. It can be used, for instance, to gather statistics on the headers fetched, or perhaps you'd like to prune some headers. I don't see why you'd want that, though.

`gnus-shell-command-separator`

> String used to separate two shell commands. The default is ';'.

`gnus-invalid-group-regexp`

> Regexp to match "invalid" group names when querying user for a group name. The default value catches some **really** invalid group names who could possibly mess up Gnus internally (like allowing ':' in a group name, which is normally used to delimit method and group).
>
> IMAP users might want to allow '/' in group names though.

9 The End

Well, that's the manual—you can get on with your life now. Keep in touch. Say hello to your cats from me.

My **ghod**—I just can't stand goodbyes. Sniffle.

Ol' Charles Reznikoff said it pretty well, so I leave the floor to him:

Te Deum

Not because of victories
I sing,
having none,
but for the common sunshine,
the breeze,
the largess of the spring.

Not for victory
but for the day's work done
as well as I was able;
not for a seat upon the dais
but at the common table.

10 Appendices

10.1 History

GNUS was written by Masanobu UMEDA. When autumn crept up in '94, Lars Magne Ingebrigtsen grew bored and decided to rewrite Gnus.

If you want to investigate the person responsible for this outrage, you can point your (feh!) web browser to `http://quimby.gnus.org/`. This is also the primary distribution point for the new and spiffy versions of Gnus, and is known as The Site That Destroys Newsrcs And Drives People Mad.

During the first extended alpha period of development, the new Gnus was called "(ding) Gnus". *(ding)* is, of course, short for *ding is not Gnus*, which is a total and utter lie, but who cares? (Besides, the "Gnus" in this abbreviation should probably be pronounced "news" as UMEDA intended, which makes it a more appropriate name, don't you think?)

In any case, after spending all that energy on coming up with a new and spunky name, we decided that the name was *too* spunky, so we renamed it back again to "Gnus". But in mixed case. "Gnus" vs. "GNUS". New vs. old.

10.1.1 Gnus Versions

The first "proper" release of Gnus 5 was done in November 1995 when it was included in the Emacs 19.30 distribution (132 (ding) Gnus releases plus 15 Gnus 5.0 releases).

In May 1996 the next Gnus generation (aka. "September Gnus" (after 99 releases)) was released under the name "Gnus 5.2" (40 releases).

On July 28th 1996 work on Red Gnus was begun, and it was released on January 25th 1997 (after 84 releases) as "Gnus 5.4" (67 releases).

On September 13th 1997, Quassia Gnus was started and lasted 37 releases. If was released as "Gnus 5.6" on March 8th 1998 (46 releases).

Gnus 5.6 begat Pterodactyl Gnus on August 29th 1998 and was released as "Gnus 5.8" (after 99 releases and a CVS repository) on December 3rd 1999.

If you happen upon a version of Gnus that has a prefixed name – "(ding) Gnus", "September Gnus", "Red Gnus", "Quassia Gnus" – don't panic. Don't let it know that you're frightened. Back away. Slowly. Whatever you do, don't run. Walk away, calmly, until you're out of its reach. Find a proper released version of Gnus and snuggle up to that instead.

10.1.2 Other Gnus Versions

In addition to the versions of Gnus which have had their releases coordinated by Lars, one major development has been Semi-gnus from Japan. It's based on a library called SEMI, which provides MIME capabilities.

These Gnusae are based mainly on Gnus 5.6 and Pterodactyl Gnus. Collectively, they are called "Semi-gnus", and different strains are called T-gnus, ET-gnus, Nana-gnus and Chaos. These provide powerful MIME and multilingualization things, especially important for Japanese users.

10.1.3 Why?

What's the point of Gnus?

I want to provide a "rad", "happening", "way cool" and "hep" newsreader, that lets you do anything you can think of. That was my original motivation, but while working on Gnus, it has become clear to me that this generation of newsreaders really belong in the stone age. Newsreaders haven't developed much since the infancy of the net. If the volume continues to rise with the current rate of increase, all current newsreaders will be pretty much useless. How do you deal with newsgroups that have thousands of new articles each day? How do you keep track of millions of people who post?

Gnus offers no real solutions to these questions, but I would very much like to see Gnus being used as a testing ground for new methods of reading and fetching news. Expanding on UMEDA-san's wise decision to separate the newsreader from the back ends, Gnus now offers a simple interface for anybody who wants to write new back ends for fetching mail and news from different sources. I have added hooks for customizations everywhere I could imagine it being useful. By doing so, I'm inviting every one of you to explore and invent.

May Gnus never be complete. `C-u 100 M-x all-hail-emacs` and `C-u 100 M-x all-hail-xemacs`.

10.1.4 Compatibility

Gnus was designed to be fully compatible with GNUS. Almost all key bindings have been kept. More key bindings have been added, of course, but only in one or two obscure cases have old bindings been changed.

Our motto is:

> In a cloud bones of steel.

All commands have kept their names. Some internal functions have changed their names.

The `gnus-uu` package has changed drastically. See Section 3.15 [Decoding Articles], page 65.

One major compatibility question is the presence of several summary buffers. All variables relevant while reading a group are buffer-local to the summary buffer they belong in. Although many important variables have their values copied into their global counterparts whenever a command is executed in the summary buffer, this change might lead to incorrect values being used unless you are careful.

All code that relies on knowledge of GNUS internals will probably fail. To take two examples: Sorting `gnus-newsrc-alist` (or changing it in any way, as a matter of fact) is strictly verboten. Gnus maintains a hash table that points to the entries in this alist (which speeds up many functions), and changing the alist directly will lead to peculiar results.

Old hilit19 code does not work at all. In fact, you should probably remove all hilit code from all Gnus hooks (`gnus-group-prepare-hook` and `gnus-summary-prepare-hook`). Gnus provides various integrated functions for highlighting. These are faster and more accurate. To make life easier for everybody, Gnus will by default remove all hilit calls from all hilit hooks. Uncleanliness! Away!

Packages like `expire-kill` will no longer work. As a matter of fact, you should probably remove all old GNUS packages (and other code) when you start using Gnus. More likely than not, Gnus already does what you have written code to make GNUS do. (Snicker.)

Even though old methods of doing things are still supported, only the new methods are documented in this manual. If you detect a new method of doing something while reading this manual, that does not mean you have to stop doing it the old way.

Gnus understands all GNUS startup files.

Overall, a casual user who hasn't written much code that depends on GNUS internals should suffer no problems. If problems occur, please let me know by issuing that magic command *M-x gnus-bug*.

If you are in the habit of sending bug reports *very* often, you may find the helpful help buffer annoying after a while. If so, set `gnus-bug-create-help-buffer` to `nil` to avoid having it pop up at you.

10.1.5 Conformity

No rebels without a clue here, ma'am. We conform to all standards known to (wo)man. Except for those standards and/or conventions we disagree with, of course.

RFC 822   There are no known breaches of this standard.

RFC 1036   There are no known breaches of this standard, either.

Son-of-RFC 1036

> We do have some breaches to this one.

> > *X-Newsreader*
> > *User-Agent*

> > > These are considered to be "vanity headers", while I consider them to be consumer information. After seeing so many badly formatted articles coming from `tin` and `Netscape` I know not to use either of those for posting articles. I would not have known that if it wasn't for the `X-Newsreader` header.

USEFOR   USEFOR is an IETF working group writing a successor to RFC 1036, based on Son-of-RFC 1036. They have produced a number of drafts proposing various changes to the format of news articles. The Gnus towers will look into implementing the changes when the draft is accepted as an RFC.

If you ever notice Gnus acting non-compliant with regards to the texts mentioned above, don't hesitate to drop a note to Gnus Towers and let us know.

10.1.6 Emacsen

Gnus should work on :
- Emacs 20.3 and up.
- XEmacs 20.4 and up.

This Gnus version will absolutely not work on any Emacsen older than that. Not reliably, at least. Older versions of Gnus may work on older Emacs versions.

There are some vague differences between Gnus on the various platforms—XEmacs features more graphics (a logo and a toolbar)—but other than that, things should look pretty much the same under all Emacsen.

10.1.7 Gnus Development

Gnus is developed in a two-phased cycle. The first phase involves much discussion on the 'ding@gnus.org' mailing list, where people propose changes and new features, post patches and new back ends. This phase is called the *alpha* phase, since the Gnusae released in this phase are *alpha releases*, or (perhaps more commonly in other circles) *snapshots*. During this phase, Gnus is assumed to be unstable and should not be used by casual users. Gnus alpha releases have names like "Red Gnus" and "Quassia Gnus".

After futzing around for 50-100 alpha releases, Gnus is declared *frozen*, and only bug fixes are applied. Gnus loses the prefix, and is called things like "Gnus 5.6.32" instead. Normal people are supposed to be able to use these, and these are mostly discussed on the 'gnu.emacs.gnus' newsgroup.

Some variable defaults differ between alpha Gnusae and released Gnusae. In particular, `mail-source-delete-incoming` defaults to `nil` in alpha Gnusae and `t` in released Gnusae. This is to prevent lossage of mail if an alpha release hiccups while handling the mail.

The division of discussion between the ding mailing list and the Gnus newsgroup is not purely based on publicity concerns. It's true that having people write about the horrible things that an alpha Gnus release can do (sometimes) in a public forum may scare people off, but more importantly, talking about new experimental features that have been introduced may confuse casual users. New features are frequently introduced, fiddled with, and judged to be found wanting, and then either discarded or totally rewritten. People reading the mailing list usually keep up with these rapid changes, while people on the newsgroup can't be assumed to do so.

10.1.8 Contributors

The new Gnus version couldn't have been done without the help of all the people on the (ding) mailing list. Every day for over a year I have gotten billions of nice bug reports from them, filling me with joy, every single one of them. Smooches. The people on the list have been tried beyond endurance, what with my "oh, that's a neat idea <type type>, yup, I'll release it right away <ship off> no wait, that doesn't work at all <type type>, yup, I'll ship that one off right away <ship off> no, wait, that absolutely does not work" policy for releases. Micro$oft—bah. Amateurs. I'm *much* worse. (Or is that "worser"? "much worser"? "worsest"?)

I would like to take this opportunity to thank the Academy for... oops, wrong show.

- Masanobu UMEDA—the writer of the original GNUS.
- Shenghuo Zhu—uudecode.el, mm-uu.el, rfc1843.el, webmail.el, nnwarchive and many, many other things connected with MIME and other types of en/decoding, as well as general bug fixing, new functionality and stuff.
- Per Abrahamsen—custom, scoring, highlighting and SOUP code (as well as numerous other things).
- Luis Fernandes—design and graphics.

- Justin Sheehy–the FAQ maintainer.
- Erik Naggum—help, ideas, support, code and stuff.
- Wes Hardaker—'`gnus-picon.el`' and the manual section on *picons* (see Section 8.16.1 [Picons], page 204).
- Kim-Minh Kaplan—further work on the picon code.
- Brad Miller—'`gnus-gl.el`' and the GroupLens manual section (see Section 7.14 [GroupLens], page 183).
- Sudish Joseph—innumerable bug fixes.
- Ilja Weis—'`gnus-topic.el`'.
- Steven L. Baur—lots and lots and lots of bugs detections and fixes.
- Vladimir Alexiev—the refcard and reference booklets.
- Felix Lee & Jamie Zawinski—I stole some pieces from the XGnus distribution by Felix Lee and JWZ.
- Scott Byer—'`nnfolder.el`' enhancements & rewrite.
- Peter Mutsaers—orphan article scoring code.
- Ken Raeburn—POP mail support.
- Hallvard B Furuseth—various bits and pieces, especially dealing with .newsrc files.
- Brian Edmonds—'`gnus-bbdb.el`'.
- David Moore—rewrite of '`nnvirtual.el`' and many other things.
- Kevin Davidson—came up with the name *ding*, so blame him.
- Franois Pinard—many, many interesting and thorough bug reports, as well as autoconf support.

This manual was proof-read by Adrian Aichner, with Ricardo Nassif, Mark Borges, and Jost Krieger proof-reading parts of the manual.

The following people have contributed many patches and suggestions:

Christopher Davis, Andrew Eskilsson, Kai Grossjohann, David Kgedal, Richard Pieri, Fabrice Popineau, Daniel Quinlan, Jason L. Tibbitts, III, and Jack Vinson.

Also thanks to the following for patches and stuff:

Jari Aalto, Adrian Aichner, Vladimir Alexiev, Russ Allbery, Peter Arius, Matt Armstrong, Marc Auslander, Miles Bader, Alexei V. Barantsev, Frank Bennett, Robert Bihlmeyer, Chris Bone, Mark Borges, Mark Boyns, Lance A. Brown, Rob Browning, Kees de Bruin, Martin Buchholz, Joe Buehler, Kevin Buhr, Alastair Burt, Joao Cachopo, Zlatko Calusic, Massimo Campostrini, Castor, David Charlap, Dan Christensen, Kevin Christian, Jae-you Chung, James H. Cloos, Jr., Laura Conrad, Michael R. Cook, Glenn Coombs, Andrew J. Cosgriff, Neil Crellin, Frank D. Cringle, Geoffrey T. Dairiki, Andre Deparade, Ulrik Dickow, Dave Disser, Rui-Tao Dong, Joev Dubach, Michael Welsh Duggan, Dave Edmondson, Paul Eggert, Mark W. Eichin, Karl Eichwalder, Enami Tsugutomo, Michael Ernst, Luc Van Eycken, Sam Falkner, Nelson Jose dos Santos Ferreira, Sigbjorn Finne, Sven Fischer, Paul Fisher, Decklin Foster, Gary D. Foster, Paul Franklin, Guy Geens, Arne Georg Gleditsch, David S. Goldberg, Michelangelo Grigni, Dale Hagglund, D. Hall, Magnus Hammerin, Kenichi Handa, Raja R. Harinath, Yoshiki Hayashi, P. E. Jareth Hein, Hisashige Kenji, Scott Hofmann, Marc Horowitz, Gunnar Horrigmo, Richard

Hoskins, Brad Howes, Miguel de Icaza, Franois Felix Ingrand, Tatsuya Ichikawa, Ishikawa Ichiro, Lee Iverson, Iwamuro Motonori, Rajappa Iyer, Andreas Jaeger, Adam P. Jenkins, Randell Jesup, Fred Johansen, Gareth Jones, Simon Josefsson, Greg Klanderman, Karl Kleinpaste, Michael Klingbeil, Peter Skov Knudsen, Shuhei Kobayashi, Petr Konecny, Koseki Yoshinori, Thor Kristoffersen, Jens Lautenbacher, Martin Larose, Seokchan Lee, Joerg Lenneis, Carsten Leonhardt, James LewisMoss, Christian Limpach, Markus Linnala, Dave Love, Mike McEwan, Tonny Madsen, Shlomo Mahlab, Nat Makarevitch, Istvan Marko, David Martin, Jason R. Mastaler, Gordon Matzigkeit, Timo Metzemakers, Richard Mlynarik, Lantz Moore, Morioka Tomohiko, Erik Toubro Nielsen, Hrvoje Niksic, Andy Norman, Fred Oberhauser, C. R. Oldham, Alexandre Oliva, Ken Olstad, Masaharu Onishi, Hideki Ono, Ettore Perazzoli, William Perry, Stephen Peters, Jens-Ulrik Holger Petersen, Ulrich Pfeifer, Matt Pharr, Andy Piper, John McClary Prevost, Bill Pringlemeir, Mike Pullen, Jim Radford, Colin Rafferty, Lasse Rasinen, Lars Balker Rasmussen, Joe Reiss, Renaud Rioboo, Roland B. Roberts, Bart Robinson, Christian von Roques, Markus Rost, Jason Rumney, Wolfgang Rupprecht, Jay Sachs, Dewey M. Sasser, Conrad Sauerwald, Loren Schall, Dan Schmidt, Ralph Schleicher, Philippe Schnoebelen, Andreas Schwab, Randal L. Schwartz, Danny Siu, Matt Simmons, Paul D. Smith, Jeff Sparkes, Toby Speight, Michael Sperber, Darren Stalder, Richard Stallman, Greg Stark, Sam Steingold, Paul Stevenson, Jonas Steverud, Paul Stodghill, Kiyokazu Suto, Kurt Swanson, Samuel Tardieu, Teddy, Chuck Thompson, Tozawa Akihiko, Philippe Troin, James Troup, Trung Tran-Duc, Jack Twilley, Aaron M. Ucko, Aki Vehtari, Didier Verna, Vladimir Volovich, Jan Vroonhof, Stefan Waldherr, Pete Ware, Barry A. Warsaw, Christoph Wedler, Joe Wells, Lee Willis, Katsumi Yamaoka and Lloyd Zusman.

For a full overview of what each person has done, the ChangeLogs included in the Gnus alpha distributions should give ample reading (550kB and counting).

Apologies to everybody that I've forgotten, of which there are many, I'm sure.

Gee, that's quite a list of people. I guess that must mean that there actually are people who are using Gnus. Who'd'a thunk it!

10.1.9 New Features

These lists are, of course, just *short* overviews of the *most* important new features. No, really. There are tons more. Yes, we have feeping creaturism in full effect.

10.1.9.1 (ding) Gnus

New features in Gnus 5.0/5.1:

- The look of all buffers can be changed by setting format-like variables (see Section 2.1 [Group Buffer Format], page 13 and see Section 3.1 [Summary Buffer Format], page 37).

- Local spool and several NNTP servers can be used at once (see Chapter 6 [Select Methods], page 105).

- You can combine groups into virtual groups (see Section 6.6.1 [Virtual Groups], page 155).

- You can read a number of different mail formats (see Section 6.3 [Getting Mail], page 114). All the mail back ends implement a convenient mail expiry scheme (see Section 6.3.9 [Expiring Mail], page 128).

- Gnus can use various strategies for gathering threads that have lost their roots (thereby gathering loose sub-threads into one thread) or it can go back and retrieve enough headers to build a complete thread (see Section 3.8.1 [Customizing Threading], page 53).

- Killed groups can be displayed in the group buffer, and you can read them as well (see Section 2.11 [Listing Groups], page 25).

- Gnus can do partial group updates—you do not have to retrieve the entire active file just to check for new articles in a few groups (see Section 1.10 [The Active File], page 9).

- Gnus implements a sliding scale of subscribedness to groups (see Section 2.6 [Group Levels], page 19).

- You can score articles according to any number of criteria (see Chapter 7 [Scoring], page 167). You can even get Gnus to find out how to score articles for you (see Section 7.6 [Adaptive Scoring], page 176).

- Gnus maintains a dribble buffer that is auto-saved the normal Emacs manner, so it should be difficult to lose much data on what you have read if your machine should go down (see Section 1.9 [Auto Save], page 9).

- Gnus now has its own startup file ('`.gnus`') to avoid cluttering up the '`.emacs`' file.

- You can set the process mark on both groups and articles and perform operations on all the marked items (see Section 8.1 [Process/Prefix], page 189).

- You can grep through a subset of groups and create a group from the results (see Section 6.6.2 [Kibozed Groups], page 156).

- You can list subsets of groups according to, well, anything (see Section 2.11 [Listing Groups], page 25).

- You can browse foreign servers and subscribe to groups from those servers (see Section 2.14 [Browse Foreign Server], page 28).

- Gnus can fetch articles, asynchronously, on a second connection to the server (see Section 3.10 [Asynchronous Fetching], page 59).

- You can cache articles locally (see Section 3.11 [Article Caching], page 61).

- The uudecode functions have been expanded and generalized (see Section 3.15 [Decoding Articles], page 65).

- You can still post uuencoded articles, which was a little-known feature of GNUS' past (see Section 3.15.5.3 [Uuencoding and Posting], page 68).

- Fetching parents (and other articles) now actually works without glitches (see Section 3.21 [Finding the Parent], page 81).

- Gnus can fetch FAQs and group descriptions (see Section 2.17.2 [Group Information], page 34).

- Digests (and other files) can be used as the basis for groups (see Section 6.5.3 [Document Groups], page 143).

- Articles can be highlighted and customized (see Section 4.3 [Customizing Articles], page 95).

- URLs and other external references can be buttonized (see Section 3.16.5 [Article Buttons], page 76).

- You can do lots of strange stuff with the Gnus window & frame configuration (see Section 8.5 [Windows Configuration], page 193).

- You can click on buttons instead of using the keyboard (see Section 8.10 [Buttons], page 199).

10.1.9.2 September Gnus

New features in Gnus 5.2/5.3:

- A new message composition mode is used. All old customization variables for `mail-mode`, `rnews-reply-mode` and `gnus-msg` are now obsolete.
- Gnus is now able to generate *sparse* threads—threads where missing articles are represented by empty nodes (see Section 3.8.1 [Customizing Threading], page 53).

      (setq gnus-build-sparse-threads 'some)

- Outgoing articles are stored on a special archive server (see Section 5.4 [Archived Messages], page 100).
- Partial thread regeneration now happens when articles are referred.
- Gnus can make use of GroupLens predictions (see Section 7.14 [GroupLens], page 183).
- Picons (personal icons) can be displayed under XEmacs (see Section 8.16.1 [Picons], page 204).
- A `trn`-like tree buffer can be displayed (see Section 3.23 [Tree Display], page 84).

      (setq gnus-use-trees t)

- An `nn`-like pick-and-read minor mode is available for the summary buffers (see Section 3.22.1 [Pick and Read], page 83).

      (add-hook 'gnus-summary-mode-hook 'gnus-pick-mode)

- In binary groups you can use a special binary minor mode (see Section 3.22.2 [Binary Groups], page 83).
- Groups can be grouped in a folding topic hierarchy (see Section 2.16 [Group Topics], page 29).

      (add-hook 'gnus-group-mode-hook 'gnus-topic-mode)

- Gnus can re-send and bounce mail (see Section 3.5.1 [Summary Mail Commands], page 44).
- Groups can now have a score, and bubbling based on entry frequency is possible (see Section 2.7 [Group Score], page 20).

      (add-hook 'gnus-summary-exit-hook 'gnus-summary-bubble-group)

- Groups can be process-marked, and commands can be performed on groups of groups (see Section 2.8 [Marking Groups], page 20).
- Caching is possible in virtual groups.
- `nndoc` now understands all kinds of digests, mail boxes, rnews news batches, ClariNet briefs collections, and just about everything else (see Section 6.5.3 [Document Groups], page 143).
- Gnus has a new back end (`nnsoup`) to create/read SOUP packets (see Section 6.5.4 [SOUP], page 145).
- The Gnus cache is much faster.
- Groups can be sorted according to many criteria (see Section 2.12 [Sorting Groups], page 26).

- New group parameters have been introduced to set list-addresses and expiry times (see Section 2.10 [Group Parameters], page 22).

- All formatting specs allow specifying faces to be used (see Section 8.4.5 [Formatting Fonts], page 192).

- There are several more commands for setting/removing/acting on process marked articles on the `M P` submap (see Section 3.6.6 [Setting Process Marks], page 51).

- The summary buffer can be limited to show parts of the available articles based on a wide range of criteria. These commands have been bound to keys on the `/` submap (see Section 3.7 [Limiting], page 51).

- Articles can be made persistent with the `*` command (see Section 3.12 [Persistent Articles], page 61).

- All functions for hiding article elements are now toggles.

- Article headers can be buttonized (see Section 3.16.4 [Article Washing], page 73).

- All mail back ends support fetching articles by `Message-ID`.

- Duplicate mail can now be treated properly (see Section 6.3.11 [Duplicates], page 131).

- All summary mode commands are available directly from the article buffer (see Section 4.4 [Article Keymap], page 97).

- Frames can be part of `gnus-buffer-configuration` (see Section 8.5 [Windows Configuration], page 193).

- Mail can be re-scanned by a daemonic process (see Section 8.11 [Daemons], page 200).

- Gnus can make use of NoCeM files to weed out spam (see Section 8.12 [NoCeM], page 200).

    ```
    (setq gnus-use-nocem t)
    ```

- Groups can be made permanently visible (see Section 2.11 [Listing Groups], page 25).

    ```
    (setq gnus-permanently-visible-groups "^nnml:")
    ```

- Many new hooks have been introduced to make customizing easier.

- Gnus respects the `Mail-Copies-To` header.

- Threads can be gathered by looking at the `References` header (see Section 3.8.1 [Customizing Threading], page 53).

    ```
    (setq gnus-summary-thread-gathering-function
          'gnus-gather-threads-by-references)
    ```

- Read articles can be stored in a special backlog buffer to avoid refetching (see Section 3.13 [Article Backlog], page 62).

    ```
    (setq gnus-keep-backlog 50)
    ```

- A clean copy of the current article is always stored in a separate buffer to allow easier treatment.

- Gnus can suggest where to save articles (see Section 3.14 [Saving Articles], page 62).

- Gnus doesn't have to do as much prompting when saving (see Section 3.14 [Saving Articles], page 62).

    ```
    (setq gnus-prompt-before-saving t)
    ```

- `gnus-uu` can view decoded files asynchronously while fetching articles (see Section 3.15.5.2 [Other Decode Variables], page 67).

```
(setq gnus-uu-grabbed-file-functions 'gnus-uu-grab-view)
```
- Filling in the article buffer now works properly on cited text (see Section 3.16.4 [Article Washing], page 73).
- Hiding cited text adds buttons to toggle hiding, and how much cited text to hide is now customizable (see Section 3.16.3 [Article Hiding], page 71).

```
(setq gnus-cited-lines-visible 2)
```
- Boring headers can be hidden (see Section 3.16.3 [Article Hiding], page 71).
- Default scoring values can now be set from the menu bar.
- Further syntax checking of outgoing articles have been added.

10.1.9.3 Red Gnus

New features in Gnus 5.4/5.5:
- 'nntp.el' has been totally rewritten in an asynchronous fashion.
- Article prefetching functionality has been moved up into Gnus (see Section 3.10 [Asynchronous Fetching], page 59).
- Scoring can now be performed with logical operators like and, or, not, and parent redirection (see Section 7.15 [Advanced Scoring], page 185).
- Article washing status can be displayed in the article mode line (see Section 4.5 [Misc Article], page 97).
- 'gnus.el' has been split into many smaller files.
- Suppression of duplicate articles based on Message-ID can be done (see Section 3.28 [Duplicate Suppression], page 91).

```
(setq gnus-suppress-duplicates t)
```
- New variables for specifying what score and adapt files are to be considered home score and adapt files (see Section 7.7 [Home Score File], page 178) have been added.
- nndoc was rewritten to be easily extendable (see Section 6.5.3.1 [Document Server Internals], page 144).
- Groups can inherit group parameters from parent topics (see Section 2.16.5 [Topic Parameters], page 32).
- Article editing has been revamped and is now actually usable.
- Signatures can be recognized in more intelligent fashions (see Section 3.16.7 [Article Signature], page 77).
- Summary pick mode has been made to look more nn-like. Line numbers are displayed and the . command can be used to pick articles (Pick and Read).
- Commands for moving the '.newsrc.eld' from one server to another have been added (see Section 1.7 [Changing Servers], page 7).
- There's a way now to specify that "uninteresting" fields be suppressed when generating lines in buffers (see Section 8.4.3 [Advanced Formatting], page 191).
- Several commands in the group buffer can be undone with M-C-_ (see Section 8.13 [Undo], page 202).
- Scoring can be done on words using the new score type w (see Section 7.4 [Score File Format], page 172).

- Adaptive scoring can be done on a Subject word-by-word basis (see Section 7.6 [Adaptive Scoring], page 176).

        (setq gnus-use-adaptive-scoring '(word))
- Scores can be decayed (see Section 7.16 [Score Decays], page 188).

        (setq gnus-decay-scores t)
- Scoring can be performed using a regexp on the Date header. The Date is normalized to compact ISO 8601 format first (see Section 7.4 [Score File Format], page 172).
- A new command has been added to remove all data on articles from the native server (see Section 1.7 [Changing Servers], page 7).
- A new command for reading collections of documents (`nndoc` with `nnvirtual` on top) has been added—`M-C-d` (see Section 3.25.4 [Really Various Summary Commands], page 89).
- Process mark sets can be pushed and popped (see Section 3.6.6 [Setting Process Marks], page 51).
- A new mail-to-news back end makes it possible to post even when the NNTP server doesn't allow posting (see Section 6.5.5 [Mail-To-News Gateways], page 148).
- A new back end for reading searches from Web search engines (*DejaNews*, *Alta Vista*, *InReference*) has been added (see Section 6.4.1 [Web Searches], page 138).
- Groups inside topics can now be sorted using the standard sorting functions, and each topic can be sorted independently (see Section 2.16.3 [Topic Sorting], page 31).
- Subsets of the groups can be sorted independently (`Sorting Groups`).
- Cached articles can be pulled into the groups (see Section 3.25.3 [Summary Generation Commands], page 88).
- Score files are now applied in a more reliable order (see Section 7.3 [Score Variables], page 169).
- Reports on where mail messages end up can be generated (see Section 6.3.3 [Splitting Mail], page 115).
- More hooks and functions have been added to remove junk from incoming mail before saving the mail (see Section 6.3.10 [Washing Mail], page 130).
- Emphasized text can be properly fontisized:

10.1.9.4 Quassia Gnus

New features in Gnus 5.6:

- New functionality for using Gnus as an offline newsreader has been added. A plethora of new commands and modes have been added. See see Section 6.7 [Gnus Unplugged], page 156 for the full story.
- The `nndraft` back end has returned, but works differently than before. All Message buffers are now also articles in the `nndraft` group, which is created automatically.
- `gnus-alter-header-function` can now be used to alter header values.
- `gnus-summary-goto-article` now accept Message-ID's.
- A new Message command for deleting text in the body of a message outside the region: `C-c C-v`.

- You can now post to component group in `nnvirtual` groups with `C-u C-c C-c`.
- `nntp-rlogin-program`—new variable to ease customization.
- `C-u C-c C-c` in `gnus-article-edit-mode` will now inhibit re-highlighting of the article buffer.
- New element in `gnus-boring-article-headers`—`long-to`.
- `M-i` symbolic prefix command. See the section "Symbolic Prefixes" in the Gnus manual for details.
- `L` and `I` in the summary buffer now take the symbolic prefix `a` to add the score rule to the "all.SCORE" file.
- `gnus-simplify-subject-functions` variable to allow greater control over simplification.
- `A T`—new command for fetching the current thread.
- `/ T`—new command for including the current thread in the limit.
- `M-`⟨RET⟩ is a new Message command for breaking cited text.
- '\\1'-expressions are now valid in `nnmail-split-methods`.
- The `custom-face-lookup` function has been removed. If you used this function in your initialization files, you must rewrite them to use `face-spec-set` instead.
- Canceling now uses the current select method. Symbolic prefix `a` forces normal posting method.
- New command to translate M******** sm*rtq**t*s into proper text—`W d`.
- For easier debugging of `nntp`, you can set `nntp-record-commands` to a non-`nil` value.
- `nntp` now uses '`~/.authinfo`', a '`.netrc`'-like file, for controlling where and how to send AUTHINFO to NNTP servers.
- A command for editing group parameters from the summary buffer has been added.
- A history of where mails have been split is available.
- A new article date command has been added—`article-date-iso8601`.
- Subjects can be simplified when threading by setting `gnus-score-thread-simplify`.
- A new function for citing in Message has been added—`message-cite-original-without-signature`.
- `article-strip-all-blank-lines`—new article command.
- A new Message command to kill to the end of the article has been added.
- A minimum adaptive score can be specified by using the `gnus-adaptive-word-minimum` variable.
- The "lapsed date" article header can be kept continually updated by the `gnus-start-date-timer` command.
- Web listserv archives can be read with the `nnlistserv` back end.
- Old dejanews archives can now be read by `nnweb`.

10.1.9.5 Pterodactyl Gnus

New features in Gnus 5.8:

- The mail-fetching functions have changed. See the manual for the many details. In particular, all procmail fetching variables are gone.

  If you used procmail like in

  ```
  (setq nnmail-use-procmail t)
  (setq nnmail-spool-file 'procmail)
  (setq nnmail-procmail-directory "~/mail/incoming/")
  (setq nnmail-procmail-suffix "\\.in")
  ```

  this now has changed to

  ```
  (setq mail-sources
        '((directory :path "~/mail/incoming/"
      :suffix ".in")))
  ```

  More information is available in the info doc at Select Methods -> Getting Mail -> Mail Sources

- Gnus is now a MIME-capable reader. This affects many parts of Gnus, and adds a slew of new commands. See the manual for details.

- Gnus has also been multilingualized. This also affects too many parts of Gnus to summarize here, and adds many new variables.

- `gnus-auto-select-first` can now be a function to be called to position point.

- The user can now decide which extra headers should be included in summary buffers and NOV files.

- `gnus-article-display-hook` has been removed. Instead, a number of variables starting with `gnus-treat-` have been added.

- The Gnus posting styles have been redone again and now works in a subtly different manner.

- New web-based back ends have been added: `nnslashdot`, `nnwarchive` and `nnultimate`. nnweb has been revamped, again, to keep up with ever-changing layouts.

- Gnus can now read IMAP mail via `nnimap`.

10.2 The Manual

This manual was generated from a TeXinfo file and then run through either `texi2dvi` to get what you hold in your hands now.

The following conventions have been used:

1. This is a 'string'

2. This is a *keystroke*

3. This is a 'file'

4. This is a symbol

So if I were to say "set `flargnoze` to 'yes'", that would mean:

```
(setq flargnoze "yes")
```

If I say "set `flumphel` to yes", that would mean:

```
(setq flumphel 'yes)
```

'yes' and yes are two *very* different things—don't ever get them confused.

10.3 On Writing Manuals

I guess most manuals are written after-the-fact; documenting a program that's already there. This is not how this manual is written. When implementing something, I write the manual entry for that something straight away. I then see that it's difficult to explain the functionality, so I write how it's supposed to be, and then I change the implementation. Writing the documentation and writing the code goes hand in hand.

This, of course, means that this manual has no, or little, flow. It documents absolutely everything in Gnus, but often not where you're looking for it. It is a reference manual, and not a guide to how to get started with Gnus.

That would be a totally different book, that should be written using the reference manual as source material. It would look quite differently.

10.4 Terminology

news This is what you are supposed to use this thing for—reading news. News is generally fetched from a nearby NNTP server, and is generally publicly available to everybody. If you post news, the entire world is likely to read just what you have written, and they'll all snigger mischievously. Behind your back.

mail Everything that's delivered to you personally is mail. Some news/mail readers (like Gnus) blur the distinction between mail and news, but there is a difference. Mail is private. News is public. Mailing is not posting, and replying is not following up.

reply Send a mail to the person who has written what you are reading.

follow up Post an article to the current newsgroup responding to the article you are reading.

back end Gnus gets fed articles from a number of back ends, both news and mail back ends. Gnus does not handle the underlying media, so to speak—this is all done by the back ends.

native Gnus will always use one method (and back end) as the *native*, or default, way of getting news.

foreign You can also have any number of foreign groups active at the same time. These are groups that use non-native non-secondary back ends for getting news.

secondary Secondary back ends are somewhere half-way between being native and being foreign, but they mostly act like they are native.

article A message that has been posted as news.

mail message
         A message that has been mailed.

message A mail message or news article

head The top part of a message, where administrative information (etc.) is put.

body The rest of an article. Everything not in the head is in the body.

header A line from the head of an article.

headers A collection of such lines, or a collection of heads. Or even a collection of NOV lines.

NOV When Gnus enters a group, it asks the back end for the headers of all unread articles in the group. Most servers support the News OverView format, which is more compact and much faster to read and parse than the normal HEAD format.

level Each group is subscribed at some *level* or other (1-9). The ones that have a lower level are "more" subscribed than the groups with a higher level. In fact, groups on levels 1-5 are considered *subscribed*; 6-7 are *unsubscribed*; 8 are *zombies*; and 9 are *killed*. Commands for listing groups and scanning for new articles will all use the numeric prefix as *working level*.

killed groups
> No information on killed groups is stored or updated, which makes killed groups much easier to handle than subscribed groups.

zombie groups
> Just like killed groups, only slightly less dead.

active file The news server has to keep track of what articles it carries, and what groups exist. All this information in stored in the active file, which is rather large, as you might surmise.

bogus groups
> A group that exists in the '.newsrc' file, but isn't known to the server (i.e., it isn't in the active file), is a *bogus group*. This means that the group probably doesn't exist (any more).

activating The act of asking the server for info on a group and computing the number of unread articles is called *activating the group*. Un-activated groups are listed with '*' in the group buffer.

server A machine one can connect to and get news (or mail) from.

select method
> A structure that specifies the back end, the server and the virtual server settings.

virtual server
> A named select method. Since a select method defines all there is to know about connecting to a (physical) server, taking the thing as a whole is a virtual server.

washing Taking a buffer and running it through a filter of some sort. The result will (more often than not) be cleaner and more pleasing than the original.

ephemeral groups
> Most groups store data on what articles you have read. *Ephemeral* groups are groups that will have no data stored—when you exit the group, it'll disappear into the aether.

solid groups
> This is the opposite of ephemeral groups. All groups listed in the group buffer are solid groups.

sparse articles
> These are article placeholders shown in the summary buffer when `gnus-build-sparse-threads` has been switched on.

threading To put responses to articles directly after the articles they respond to—in a hierarchical fashion.

root The first article in a thread is the root. It is the ancestor of all articles in the thread.

parent An article that has responses.

child An article that responds to a different article—its parent.

digest   A collection of messages in one file. The most common digest format is specified by RFC 1153.

10.5 Customization

All variables are properly documented elsewhere in this manual. This section is designed to give general pointers on how to customize Gnus for some quite common situations.

10.5.1 Slow/Expensive NNTP Connection

If you run Emacs on a machine locally, and get your news from a machine over some very thin strings, you want to cut down on the amount of data Gnus has to get from the NNTP server.

`gnus-read-active-file`

> Set this to `nil`, which will inhibit Gnus from requesting the entire active file from the server. This file is often very large. You also have to set `gnus-check-new-newsgroups` and `gnus-check-bogus-newsgroups` to `nil` to make sure that Gnus doesn't suddenly decide to fetch the active file anyway.

`gnus-nov-is-evil`

> This one has to be `nil`. If not, grabbing article headers from the NNTP server will not be very fast. Not all NNTP servers support XOVER; Gnus will detect this by itself.

10.5.2 Slow Terminal Connection

Let's say you use your home computer for dialing up the system that runs Emacs and Gnus. If your modem is slow, you want to reduce (as much as possible) the amount of data sent over the wires.

`gnus-auto-center-summary`

> Set this to `nil` to inhibit Gnus from re-centering the summary buffer all the time. If it is `vertical`, do only vertical re-centering. If it is neither `nil` nor `vertical`, do both horizontal and vertical recentering.

`gnus-visible-headers`

> Cut down on the headers included in the articles to the minimum. You can, in fact, make do without them altogether—most of the useful data is in the summary buffer, anyway. Set this variable to '`^NEVVVVER`' or '`From:`', or whatever you feel you need.
>
> Set this hook to all the available hiding commands:
>
> ```
> (setq gnus-treat-hide-headers 'head
>       gnus-treat-hide-signature t
>       gnus-treat-hide-citation t)
> ```

`gnus-use-full-window`

> By setting this to `nil`, you can make all the windows smaller. While this doesn't really cut down much generally, it means that you have to see smaller portions of articles before deciding that you didn't want to read them anyway.

`gnus-thread-hide-subtree`

> If this is non-`nil`, all threads in the summary buffer will be hidden initially.

`gnus-updated-mode-lines`

> If this is `nil`, Gnus will not put information in the buffer mode lines, which might save some time.

10.5.3 Little Disk Space

The startup files can get rather large, so you may want to cut their sizes a bit if you are running out of space.

`gnus-save-newsrc-file`

> If this is `nil`, Gnus will never save '`.newsrc`'—it will only save '`.newsrc.eld`'. This means that you will not be able to use any other newsreaders than Gnus. This variable is `t` by default.

`gnus-read-newsrc-file`

> If this is `nil`, Gnus will never read '`.newsrc`'—it will only read '`.newsrc.eld`'. This means that you will not be able to use any other newsreaders than Gnus. This variable is `t` by default.

`gnus-save-killed-list`

> If this is `nil`, Gnus will not save the list of dead groups. You should also set `gnus-check-new-newsgroups` to `ask-server` and `gnus-check-bogus-newsgroups` to `nil` if you set this variable to `nil`. This variable is `t` by default.

10.5.4 Slow Machine

If you have a slow machine, or are just really impatient, there are a few things you can do to make Gnus run faster.

Set `gnus-check-new-newsgroups` and `gnus-check-bogus-newsgroups` to `nil` to make startup faster.

Set `gnus-show-threads`, `gnus-use-cross-reference` and `gnus-nov-is-evil` to `nil` to make entering and exiting the summary buffer faster.

10.6 Troubleshooting

Gnus works *so* well straight out of the box—I can't imagine any problems, really.

Ahem.

1. Make sure your computer is switched on.

2. Make sure that you really load the current Gnus version. If you have been running GNUS, you need to exit Emacs and start it up again before Gnus will work.

3. Try doing an `M-x gnus-version`. If you get something that looks like '`Gnus v5.46; nntp 4.0`' you have the right files loaded. If, on the other hand, you get something like '`NNTP 3.x`' or '`nntp flee`', you have some old '`.el`' files lying around. Delete these.

4. Read the help group (`G h` in the group buffer) for a FAQ and a how-to.

5. Gnus works on many recursive structures, and in some extreme (and very rare) cases Gnus may recurse down "too deeply" and Emacs will beep at you. If this happens to you, set `max-lisp-eval-depth` to 500 or something like that.

If all else fails, report the problem as a bug.

If you find a bug in Gnus, you can report it with the `M-x gnus-bug` command. `M-x set-variable` ⟨RET⟩ `debug-on-error` ⟨RET⟩ `t` ⟨RET⟩, and send me the backtrace. I will fix bugs, but I can only fix them if you send me a precise description as to how to reproduce the bug.

You really can never be too detailed in a bug report. Always use the `M-x gnus-bug` command when you make bug reports, even if it creates a 10Kb mail each time you use it, and even if you have sent me your environment 500 times before. I don't care. I want the full info each time.

It is also important to remember that I have no memory whatsoever. If you send a bug report, and I send you a reply, and then you just send back "No, it's not! Moron!", I will have no idea what you are insulting me about. Always over-explain everything. It's much easier for all of us—if I don't have all the information I need, I will just mail you and ask for more info, and everything takes more time.

If the problem you're seeing is very visual, and you can't quite explain it, copy the Emacs window to a file (with `xwd`, for instance), put it somewhere it can be reached, and include the URL of the picture in the bug report.

If you just need help, you are better off asking on '`gnu.emacs.gnus`'. I'm not very helpful.

You can also ask on the ding mailing list—'`ding@gnus.org`'. Write to '`ding-request@gnus.org`' to subscribe.

10.7 Gnus Reference Guide

It is my hope that other people will figure out smart stuff that Gnus can do, and that other people will write those smart things as well. To facilitate that I thought it would be a good idea to describe the inner workings of Gnus. And some of the not-so-inner workings, while I'm at it.

You can never expect the internals of a program not to change, but I will be defining (in some details) the interface between Gnus and its back ends (this is written in stone), the format of the score files (ditto), data structures (some are less likely to change than others) and general methods of operation.

10.7.1 Gnus Utility Functions

When writing small functions to be run from hooks (and stuff), it's vital to have access to the Gnus internal functions and variables. Below is a list of the most common ones.

`gnus-newsgroup-name`
> This variable holds the name of the current newsgroup.

`gnus-find-method-for-group`
> A function that returns the select method for *group*.

`gnus-group-real-name`
> Takes a full (prefixed) Gnus group name, and returns the unprefixed name.

`gnus-group-prefixed-name`
> Takes an unprefixed group name and a select method, and returns the full (prefixed) Gnus group name.

`gnus-get-info`
> Returns the group info list for *group*.

`gnus-group-unread`
> The number of unread articles in *group*, or `t` if that is unknown.

`gnus-active`
> The active entry for *group*.

`gnus-set-active`
> Set the active entry for *group*.

`gnus-add-current-to-buffer-list`
> Adds the current buffer to the list of buffers to be killed on Gnus exit.

`gnus-continuum-version`
> Takes a Gnus version string as a parameter and returns a floating point number. Earlier versions will always get a lower number than later versions.

`gnus-group-read-only-p`
> Says whether *group* is read-only or not.

`gnus-news-group-p`
> Says whether *group* came from a news back end.

`gnus-ephemeral-group-p`
> Says whether *group* is ephemeral or not.

`gnus-server-to-method`
> Returns the select method corresponding to *server*.

`gnus-server-equal`
> Says whether two virtual servers are equal.

`gnus-group-native-p`
> Says whether *group* is native or not.

`gnus-group-secondary-p`
> Says whether *group* is secondary or not.

`gnus-group-foreign-p`
> Says whether *group* is foreign or not.

`group-group-find-parameter`
> Returns the parameter list of *group*. If given a second parameter, returns the value of that parameter for *group*.

`gnus-group-set-parameter`
> Takes three parameters; *group*, *parameter* and *value*.

`gnus-narrow-to-body`
> Narrows the current buffer to the body of the article.

`gnus-check-backend-function`
> Takes two parameters, *function* and *group*. If the back end *group* comes from supports *function*, return non-`nil`.
>
> ```
> (gnus-check-backend-function "request-scan" "nnml:misc")
> ⇒ t
> ```

`gnus-read-method`
> Prompts the user for a select method.

10.7.2 Back End Interface

Gnus doesn't know anything about NNTP, spools, mail or virtual groups. It only knows how to talk to *virtual servers*. A virtual server is a *back end* and some *back end variables*. As examples of the first, we have `nntp`, `nnspool` and `nnmbox`. As examples of the latter we have `nntp-port-number` and `nnmbox-directory`.

When Gnus asks for information from a back end—say `nntp`—on something, it will normally include a virtual server name in the function parameters. (If not, the back end should use the "current" virtual server.) For instance, `nntp-request-list` takes a virtual server as its only (optional) parameter. If this virtual server hasn't been opened, the function should fail.

Note that a virtual server name has no relation to some physical server name. Take this example:

```
(nntp "odd-one"
      (nntp-address "ifi.uio.no")
```

```
        (nntp-port-number 4324))
```

Here the virtual server name is '`odd-one`' while the name of the physical server is '`ifi.uio.no`'.

The back ends should be able to switch between several virtual servers. The standard back ends implement this by keeping an alist of virtual server environments that they pull down/push up when needed.

There are two groups of interface functions: *required functions*, which must be present, and *optional functions*, which Gnus will always check for presence before attempting to call 'em.

All these functions are expected to return data in the buffer `nntp-server-buffer` ('`*nntpd*`'), which is somewhat unfortunately named, but we'll have to live with it. When I talk about *resulting data*, I always refer to the data in that buffer. When I talk about *return value*, I talk about the function value returned by the function call. Functions that fail should return `nil` as the return value.

Some back ends could be said to be *server-forming* back ends, and some might be said not to be. The latter are back ends that generally only operate on one group at a time, and have no concept of "server" – they have a group, and they deliver info on that group and nothing more.

In the examples and definitions I will refer to the imaginary back end `nnchoke`.

10.7.2.1 Required Back End Functions

(`nnchoke-retrieve-headers ARTICLES &optional GROUP SERVER FETCH-OLD`)

> *articles* is either a range of article numbers or a list of `Message-ID`s. Current back ends do not fully support either—only sequences (lists) of article numbers, and most back ends do not support retrieval of `Message-ID`s. But they should try for both.

> The result data should either be HEADs or NOV lines, and the result value should either be `headers` or `nov` to reflect this. This might later be expanded to `various`, which will be a mixture of HEADs and NOV lines, but this is currently not supported by Gnus.

> If *fetch-old* is non-`nil` it says to try fetching "extra headers", in some meaning of the word. This is generally done by fetching (at most) *fetch-old* extra headers less than the smallest article number in `articles`, and filling the gaps as well. The presence of this parameter can be ignored if the back end finds it cumbersome to follow the request. If this is non-`nil` and not a number, do maximum fetches.

> Here's an example HEAD:

```
        221 1056 Article retrieved.
        Path: ifi.uio.no!sturles
        From: sturles@ifi.uio.no (Sturle Sunde)
        Newsgroups: ifi.discussion
        Subject: Re: Something very droll
        Date: 27 Oct 1994 14:02:57 +0100
        Organization: Dept. of Informatics, University of Oslo, Norway
```

```
Lines: 26
Message-ID: <38o8e1$a0o@holmenkollen.ifi.uio.no>
References: <38jdmq$4qu@visbur.ifi.uio.no>
NNTP-Posting-Host: holmenkollen.ifi.uio.no
.
```

So a `headers` return value would imply that there's a number of these in the data buffer.

Here's a BNF definition of such a buffer:

```
headers        = *head
head           = error / valid-head
error-message  = [ "4" / "5" ] 2number " " <error message> eol
valid-head     = valid-message *header "." eol
valid-message  = "221 " <number> " Article retrieved." eol
header         = <text> eol
```

If the return value is `nov`, the data buffer should contain *network overview database* lines. These are basically fields separated by tabs.

```
nov-buffer = *nov-line
nov-line   = 8*9 [ field <TAB> ] eol
field      = <text except TAB>
```

For a closer look at what should be in those fields, see Section 10.7.4 [Headers], page 248.

(nnchoke-open-server SERVER &optional DEFINITIONS)
server is here the virtual server name. *definitions* is a list of (`VARIABLE VALUE`) pairs that define this virtual server.

If the server can't be opened, no error should be signaled. The back end may then choose to refuse further attempts at connecting to this server. In fact, it should do so.

If the server is opened already, this function should return a non-`nil` value. There should be no data returned.

(nnchoke-close-server &optional SERVER)
Close connection to *server* and free all resources connected to it. Return `nil` if the server couldn't be closed for some reason.

There should be no data returned.

(nnchoke-request-close)
Close connection to all servers and free all resources that the back end have reserved. All buffers that have been created by that back end should be killed. (Not the `nntp-server-buffer`, though.) This function is generally only called when Gnus is shutting down.

There should be no data returned.

(nnchoke-server-opened &optional SERVER)
If *server* is the current virtual server, and the connection to the physical server is alive, then this function should return a non-`nil` vlue. This function should under no circumstances attempt to reconnect to a server we have lost connection to.

There should be no data returned.

`(nnchoke-status-message &optional SERVER)`

This function should return the last error message from *server*.

There should be no data returned.

`(nnchoke-request-article ARTICLE &optional GROUP SERVER TO-BUFFER)`

The result data from this function should be the article specified by *article*. This might either be a `Message-ID` or a number. It is optional whether to implement retrieval by `Message-ID`, but it would be nice if that were possible.

If *to-buffer* is non-`nil`, the result data should be returned in this buffer instead of the normal data buffer. This is to make it possible to avoid copying large amounts of data from one buffer to another, while Gnus mainly requests articles to be inserted directly into its article buffer.

If it is at all possible, this function should return a cons cell where the `car` is the group name the article was fetched from, and the `cdr` is the article number. This will enable Gnus to find out what the real group and article numbers are when fetching articles by `Message-ID`. If this isn't possible, `t` should be returned on successful article retrieval.

`(nnchoke-request-group GROUP &optional SERVER FAST)`

Get data on *group*. This function also has the side effect of making *group* the current group.

If *fast*, don't bother to return useful data, just make *group* the current group.

Here's an example of some result data and a definition of the same:

```
211 56 1000 1059 ifi.discussion
```

The first number is the status, which should be 211. Next is the total number of articles in the group, the lowest article number, the highest article number, and finally the group name. Note that the total number of articles may be less than one might think while just considering the highest and lowest article numbers, but some articles may have been canceled. Gnus just discards the total-number, so whether one should take the bother to generate it properly (if that is a problem) is left as an exercise to the reader.

```
group-status = [ error / info ] eol
error        = [ "4" / "5" ] 2<number> " " <Error message>
info         = "211 " 3* [ <number> " " ] <string>
```

`(nnchoke-close-group GROUP &optional SERVER)`

Close *group* and free any resources connected to it. This will be a no-op on most back ends.

There should be no data returned.

`(nnchoke-request-list &optional SERVER)`

Return a list of all groups available on *server*. And that means *all*.

Here's an example from a server that only carries two groups:

```
ifi.test 0000002200 0000002000 y
ifi.discussion 3324 3300 n
```

On each line we have a group name, then the highest article number in that group, the lowest article number, and finally a flag.

```
active-file = *active-line
active-line = name " " <number> " " <number> " " flags eol
name        = <string>
flags       = "n" / "y" / "m" / "x" / "j" / "=" name
```

The flag says whether the group is read-only ('n'), is moderated ('m'), is dead ('x'), is aliased to some other group ('=other-group') or none of the above ('y').

`(nnchoke-request-post &optional SERVER)`

> This function should post the current buffer. It might return whether the posting was successful or not, but that's not required. If, for instance, the posting is done asynchronously, it has generally not been completed by the time this function concludes. In that case, this function should set up some kind of sentinel to beep the user loud and clear if the posting could not be completed.
>
> There should be no result data from this function.

10.7.2.2 Optional Back End Functions

`(nnchoke-retrieve-groups GROUPS &optional SERVER)`

> *groups* is a list of groups, and this function should request data on all those groups. How it does it is of no concern to Gnus, but it should attempt to do this in a speedy fashion.
>
> The return value of this function can be either `active` or `group`, which says what the format of the result data is. The former is in the same format as the data from `nnchoke-request-list`, while the latter is a buffer full of lines in the same format as `nnchoke-request-group` gives.
>
> ```
> group-buffer = *active-line / *group-status
> ```

`(nnchoke-request-update-info GROUP INFO &optional SERVER)`

> A Gnus group info (see Section 10.7.6 [Group Info], page 250) is handed to the back end for alterations. This comes in handy if the back end really carries all the information (as is the case with virtual and imap groups). This function should destructively alter the info to suit its needs, and should return the (altered) group info.
>
> There should be no result data from this function.

`(nnchoke-request-type GROUP &optional ARTICLE)`

> When the user issues commands for "sending news" (`F` in the summary buffer, for instance), Gnus has to know whether the article the user is following up on is news or mail. This function should return `news` if *article* in *group* is news, `mail` if it is mail and `unknown` if the type can't be decided. (The *article* parameter is necessary in `nnvirtual` groups which might very well combine mail groups and news groups.) Both *group* and *article* may be `nil`.
>
> There should be no result data from this function.

`(nnchoke-request-set-mark GROUP ACTION &optional SERVER)`

>Set/remove/add marks on articles. Normally Gnus handles the article marks (such as read, ticked, expired etc) internally, and store them in `~/.newsrc.eld`. Some back ends (such as IMAP) however carry all information about the articles on the server, so Gnus need to propagate the mark information to the server.

>ACTION is a list of mark setting requests, having this format:

>>`(RANGE ACTION MARK)`

>Range is a range of articles you wish to update marks on. Action is `set`, `add` or `del`, respectively used for removing all existing marks and setting them as specified, adding (preserving the marks not mentioned) mark and removing (preserving the marks not mentioned) marks. Mark is a list of marks; where each mark is a symbol. Currently used marks are `read`, `tick`, `reply`, `expire`, `killed`, `dormant`, `save`, `download` and `unsend`, but your back end should, if possible, not limit itself to these.

>Given contradictory actions, the last action in the list should be the effective one. That is, if your action contains a request to add the `tick` mark on article 1 and, later in the list, a request to remove the mark on the same article, the mark should in fact be removed.

>An example action list:

>>```
(((5 12 30) 'del '(tick))
 ((10 . 90) 'add '(read expire))
 ((92 94) 'del '(read)))
```

>The function should return a range of articles it wasn't able to set the mark on (currently not used for anything).

>There should be no result data from this function.

`(nnchoke-request-update-mark GROUP ARTICLE MARK)`

>If the user tries to set a mark that the back end doesn't like, this function may change the mark. Gnus will use whatever this function returns as the mark for *article* instead of the original *mark*. If the back end doesn't care, it must return the original *mark*, and not `nil` or any other type of garbage.

>The only use for this I can see is what `nnvirtual` does with it—if a component group is auto-expirable, marking an article as read in the virtual group should result in the article being marked as expirable.

>There should be no result data from this function.

`(nnchoke-request-scan &optional GROUP SERVER)`

>This function may be called at any time (by Gnus or anything else) to request that the back end check for incoming articles, in one way or another. A mail back end will typically read the spool file or query the POP server when this function is invoked. The *group* doesn't have to be heeded—if the back end decides that it is too much work just scanning for a single group, it may do a total scan of all groups. It would be nice, however, to keep things local if that's practical.

>There should be no result data from this function.

`(nnchoke-request-group-description GROUP &optional SERVER)`

> The result data from this function should be a description of *group*.
>
> ```
> description-line = name <TAB> description eol
> name             = <string>
> description      = <text>
> ```

`(nnchoke-request-list-newsgroups &optional SERVER)`

> The result data from this function should be the description of all groups available on the server.
>
> ```
> description-buffer = *description-line
> ```

`(nnchoke-request-newgroups DATE &optional SERVER)`

> The result data from this function should be all groups that were created after 'date', which is in normal human-readable date format. The data should be in the active buffer format.

`(nnchoke-request-create-group GROUP &optional SERVER)`

> This function should create an empty group with name *group*.
>
> There should be no return data.

`(nnchoke-request-expire-articles ARTICLES &optional GROUP SERVER FORCE)`

> This function should run the expiry process on all articles in the *articles* range (which is currently a simple list of article numbers.) It is left up to the back end to decide how old articles should be before they are removed by this function. If *force* is non-`nil`, all *articles* should be deleted, no matter how new they are.
>
> This function should return a list of articles that it did not/was not able to delete.
>
> There should be no result data returned.

`(nnchoke-request-move-article ARTICLE GROUP SERVER ACCEPT-FORM`
`        &optional LAST)`

> This function should move *article* (which is a number) from *group* by calling *accept-form*.
>
> This function should ready the article in question for moving by removing any header lines it has added to the article, and generally should "tidy up" the article. Then it should `eval` *accept-form* in the buffer where the "tidy" article is. This will do the actual copying. If this `eval` returns a non-`nil` value, the article should be removed.
>
> If *last* is `nil`, that means that there is a high likelihood that there will be more requests issued shortly, so that allows some optimizations.
>
> The function should return a cons where the `car` is the group name and the `cdr` is the article number that the article was entered as.
>
> There should be no data returned.

`(nnchoke-request-accept-article GROUP &optional SERVER LAST)`

> This function takes the current buffer and inserts it into *group*. If *last* in `nil`, that means that there will be more calls to this function in short order.
>
> The function should return a cons where the `car` is the group name and the `cdr` is the article number that the article was entered as.

There should be no data returned.

`(nnchoke-request-replace-article ARTICLE GROUP BUFFER)`

> This function should remove *article* (which is a number) from *group* and insert *buffer* there instead.
>
> There should be no data returned.

`(nnchoke-request-delete-group GROUP FORCE &optional SERVER)`

> This function should delete *group*. If *force*, it should really delete all the articles in the group, and then delete the group itself. (If there is such a thing as "the group itself".)
>
> There should be no data returned.

`(nnchoke-request-rename-group GROUP NEW-NAME &optional SERVER)`

> This function should rename *group* into *new-name*. All articles in *group* should move to *new-name*.
>
> There should be no data returned.

10.7.2.3 Error Messaging

The back ends should use the function `nnheader-report` to report error conditions—they should not raise errors when they aren't able to perform a request. The first argument to this function is the back end symbol, and the rest are interpreted as arguments to `format` if there are multiple of them, or just a string if there is one of them. This function must always returns `nil`.

```
(nnheader-report 'nnchoke "You did something totally bogus")

(nnheader-report 'nnchoke "Could not request group %s" group)
```

Gnus, in turn, will call `nnheader-get-report` when it gets a `nil` back from a server, and this function returns the most recently reported message for the back end in question. This function takes one argument—the server symbol.

Internally, these functions access *back-end*-`status-string`, so the `nnchoke` back end will have its error message stored in `nnchoke-status-string`.

10.7.2.4 Writing New Back Ends

Many back ends are quite similar. `nnml` is just like `nnspool`, but it allows you to edit the articles on the server. `nnmh` is just like `nnml`, but it doesn't use an active file, and it doesn't maintain overview databases. `nndir` is just like `nnml`, but it has no concept of "groups", and it doesn't allow editing articles.

It would make sense if it were possible to "inherit" functions from back ends when writing new back ends. And, indeed, you can do that if you want to. (You don't have to if you don't want to, of course.)

All the back ends declare their public variables and functions by using a package called `nnoo`.

To inherit functions from other back ends (and allow other back ends to inherit functions from the current back end), you should use the following macros:

nnoo-declare
:   This macro declares the first parameter to be a child of the subsequent parameters. For instance:

    ```
    (nnoo-declare nndir
      nnml nnmh)
    ```

    `nndir` has declared here that it intends to inherit functions from both `nnml` and `nnmh`.

defvoo
:   This macro is equivalent to `defvar`, but registers the variable as a public server variable. Most state-oriented variables should be declared with `defvoo` instead of `defvar`.

    In addition to the normal `defvar` parameters, it takes a list of variables in the parent back ends to map the variable to when executing a function in those back ends.

    ```
    (defvoo nndir-directory nil
      "Where nndir will look for groups."
      nnml-current-directory nnmh-current-directory)
    ```

    This means that `nnml-current-directory` will be set to `nndir-directory` when an `nnml` function is called on behalf of `nndir`. (The same with `nnmh`.)

nnoo-define-basics
:   This macro defines some common functions that almost all back ends should have.

    ```
    (nnoo-define-basics nndir)
    ```

deffoo
:   This macro is just like `defun` and takes the same parameters. In addition to doing the normal `defun` things, it registers the function as being public so that other back ends can inherit it.

nnoo-map-functions
:   This macro allows mapping of functions from the current back end to functions from the parent back ends.

    ```
    (nnoo-map-functions nndir
      (nnml-retrieve-headers 0 nndir-current-group 0 0)
      (nnmh-request-article 0 nndir-current-group 0 0))
    ```

    This means that when `nndir-retrieve-headers` is called, the first, third, and fourth parameters will be passed on to `nnml-retrieve-headers`, while the second parameter is set to the value of `nndir-current-group`.

nnoo-import
:   This macro allows importing functions from back ends. It should be the last thing in the source file, since it will only define functions that haven't already been defined.

    ```
    (nnoo-import nndir
      (nnmh
       nnmh-request-list
       nnmh-request-newgroups)
      (nnml))
    ```

> This means that calls to `nndir-request-list` should just be passed on to
> `nnmh-request-list`, while all public functions from `nnml` that haven't been
> defined in `nndir` yet should be defined now.

Below is a slightly shortened version of the `nndir` back end.

```
;;; nndir.el --- single directory newsgroup access for Gnus
;; Copyright (C) 1995,96 Free Software Foundation, Inc.

;;; Code:

(require 'nnheader)
(require 'nnmh)
(require 'nnml)
(require 'nnoo)
(eval-when-compile (require 'cl))

(nnoo-declare nndir
  nnml nnmh)

(defvoo nndir-directory nil
  "Where nndir will look for groups."
  nnml-current-directory nnmh-current-directory)

(defvoo nndir-nov-is-evil nil
  "*Non-nil means that nndir will never retrieve NOV headers."
  nnml-nov-is-evil)

(defvoo nndir-current-group ""
  nil
  nnml-current-group nnmh-current-group)
(defvoo nndir-top-directory nil nil nnml-directory nnmh-directory)
(defvoo nndir-get-new-mail nil nil nnml-get-new-mail nnmh-get-new-mail)

(defvoo nndir-status-string "" nil nnmh-status-string)
(defconst nndir-version "nndir 1.0")

;;; Interface functions.

(nnoo-define-basics nndir)

(deffoo nndir-open-server (server &optional defs)
  (setq nndir-directory
        (or (cadr (assq 'nndir-directory defs))
            server))
  (unless (assq 'nndir-directory defs)
    (push `(nndir-directory ,server) defs))
  (push `(nndir-current-group
          ,(file-name-nondirectory
             (directory-file-name nndir-directory)))
```

```
            defs)
      (push '(nndir-top-directory
              ,(file-name-directory (directory-file-name nndir-directory)))
            defs)
      (nnoo-change-server 'nndir server defs))

  (nnoo-map-functions nndir
    (nnml-retrieve-headers 0 nndir-current-group 0 0)
    (nnmh-request-article 0 nndir-current-group 0 0)
    (nnmh-request-group nndir-current-group 0 0)
    (nnmh-close-group nndir-current-group 0))

  (nnoo-import nndir
    (nnmh
     nnmh-status-message
     nnmh-request-list
     nnmh-request-newgroups))

  (provide 'nndir)
```

10.7.2.5 Hooking New Back Ends Into Gnus

Having Gnus start using your new back end is rather easy—you just declare it with
the `gnus-declare-backend` functions. This will enter the back end into the `gnus-valid-select-methods` variable.

`gnus-declare-backend` takes two parameters—the back end name and an arbitrary
number of *abilities*.

Here's an example:

```
(gnus-declare-backend "nnchoke" 'mail 'respool 'address)
```

The abilities can be:

mail   This is a mailish back end—followups should (probably) go via mail.

post   This is a newsish back end—followups should (probably) go via news.

post-mail
            This back end supports both mail and news.

none   This is neither a post nor mail back end—it's something completely different.

respool   It supports respooling—or rather, it is able to modify its source articles and
            groups.

address   The name of the server should be in the virtual server name. This is true for
            almost all back ends.

prompt-address
            The user should be prompted for an address when doing commands like B in the
            group buffer. This is true for back ends like `nntp`, but not `nnmbox`, for instance.

10.7.2.6 Mail-like Back Ends

One of the things that separate the mail back ends from the rest of the back ends is the heavy dependence by the mail back ends on common functions in 'nnmail.el'. For instance, here's the definition of nnml-request-scan:

```
(deffoo nnml-request-scan (&optional group server)
  (setq nnml-article-file-alist nil)
  (nnmail-get-new-mail 'nnml 'nnml-save-nov nnml-directory group))
```

It simply calls nnmail-get-new-mail with a few parameters, and nnmail takes care of all the moving and splitting of the mail.

This function takes four parameters.

method   This should be a symbol to designate which back end is responsible for the call.

exit-function
            This function should be called after the splitting has been performed.

temp-directory
            Where the temporary files should be stored.

group   This optional argument should be a group name if the splitting is to be performed for one group only.

nnmail-get-new-mail will call *back-end*-save-mail to save each article. *back-end*-active-number will be called to find the article number assigned to this article.

The function also uses the following variables: *back-end*-get-new-mail (to see whether to get new mail for this back end); and *back-end*-group-alist and *back-end*-active-file to generate the new active file. *back-end*-group-alist should be a group-active alist, like this:

```
(("a-group" (1 . 10))
 ("some-group" (34 . 39)))
```

10.7.3 Score File Syntax

Score files are meant to be easily parseable, but yet extremely mallable. It was decided that something that had the same read syntax as an Emacs Lisp list would fit that spec.

Here's a typical score file:

```
(("summary"
  ("win95" -10000 nil s)
  ("Gnus"))
 ("from"
  ("Lars" -1000))
 (mark -100))
```

BNF definition of a score file:

```
score-file      = "" / "(" *element ")"
element         = rule / atom
rule            = string-rule / number-rule / date-rule
string-rule     = "(" quote string-header quote space *string-match ")"
number-rule     = "(" quote number-header quote space *number-match ")"
```

```
date-rule       = "(" quote date-header quote space *date-match ")"
quote           = <ascii 34>
string-header   = "subject" / "from" / "references" / "message-id" /
                  "xref" / "body" / "head" / "all" / "followup"
number-header   = "lines" / "chars"
date-header     = "date"
string-match    = "(" quote <string> quote [ "" / [ space score [ "" /
                  space date [ "" / [ space string-match-t ] ] ] ] ] ")"
score           = "nil" / <integer>
date            = "nil" / <natural number>
string-match-t  = "nil" / "s" / "substring" / "S" / "Substring" /
                  "r" / "regex" / "R" / "Regex" /
                  "e" / "exact" / "E" / "Exact" /
                  "f" / "fuzzy" / "F" / "Fuzzy"
number-match    = "(" <integer> [ "" / [ space score [ "" /
                  space date [ "" / [ space number-match-t ] ] ] ] ] ")"
number-match-t  = "nil" / "=" / "<" / ">" / ">=" / "<="
date-match      = "(" quote <string> quote [ "" / [ space score [ "" /
                  space date [ "" / [ space date-match-t ] ] ] ] ")"
date-match-t    = "nil" / "at" / "before" / "after"
atom            = "(" [ required-atom / optional-atom ] ")"
required-atom   = mark / expunge / mark-and-expunge / files /
                  exclude-files / read-only / touched
optional-atom   = adapt / local / eval
mark            = "mark" space nil-or-number
nil-or-number   = "nil" / <integer>
expunge         = "expunge" space nil-or-number
mark-and-expunge = "mark-and-expunge" space nil-or-number
files           = "files" *[ space <string> ]
exclude-files   = "exclude-files" *[ space <string> ]
read-only       = "read-only" [ space "nil" / space "t" ]
adapt         = "adapt" [ space "ignore" / space "t" / space adapt-rule ]
adapt-rule      = "(" *[ <string> *[ "(" <string> <integer> ")" ] ")"
local           = "local" *[ space "(" <string> space <form> ")" ]
eval            = "eval" space <form>
space           = *[ " " / <TAB> / <NEWLINE> ]
```

Any unrecognized elements in a score file should be ignored, but not discarded.

As you can see, white space is needed, but the type and amount of white space is irrelevant. This means that formatting of the score file is left up to the programmer—if it's simpler to just spew it all out on one looong line, then that's ok.

The meaning of the various atoms are explained elsewhere in this manual (see Section 7.4 [Score File Format], page 172).

10.7.4 Headers

Internally Gnus uses a format for storing article headers that corresponds to the NOV format in a mysterious fashion. One could almost suspect that the author looked at the NOV specification and just shamelessly *stole* the entire thing, and one would be right.

Header is a severely overloaded term. "Header" is used in RFC 1036 to talk about lines in the head of an article (e.g., `From`). It is used by many people as a synonym for "head"— "the header and the body". (That should be avoided, in my opinion.) And Gnus uses a format internally that it calls "header", which is what I'm talking about here. This is a 9-element vector, basically, with each header (ouch) having one slot.

These slots are, in order: `number`, `subject`, `from`, `date`, `id`, `references`, `chars`, `lines`, `xref`, and `extra`. There are macros for accessing and setting these slots—they all have predictable names beginning with `mail-header-` and `mail-header-set-`, respectively.

All these slots contain strings, except the `extra` slot, which contains an alist of header/value pairs (see Section 3.1.2 [To From Newsgroups], page 39).

10.7.5 Ranges

GNUS introduced a concept that I found so useful that I've started using it a lot and have elaborated on it greatly.

The question is simple: If you have a large amount of objects that are identified by numbers (say, articles, to take a *wild* example) that you want to qualify as being "included", a normal sequence isn't very useful. (A 200,000 length sequence is a bit long-winded.)

The solution is as simple as the question: You just collapse the sequence.

```
(1 2 3 4 5 6 10 11 12)
```

is transformed into

```
((1 . 6) (10 . 12))
```

To avoid having those nasty '`(13 . 13)`' elements to denote a lonesome object, a '`13`' is a valid element:

```
((1 . 6) 7 (10 . 12))
```

This means that comparing two ranges to find out whether they are equal is slightly tricky:

```
((1 . 5) 7 8 (10 . 12))
```

and

```
((1 . 5) (7 . 8) (10 . 12))
```

are equal. In fact, any non-descending list is a range:

```
(1 2 3 4 5)
```

is a perfectly valid range, although a pretty long-winded one. This is also valid:

```
(1 . 5)
```

and is equal to the previous range.

Here's a BNF definition of ranges. Of course, one must remember the semantic requirement that the numbers are non-descending. (Any number of repetition of the same number is allowed, but apt to disappear in range handling.)

```
range            = simple-range / normal-range
simple-range     = "(" number " . " number ")"
normal-range     = "(" start-contents ")"
contents         = "" / simple-range *[ " " contents ] /
                   number *[ " " contents ]
```

Gnus currently uses ranges to keep track of read articles and article marks. I plan on implementing a number of range operators in C if The Powers That Be are willing to let me. (I haven't asked yet, because I need to do some more thinking on what operators I need to make life totally range-based without ever having to convert back to normal sequences.)

10.7.6 Group Info

Gnus stores all permanent info on groups in a *group info* list. This list is from three to six elements (or more) long and exhaustively describes the group.

Here are two example group infos; one is a very simple group while the second is a more complex one:

```
("no.group" 5 ((1 . 54324)))

("nnml:my.mail" 3 ((1 . 5) 9 (20 . 55))
                  ((tick (15 . 19)) (replied 3 6 (19 . 3)))
                  (nnml "")
                  ((auto-expire . t) (to-address . "ding@gnus.org")))
```

The first element is the *group name*—as Gnus knows the group, anyway. The second element is the *subscription level*, which normally is a small integer. (It can also be the *rank*, which is a cons cell where the `car` is the level and the `cdr` is the score.) The third element is a list of ranges of read articles. The fourth element is a list of lists of article marks of various kinds. The fifth element is the select method (or virtual server, if you like). The sixth element is a list of *group parameters*, which is what this section is about.

Any of the last three elements may be missing if they are not required. In fact, the vast majority of groups will normally only have the first three elements, which saves quite a lot of cons cells.

Here's a BNF definition of the group info format:

```
info          = "(" group space ralevel space read
                [ "" / [ space marks-list [ "" / [ space method [ "" /
                space parameters ] ] ] ] ] ")"
group         = quote <string> quote
ralevel       = rank / level
level         = <integer in the range of 1 to inf>
rank          = "(" level "." score ")"
score         = <integer in the range of 1 to inf>
read          = range
marks-lists   = nil / "(" *marks ")"
marks         = "(" <string> range ")"
method        = "(" <string> *elisp-forms ")"
parameters    = "(" *elisp-forms ")"
```

Actually that 'marks' rule is a fib. A 'marks' is a '`<string>`' consed on to a '`range`', but that's a bitch to say in pseudo-BNF.

If you have a Gnus info and want to access the elements, Gnus offers a series of macros for getting/setting these elements.

```
gnus-info-group
gnus-info-set-group
```
            Get/set the group name.

```
gnus-info-rank
gnus-info-set-rank
```
            Get/set the group rank (see Section 2.7 [Group Score], page 20).

```
gnus-info-level
gnus-info-set-level
```
            Get/set the group level.

```
gnus-info-score
gnus-info-set-score
```
            Get/set the group score (see Section 2.7 [Group Score], page 20).

```
gnus-info-read
gnus-info-set-read
```
            Get/set the ranges of read articles.

```
gnus-info-marks
gnus-info-set-marks
```
            Get/set the lists of ranges of marked articles.

```
gnus-info-method
gnus-info-set-method
```
            Get/set the group select method.

```
gnus-info-params
gnus-info-set-params
```
            Get/set the group parameters.

All the getter functions take one parameter—the info list. The setter functions take two parameters—the info list and the new value.

The last three elements in the group info aren't mandatory, so it may be necessary to extend the group info before setting the element. If this is necessary, you can just pass on a non-`nil` third parameter to the three final setter functions to have this happen automatically.

10.7.7 Extended Interactive

Gnus extends the standard Emacs `interactive` specification slightly to allow easy use of the symbolic prefix (see Section 8.3 [Symbolic Prefixes], page 190). Here's an example of how this is used:

```
(defun gnus-summary-increase-score (&optional score symp)
  (interactive (gnus-interactive "P\ny"))
  ...
  )
```

The best thing to do would have been to implement `gnus-interactive` as a macro which would have returned an `interactive` form, but this isn't possible since Emacs checks whether a function is interactive or not by simply doing an `assq` on the lambda form. So,

instead we have `gnus-interactive` function that takes a string and returns values that are usable to `interactive`.

This function accepts (almost) all normal `interactive` specs, but adds a few more.

'y'   The current symbolic prefix—the `gnus-current-prefix-symbol` variable.

'Y'   A list of the current symbolic prefixes—the `gnus-current-prefix-symbol` variable.

'A'   The current article number—the `gnus-summary-article-number` function.

'H'   The current article header—the `gnus-summary-article-header` function.

'g'   The current group name—the `gnus-group-group-name` function.

10.7.8 Emacs/XEmacs Code

While Gnus runs under Emacs, XEmacs and Mule, I decided that one of the platforms must be the primary one. I chose Emacs. Not because I don't like XEmacs or Mule, but because it comes first alphabetically.

This means that Gnus will byte-compile under Emacs with nary a warning, while XEmacs will pump out gigabytes of warnings while byte-compiling. As I use byte-compilation warnings to help me root out trivial errors in Gnus, that's very useful.

I've also consistently used Emacs function interfaces, but have used Gnusey aliases for the functions. To take an example: Emacs defines a `run-at-time` function while XEmacs defines a `start-itimer` function. I then define a function called `gnus-run-at-time` that takes the same parameters as the Emacs `run-at-time`. When running Gnus under Emacs, the former function is just an alias for the latter. However, when running under XEmacs, the former is an alias for the following function:

```
(defun gnus-xmas-run-at-time (time repeat function &rest args)
  (start-itimer
   "gnus-run-at-time"
   `(lambda ()
      (,function ,@args))
   time repeat))
```

This sort of thing has been done for bunches of functions. Gnus does not redefine any native Emacs functions while running under XEmacs—it does this `defalias` thing with Gnus equivalents instead. Cleaner all over.

In the cases where the XEmacs function interface was obviously cleaner, I used it instead. For example `gnus-region-active-p` is an alias for `region-active-p` in XEmacs, whereas in Emacs it is a function.

Of course, I could have chosen XEmacs as my native platform and done mapping functions the other way around. But I didn't. The performance hit these indirections impose on Gnus under XEmacs should be slight.

10.7.9 Various File Formats

10.7.9.1 Active File Format

The active file lists all groups available on the server in question. It also lists the highest and lowest current article numbers in each group.

Here's an excerpt from a typical active file:

```
soc.motss 296030 293865 y
alt.binaries.pictures.fractals 3922 3913 n
comp.sources.unix 1605 1593 m
comp.binaries.ibm.pc 5097 5089 y
no.general 1000 900 y
```

Here's a pseudo-BNF definition of this file:

```
active      = *group-line
group-line  = group spc high-number spc low-number spc flag <NEWLINE>
group       = <non-white-space string>
spc         = " "
high-number = <non-negative integer>
low-number  = <positive integer>
flag        = "y" / "n" / "m" / "j" / "x" / "=" group
```

For a full description of this file, see the manual pages for 'innd', in particular 'active(5)'.

10.7.9.2 Newsgroups File Format

The newsgroups file lists groups along with their descriptions. Not all groups on the server have to be listed, and not all groups in the file have to exist on the server. The file is meant purely as information to the user.

The format is quite simple; a group name, a tab, and the description. Here's the definition:

```
newsgroups   = *line
line         = group tab description <NEWLINE>
group        = <non-white-space string>
tab          = <TAB>
description   = <string>
```

10.8 Emacs for Heathens

Believe it or not, but some people who use Gnus haven't really used Emacs much before they embarked on their journey on the Gnus Love Boat. If you are one of those unfortunates whom "`M-C-a`", "kill the region", and "set `gnus-flargblossen` to an alist where the key is a regexp that is used for matching on the group name" are magical phrases with little or no meaning, then this appendix is for you. If you are already familiar with Emacs, just ignore this and go fondle your cat instead.

10.8.1 Keystrokes

- Q: What is an experienced Emacs user?
- A: A person who wishes that the terminal had pedals.

Yes, when you use Emacs, you are apt to use the control key, the shift key and the meta key a lot. This is very annoying to some people (notably `vile` users), and the rest of us just love the hell out of it. Just give up and submit. Emacs really does stand for "Escape-Meta-Alt-Control-Shift", and not "Editing Macros", as you may have heard from other disreputable sources (like the Emacs author).

The shift keys are normally located near your pinky fingers, and are normally used to get capital letters and stuff. You probably use it all the time. The control key is normally marked "CTRL" or something like that. The meta key is, funnily enough, never marked as such on any keyboard. The one I'm currently at has a key that's marked "Alt", which is the meta key on this keyboard. It's usually located somewhere to the left hand side of the keyboard, usually on the bottom row.

Now, us Emacs people don't say "press the meta-control-m key", because that's just too inconvenient. We say "press the `M-C-m` key". `M-` is the prefix that means "meta" and "C-" is the prefix that means "control". So "press `C-k`" means "press down the control key, and hold it down while you press `k`". "Press `M-C-k`" means "press down and hold down the meta key and the control key and then press `k`". Simple, ay?

This is somewhat complicated by the fact that not all keyboards have a meta key. In that case you can use the "escape" key. Then `M-k` means "press escape, release escape, press `k`". That's much more work than if you have a meta key, so if that's the case, I respectfully suggest you get a real keyboard with a meta key. You can't live without it.

10.8.2 Emacs Lisp

Emacs is the King of Editors because it's really a Lisp interpreter. Each and every key you tap runs some Emacs Lisp code snippet, and since Emacs Lisp is an interpreted language, that means that you can configure any key to run any arbitrary code. You just, like, do it.

Gnus is written in Emacs Lisp, and is run as a bunch of interpreted functions. (These are byte-compiled for speed, but it's still interpreted.) If you decide that you don't like the way Gnus does certain things, it's trivial to have it do something a different way. (Well, at least if you know how to write Lisp code.) However, that's beyond the scope of this manual, so we are simply going to talk about some common constructs that you normally use in your '`.emacs`' file to customize Gnus.

If you want to set the variable `gnus-florgbnize` to four (4), you write the following:

```
(setq gnus-florgbnize 4)
```

This function (really "special form") `setq` is the one that can set a variable to some value. This is really all you need to know. Now you can go and fill your `.emacs` file with lots of these to change how Gnus works.

If you have put that thing in your `.emacs` file, it will be read and `eval`ed (which is lisp-ese for "run") the next time you start Emacs. If you want to change the variable right away, simply say `C-x C-e` after the closing parenthesis. That will `eval` the previous "form", which is a simple `setq` statement here.

Go ahead—just try it, if you're located at your Emacs. After you `C-x C-e`, you will see '4' appear in the echo area, which is the return value of the form you `eval`ed.

Some pitfalls:

If the manual says "set `gnus-read-active-file` to `some`", that means:

```
(setq gnus-read-active-file 'some)
```

On the other hand, if the manual says "set `gnus-nntp-server` to '`nntp.ifi.uio.no`'", that means:

```
(setq gnus-nntp-server "nntp.ifi.uio.no")
```

So be careful not to mix up strings (the latter) with symbols (the former). The manual is unambiguous, but it can be confusing.

10.9 Frequently Asked Questions

This is the Gnus Frequently Asked Questions list. If you have a Web browser, the official hypertext version is at '`http://www.ccs.neu.edu/software/gnus/`', and has probably been updated since you got this manual.

10.9.1 Installation

- Q1.1 What is the latest version of Gnus?

  The latest (and greatest) version is 5.0.10. You might also run across something called *September Gnus*. September Gnus is the alpha version of the next major release of Gnus. It is currently not stable enough to run unless you are prepared to debug lisp.

- Q1.2 Where do I get Gnus?

  Any of the following locations:
  - '`ftp://ftp.ifi.uio.no/pub/emacs/gnus/gnus.tar.gz`'
  - '`ftp://ftp.pilgrim.umass.edu/pub/misc/ding/`'
  - '`gopher://gopher.pilgrim.umass.edu/11/pub/misc/ding/`'
  - '`ftp://aphrodite.nectar.cs.cmu.edu/pub/ding-gnus/`'
  - '`ftp://ftp.solace.mh.se:/pub/gnu/elisp/`'

- Q1.3 Which version of Emacs do I need?

  At least GNU Emacs 19.28, or XEmacs 19.12 is recommended. GNU Emacs 19.25 has been reported to work under certain circumstances, but it doesn't *officially* work on it. 19.27 has also been reported to work. Gnus has been reported to work under OS/2 as well as Unix.

- Q1.4 Where is timezone.el?

  Upgrade to XEmacs 19.13. In earlier versions of XEmacs this file was placed with Gnus 4.1.3, but that has been corrected.

- Q1.5 When I run Gnus on XEmacs 19.13 I get weird error messages.

  You're running an old version of Gnus. Upgrade to at least version 5.0.4.

- Q1.6 How do I unsubscribe from the Mailing List?

  Send an e-mail message to '`ding-request@ifi.uio.no`' with the magic word *unsubscribe* somewhere in it, and you will be removed.

  If you are reading the digest version of the list, send an e-mail message to '`ding-rn-digests-d-request@moe.shore.net`' with *unsubscribe* as the subject and you will be removed.

- Q1.7 How do I run Gnus on both Emacs and XEmacs?

  The basic answer is to byte-compile under XEmacs, and then you can run under either Emacsen. There is, however, a potential version problem with easymenu.el with Gnu Emacs prior to 19.29.

  Per Abrahamsen <abraham@dina.kvl.dk> writes :
  The internal easymenu.el interface changed between 19.28 and 19.29 in order to make it possible to create byte compiled files that can be shared between Gnu Emacs and XEmacs. The change is upward compatible, but not downward compatible. This gives the following compatibility table:

```
Compiled with:  | Can be used with:
----------------+------------------------------------
19.28           | 19.28      19.29
19.29           |            19.29          XEmacs
XEmacs          |            19.29          XEmacs
```

If you have Gnu Emacs 19.28 or earlier, or XEmacs 19.12 or earlier, get a recent version of auc-menu.el from '`ftp://ftp.iesd.auc.dk/pub/emacs-lisp/auc-menu.el`', and install it under the name easymenu.el somewhere early in your load path.

- Q1.8 What resources are available?

  There is the newsgroup Gnu.emacs.gnus. Discussion of Gnus 5.x is now taking place there. There is also a mailing list, send mail to '`ding-request@ifi.uio.no`' with the magic word *subscribe* somewhere in it.

  NOTE: the traffic on this list is heavy so you may not want to be on it (unless you use Gnus as your mailer reader, that is). The mailing list is mainly for developers and testers.

  Gnus has a home World Wide Web page at
  '`http://www.gnus.org/`'.

  Gnus has a write up in the X Applications FAQ at
  '`http://www.ee.ryerson.ca:8080/~elf/xapps/Q-III.html`'.

  The Gnus manual is also available on the World Wide Web. The canonical source is in Norway at
  '`http://www.gnus.org/manual/gnus_toc.html`'.

  There are three mirrors in the United States:

  1. '`http://www.miranova.com/gnus-man/`'
  2. '`http://www.pilgrim.umass.edu/pub/misc/ding/manual/gnus_toc.html`'
  3. '`http://www.rtd.com/~woo/gnus/`'

  PostScript copies of the Gnus Reference card are available from
  '`ftp://ftp.cs.ualberta.ca/pub/oolog/gnus/`'. They are mirrored at
  '`ftp://ftp.pilgrim.umass.edu/pub/misc/ding/refcard/`' in the United States. And
  '`ftp://marvin.fkphy.uni-duesseldorf.de/pub/gnus/`' in Germany.

  An online version of the Gnus FAQ is available at
  '`http://www.miranova.com/~steve/gnus-faq.html`'. Off-line formats are also available:
  ASCII: '`ftp://ftp.miranova.com/pub/gnus/gnus-faq`'
  PostScript: '`ftp://ftp.miranova.com/pub/gnus/gnus-faq.ps`'.

- Q1.9 Gnus hangs on connecting to NNTP server

  I am running XEmacs on SunOS and Gnus prints a message about Connecting to NNTP server and then just hangs.

  Ben Wing <wing@netcom.com> writes :
  I wonder if you're hitting the infamous *libresolv* problem. The basic problem is that under SunOS you can compile either with DNS or NIS name lookup libraries but not both. Try substituting the IP address and see if that works; if so, you need to download the sources and recompile.

- Q1.10 Mailcrypt 3.4 doesn't work

  This problem is verified to still exist in Gnus 5.0.9 and Mailcrypt 3.4. The answer comes from Peter Arius <arius@immd2.informatik.uni-erlangen.de>.

  I found out that mailcrypt uses `gnus-eval-in-buffer-window`, which is a macro. It seems as if you have compiled mailcrypt with plain old GNUS in load path, and the XEmacs byte compiler has inserted that macro definition into '`mc-toplev.elc`'. The solution is to recompile '`mc-toplev.el`' with Gnus 5 in load-path, and it works fine.

  Steve Baur <steve@miranova.com> adds :
  The problem also manifests itself if neither GNUS 4 nor Gnus 5 is in the load-path.

- Q1.11 What other packages work with Gnus?

  – Mailcrypt.

    Mailcrypt is an Emacs interface to PGP. It works, it installs without hassle, and integrates very easily. Mailcrypt can be obtained from '`ftp://cag.lcs.mit.edu/pub/patl/mailcrypt-3.4.tar.gz`'.

  – Tools for Mime.

    Tools for Mime is an Emacs MUA interface to MIME. Installation is a two-step process unlike most other packages, so you should be prepared to move the byte-compiled code somewhere. There are currently two versions of this package available. It can be obtained from '`ftp://ftp.jaist.ac.jp/pub/GNU/elisp/`'. Be sure to apply the supplied patch. It works with Gnus through version 5.0.9. In order for all dependencies to work correctly the load sequence is as follows:

    ```
    (load "tm-setup")
    (load "gnus")
    (load "mime-compose")
    ```

    NOTE: Loading the package disables citation highlighting by default. To get the old behavior back, use the `M-t` command.

10.9.2 Customization

- Q2.1 Custom Edit does not work under XEmacs

  The custom package has not been ported to XEmacs.

- Q2.2 How do I quote messages?

  I see lots of messages with quoted material in them. I am wondering how to have Gnus do it for me.

  This is Gnus, so there are a number of ways of doing this. You can use the built-in commands to do this. There are the `F` and `R` keys from the summary buffer which automatically include the article being responded to. These commands are also selectable as *Followup and Yank* and *Reply and Yank* in the Post menu.

  `C-c C-y` grabs the previous message and prefixes each line with `ail-indentation-spaces` spaces or `mail-yank-prefix` if that is non-nil, unless you have set your own `mail-citation-hook`, which will be called to do the job.

  You might also consider the Supercite package, which allows for pretty arbitrarily complex quoting styles. Some people love it, some people hate it.

- Q2.3 How can I keep my nnvirtual:* groups sorted?

  How can I most efficiently arrange matters so as to keep my nnvirtual:* (etc) groups at the top of my group selection buffer, whilst keeping everything sorted in alphabetical order.

  If you don't subscribe often to new groups then the easiest way is to first sort the groups and then manually kill and yank the virtuals wherever you want them.

- Q2.4 Any good suggestions on stuff for an all.SCORE file?

  Here is a collection of suggestions from the Gnus mailing list.

  1. From "Dave Disser" <disser@sdd.hp.com>
     I like blasting anything without lowercase letters. Weeds out most of the make $$ fast, as well as the lame titles like "IBM" and "HP-UX" with no further description.

     ```
     (("Subject"
       ("^\\(Re: \\)?[^a-z]*$" -200 nil R)))
     ```

  2. From "Peter Arius" <arius@immd2.informatik.uni-erlangen.de>
     The most vital entries in my (still young) all.SCORE:

     ```
     (("xref"
       ("alt.fan.oj-simpson" -1000 nil s))
      ("subject"
       (concat "\\<\\(make\\|fast\\|big\\)\\s-*"
               "\\(money\\|cash\\|bucks?\\)\\>"
        -1000 nil r)
       ("$$$$" -1000 nil s)))
     ```

  3. From "Per Abrahamsen" <abraham@dina.kvl.dk>

     ```
     (("subject"
       ;; CAPS OF THE WORLD, UNITE
       ("^..[^a-z]+$" -1 nil R)
       ;; $$$ Make Money $$$ (Try work)
       ("$" -1 nil s)
       ;; I'm important! And I have exclamation marks to prove it!
       ("!" -1 nil s)))
     ```

  4. From "heddy boubaker" <boubaker@cenatls.cena.dgac.fr>
     I would like to contribute with mine.

     ```
     (
      (read-only t)
      ("subject"
       ;; ALL CAPS SUBJECTS
       ("^\\([Rr][Ee]: +\\)?[^a-z]+$" -1 nil R)
       ;; $$$ Make Money $$$
       ("$$" -10 nil s)
       ;; Empty subjects are worthless!
       ("^ *\\([(<]none[>)]\\|(no subject\\( given\\)?)\\)? *$"
        -10 nil r)
       ;; Sometimes interesting announces occur!
       ("ANN?OU?NC\\(E\\|ING\\)" +10 nil r)
       ;; Some people think they're on mailing lists
     ```

```
          ("\\(un\\)?sub?scribe" -100 nil r)
          ;; Stop Micro$oft NOW!!
          ;; ("concat" used to avoid overfull box.)
          (concat "\\(m\\(icro\\)?[s$]\\(oft\\|lot\\)?-?\\)?"
                  "wind?\\(ows\\|aube\\|oze\\)?[- ]*"
                  "\\('?95\\|NT\\|3[.]1\\|32\\)" -1001 nil r)
          ;; I've nothing to buy
          ("\\(for\\|4\\)[- ]*sale" -100 nil r)
          ;; SELF-DISCIPLINED people
          ("\\[[^a-z0-9 \t\n][^a-z0-9 \t\n]\\]" +100 nil r)
          )
         ("from"
          ;; To keep track of posters from my site
          (".dgac.fr" +1000 nil s))
         ("followup"
          ;; Keep track of answers to my posts
          ("boubaker" +1000 nil s))
         ("lines"
          ;; Some people have really nothing to say!!
          (1 -10 nil <=))
         (mark -100)
         (expunge -1000)
         )
```

5. From "Christopher Jones" <cjones@au.oracle.com>
   The sample 'all.SCORE' files from Per and boubaker could be augmented with:

```
        (("subject"
          ;; No junk mail please!
          ("please ignore" -500 nil s)
          ("test" -500 nil e))
         )
```

6. From "Brian Edmonds" <edmonds@cs.ubc.ca>
   Augment any of the above with a fast method of scoring down excessively cross
   posted articles.

```
          ("xref"
           ;; the more cross posting, the exponentially worse the article
           ("^xref: \\S-+ \\S-+ \\S-+ \\S-+" -1 nil r)
           ("^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+" -2 nil r)
           ("^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+" -4 nil r)
           ("^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+" -8 nil r)
           ("^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
             -16 nil r)
           (concat "^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
                   " \\S-+ \\S-+"
             -32 nil r)
           (concat "^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
                   " \\S-+ \\S-+ \\S-+" -64 nil r)
           (concat "^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
                   " \\S-+ \\S-+ \\S-+ \\S-+" -128 nil r)
```

```
                    (concat "^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
                            " \\S-+ \\S-+ \\S-+ \\S-+ \\S-+" -256 nil r)
                    (concat "^xref: \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+ \\S-+"
                            " \\S-+" \\S-+ \\S-+ \\S-+ \\S-+ \\S-+" -512 nil r))
```

- Q2.5 What do I use to yank-through when replying?

  You should probably reply and followup with `R` and `F`, instead of `r` and `f`, which solves your problem. But you could try something like:

  ```
  (defconst mail-yank-ignored-headers
    "^.*:"
    "Delete these headers from message when it's inserted in reply.")
  ```

- Q2.6 I don't like the default WWW browser

  Now when choosing an URL Gnus starts up a W3 buffer, I would like it to always use Netscape (I don't browse in text-mode ;-).

  1. Activate 'Customize...' from the 'Help' menu.
  2. Scroll down to the 'WWW Browser' field.
  3. Click 'mouse-2' on 'WWW Browser'.
  4. Select 'Netscape' from the pop up menu.
  5. Press 'C-c C-c'

  If you are using XEmacs then to specify Netscape do

  ```
  (setq gnus-button-url 'gnus-netscape-open-url)
  ```

- Q2.7 What, if any, relation is between "ask-server" and "(setq gnus-read-active-file 'some)"?

  In order for Gnus to show you the complete list of newsgroups, it will either have to either store the list locally, or ask the server to transmit the list. You enable the first with

  ```
  (setq gnus-save-killed-list t)
  ```

  and the second with

  ```
  (setq gnus-read-active-file t)
  ```

  If both are disabled, Gnus will not know what newsgroups exists. There is no option to get the list by casting a spell.

- Q2.8 Moving between groups is slow.

  Per Abrahamsen <abraham@dina.kvl.dk> writes:

  Do you call `define-key` or something like that in one of the summary mode hooks? This would force Emacs to recalculate the keyboard shortcuts. Removing the call should speed up `M-x gnus-summary-mode RET` by a couple of orders of magnitude. You can use

  ```
  (define-key gnus-summary-mode-map KEY COMMAND)
  ```

  in your '.gnus' instead.

10.9.3 Reading News

- Q3.1 How do I convert my kill files to score files?

  Ethan Bradford (`ethanb@ptolemy.astro.washington.edu`) write a kill-to-score translator. It is available from
  '`http://baugi.ifi.uio.no/~larsi/ding-various/gnus-kill-to-score.el`'.

- Q3.2 My news server has a lot of groups, and killing groups is painfully slow.

  Don't do that then. The best way to get rid of groups that should be dead is to edit your newsrc directly. This problem will be addressed in the near future.

- Q3.3 How do I use an NNTP server with authentication?

  Put the following into your .gnus:

      (add-hook 'nntp-server-opened-hook 'nntp-send-authinfo)

- Q3.4 Not reading the first article.

  How do I avoid reading the first article when a group is selected?

  1. Use *RET* to select the group instead of *SPC*.

  2. `(setq gnus-auto-select first nil)`

  3. Luis Fernandes <elf@mailhost.ee.ryerson.ca>writes:
     This is what I use...customize as necessary...

     ```
     ;;; Don't auto-select first article if reading sources, or
     ;;; archives or jobs postings, etc. and just display the
     ;;; summary buffer
     (add-hook 'gnus-select-group-hook
       (function
        (lambda ()
          (cond ((string-match "sources" gnus-newsgroup-name)
           (setq gnus-auto-select-first nil))
          ((string-match "jobs" gnus-newsgroup-name)
     (setq gnus-auto-select-first nil))
          ((string-match "comp\\.archives" gnus-newsgroup-name)
           (setq gnus-auto-select-first nil))
          ((string-match "reviews" gnus-newsgroup-name)
           (setq gnus-auto-select-first nil))
          ((string-match "announce" gnus-newsgroup-name)
           (setq gnus-auto-select-first nil))
          ((string-match "binaries" gnus-newsgroup-name)
           (setq gnus-auto-select-first nil))
          (t
           (setq gnus-auto-select-first t))))))
     ```

  4. Per Abrahamsen <abraham@dina.kvl.dk> writes:
     Another possibility is to create an '`all.binaries.all.SCORE`' file like this:

     ```
     ((local
        (gnus-auto-select-first nil)))
     ```

     and insert

     ```
     (setq gnus-auto-select-first t)
     ```

     in your '**.gnus**'.

- Q3.5 Why aren't BBDB known posters marked in the summary buffer?

  Brian Edmonds `<edmonds@cs.ubc.ca>` writes:

  Due to changes in Gnus 5.0, '`bbdb-gnus.el`' no longer marks known posters in the summary buffer. An updated version, '`gnus-bbdb.el`' is available at the locations listed below. This package also supports autofiling of incoming mail to folders specified in the BBDB. Extensive instructions are included as comments in the file.

  Send mail to '`majordomo@edmonds.home.cs.ubc.ca`' with the following line in the body of the message: *get misc gnus-bbdb.el*.

  Or get it from the World Wide Web:
  '`http://www.cs.ubc.ca/spider/edmonds/gnus-bbdb.el`'.

10.9.4 Reading Mail

- Q4.1 What does the message "Buffer has changed on disk" mean in a mail group?

  Your filter program should not deliver mail directly to your folders, instead it should put the mail into spool files. Gnus will then move the mail safely from the spool files into the folders. This will eliminate the problem. Look it up in the manual, in the section entitled "Mail & Procmail".

- Q4.2 How do you make articles un-expirable?

  I am using nnml to read news and have used `gnus-auto-expirable-newsgroups` to automagically expire articles in some groups (Gnus being one of them). Sometimes there are interesting articles in these groups that I want to keep. Is there any way of explicitly marking an article as un-expirable - that is mark it as read but not expirable?

  Use `u`, `!`, `d` or `M-u` in the summary buffer. You just remove the `E` mark by setting some other mark. It's not necessary to tick the articles.

- Q4.3 How do I delete bogus nnml: groups?

  My problem is that I have various mail (nnml) groups generated while experimenting with Gnus. How do I remove them now? Setting the level to 9 does not help. Also `gnus-group-check-bogus-groups` does not recognize them.

  Removing mail groups is tricky at the moment. (It's on the to-do list, though.) You basically have to kill the groups in Gnus, shut down Gnus, edit the active file to exclude these groups, and probably remove the nnml directories that contained these groups as well. Then start Gnus back up again.

- Q4.4 What happened to my new mail groups?

  I got new mail, but I have never seen the groups they should have been placed in.

  They are probably there, but as zombies. Press `A z` to list zombie groups, and then subscribe to the groups you want with `u`. This is all documented quite nicely in the user's manual.

- Q4.5 Not scoring mail groups

  How do you *totally* turn off scoring in mail groups?

  Use an nnbabyl:all.SCORE (or nnmh, or nnml, or whatever) file containing:

  ```
  ((adapt ignore)
   (local (gnus-use-scoring nil))
   (exclude-files "all.SCORE"))
  ```

11 Index

%

(

*

.

/

>

<

A

B

C

M

N

T

U

V

W

X

Z

12 Key Index

B

C

D

E

F

G

H

J

K

L

M

R

S

T

U

Short Contents

Table of Contents