

Version Management with CVS

for cvs 1.9.10

Per Cederqvist et al

Copyright © 1992, 1993 Signum Support AB

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the section entitled “GNU General Public License” is included exactly as in the original, and provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that the section entitled “GNU General Public License” and this permission notice may be included in translations approved by the Free Software Foundation instead of in the original English.

About this manual

Up to this point, one of the weakest parts of CVS has been the documentation. CVS is a complex program. Previous versions of the manual were written in the manual page format, which is not really well suited for such a complex program.

When writing this manual, I had several goals in mind:

- No knowledge of RCS should be necessary.
- No previous knowledge of revision control software should be necessary. All terms, such as *revision numbers*, *revision trees* and *merging* are explained as they are introduced.
- The manual should concentrate on the things CVS users want to do, instead of what the CVS commands can do. The first part of this manual leads you through things you might want to do while doing development, and introduces the relevant CVS commands as they are needed.
- Information should be easy to find. In the reference manual in the appendices almost all information about every CVS command is gathered together. There is also an extensive index, and a lot of cross references.

Checklist for the impatient reader

CVS is a complex system. You will need to read the manual to be able to use all of its capabilities. There are dangers that can easily be avoided if you know about them, and this manual tries to warn you about them. This checklist is intended to help you avoid the dangers without reading the entire manual. If you intend to read the entire manual you can skip this table.

Binary files

CVS can handle binary files, but you must have RCS release 5.5 or later and a release of GNU diff that supports the '-a' flag (release 1.15 and later are OK). You must also configure both RCS and CVS to handle binary files when you install them.

Keyword substitution can be a source of trouble with binary files. See Chapter 16 [Keyword substitution], page 73, for solutions.

The `admin` command

Careless use of the `admin` command can cause CVS to cease working. See Section A.6 [admin], page 91, before trying to use it.

Credits

Roland Pesch, then of Cygnus Support <roland@wrs.com> wrote the manual pages which were distributed with CVS 1.3. Much of their text was copied into this manual. He also read an early draft of this manual and contributed many ideas and corrections.

The mailing-list `info-cvs` is sometimes informative. I have included information from postings made by the following persons: David G. Grubbs <dgg@think.com>.

Some text has been extracted from the man pages for RCS.

The CVS FAQ by David G. Grubbs has provided useful material. The FAQ is no longer maintained, however, and this manual is about the closest thing there is to a successor (with respect to documenting how to use CVS, at least).

In addition, the following persons have helped by telling me about mistakes I've made:

Roxanne Brunskill <rbrunski@datap.ca>,
 Kathy Dyer <dyer@phoenix.ocf.llnl.gov>,
 Karl Pingle <pingle@acuson.com>,
 Thomas A Peterson <tap@src.honeywell.com>,
 Inge Wallin <ingwa@signum.se>,
 Dirk Koschuetzki <koschuet@fmi.uni-passau.de>
 and Michael Brown <brown@wi.extrel.com>.

The list of contributors here is not comprehensive; for a more complete list of who has contributed to this manual see the file 'doc/ChangeLog' in the CVS source distribution.

BUGS

Neither CVS nor this manual is perfect, and they probably never will be. If you are having trouble using CVS, or think you have found a bug, there are a number of things you can do about it. Note that if the manual is unclear, that can be considered a bug in the manual, so these problems are often worth doing something about as well as problems with CVS itself.

- If you want someone to help you and fix bugs that you report, there are companies which will do that for a fee. Two such companies are:

Signum Support AB
 Box 2044
 S-580 02 Linköping
 Sweden
 Email: info@signum.se
 Phone: +46 (0)13 - 21 46 00
 Fax: +46 (0)13 - 21 47 00
<http://www.signum.se/>

Cyclic Software
 United States of America
<http://www.cyclic.com/>
info@cyclic.com

- If you got CVS through a distributor, such as an operating system vendor or a vendor of freeware CD-ROMs, you may wish to see whether the distributor provides support. Often, they will provide no support or minimal support, but this may vary from distributor to distributor.
- If you have the skills and time to do so, you may wish to fix the bug yourself. If you wish to submit your fix for inclusion in future releases of CVS, see the file HACKING in the CVS source distribution. It contains much more information on the process of submitting fixes.
- There may be resources on the net which can help. Two good places to start are:

<http://www.cyclic.com>
particularly the Unsupported Resources page
<http://www.loria.fr/~molli/cvs-index.html>

If you are so inspired, increasing the information available on the net is likely to be appreciated. For example, before the standard CVS distribution worked on Windows 95, there was a web page with some explanation and patches for running CVS on Windows 95, and various people helped out by mentioning this page on mailing lists or newsgroups when the subject came up.

- It is also possible to report bugs to **bug-cvs**. Note that someone may or may not want to do anything with your bug report—if you need a solution consider one of the options mentioned above. People probably do want to hear about bugs which are particularly severe in consequences and/or easy to fix, however. You can also increase your odds by being as clear as possible about the exact nature of the bug and any other relevant information. The way to report bugs is to send email to **bug-cvs@prep.ai.mit.edu**. Note that submissions to **bug-cvs** may be distributed under the terms of the GNU Public License, so if you don't like this, don't submit them. There is usually no justification for sending mail directly to one of the CVS maintainers rather than to **bug-cvs**; those maintainers who want to hear about such bug reports read **bug-cvs**. Also note that sending a bug report to other mailing lists or newsgroups is *not* a substitute for sending it to **bug-cvs**. It is fine to discuss CVS bugs on whatever forum you prefer, but there are not necessarily any maintainers reading bug reports sent anywhere except **bug-cvs**.

People often ask if there is a list of known bugs or whether a particular bug is a known one. The file **BUGS** in the CVS source distribution is one list of known bugs, but it doesn't necessarily try to be comprehensive. Perhaps there will never be a comprehensive, detailed list of known bugs.

1 What is CVS?

CVS is a version control system. Using it, you can record the history of your source files.

For example, bugs sometimes creep in when software is modified, and you might not detect the bug until a long time after you make the modification. With CVS, you can easily retrieve old versions to see exactly which change caused the bug. This can sometimes be a big help.

You could of course save every version of every file you have ever created. This would however waste an enormous amount of disk space. CVS stores all the versions of a file in a single file in a clever way that only stores the differences between versions.

CVS also helps you if you are part of a group of people working on the same project. It is all too easy to overwrite each others' changes unless you are extremely careful. Some editors, like GNU Emacs, try to make sure that the same file is never modified by two people at the same time. Unfortunately, if someone is using another editor, that safeguard will not work. CVS solves this problem by insulating the different developers from each other. Every developer works in his own directory, and CVS merges the work when each developer is done.

CVS started out as a bunch of shell scripts written by Dick Grune, posted to the newsgroup `comp.sources.unix` in the volume 6 release of December, 1986. While no actual code from these shell scripts is present in the current version of CVS much of the CVS conflict resolution algorithms come from them.

In April, 1989, Brian Berliner designed and coded CVS. Jeff Polk later helped Brian with the design of the CVS module and vendor branch support.

You can get CVS via anonymous FTP from a number of sites; for example see

<http://www.gnu.ai.mit.edu/order/ftp.html>

for a list of the GNU FTP sites.

There is a mailing list, known as `info-cvs`, devoted to CVS. To subscribe or unsubscribe send a message to `info-cvs-request@prep.ai.mit.edu`. Please be specific about your email address. As of May 1996, subscription requests are handled by a busy human being, so you cannot expect to be added or removed immediately. If you prefer a usenet group, the right group is `comp.software.config-mgmt` which is for CVS discussions (along with other configuration management systems). In the future, it might be possible to create a `comp.software.config-mgmt.cvs`, but probably only if there is sufficient CVS traffic on `comp.software.config-mgmt`.

You can also subscribe to the bug-cvs mailing list, described in more detail in [BUGS], page 2. To subscribe send mail to `bug-cvs-request@prep.ai.mit.edu`.

CVS is not...

CVS can do a lot of things for you, but it does not try to be everything for everyone.

CVS is not a build system.

Though the structure of your repository and modules file interact with your build system (e.g. 'Makefile's), they are essentially independent.

CVS does not dictate how you build anything. It merely stores files for retrieval in a tree structure you devise.

CVS does not dictate how to use disk space in the checked out working directories. If you write your ‘Makefile’s or scripts in every directory so they have to know the relative positions of everything else, you wind up requiring the entire repository to be checked out.

If you modularize your work, and construct a build system that will share files (via links, mounts, `VPATH` in ‘Makefile’s, etc.), you can arrange your disk usage however you like.

But you have to remember that *any* such system is a lot of work to construct and maintain. CVS does not address the issues involved.

Of course, you should place the tools created to support such a build system (scripts, ‘Makefile’s, etc) under CVS.

Figuring out what files need to be rebuilt when something changes is, again, something to be handled outside the scope of CVS. One traditional approach is to use `make` for building, and use some automated tool for generating the dependencies which `make` uses.

See Chapter 18 [Builds], page 79, for more information on doing builds in conjunction with CVS.

CVS is not a substitute for management.

Your managers and project leaders are expected to talk to you frequently enough to make certain you are aware of schedules, merge points, branch names and release dates. If they don’t, CVS can’t help.

CVS is an instrument for making sources dance to your tune. But you are the piper and the composer. No instrument plays itself or writes its own music.

CVS is not a substitute for developer communication.

When faced with conflicts within a single file, most developers manage to resolve them without too much effort. But a more general definition of “conflict” includes problems too difficult to solve without communication between developers.

CVS cannot determine when simultaneous changes within a single file, or across a whole collection of files, will logically conflict with one another. Its concept of a *conflict* is purely textual, arising when two changes to the same base file are near enough to spook the merge (i.e. `diff3`) command.

CVS does not claim to help at all in figuring out non-textual or distributed conflicts in program logic.

For example: Say you change the arguments to function `X` defined in file ‘A’. At the same time, someone edits file ‘B’, adding new calls to function `X` using the old arguments. You are outside the realm of CVS’s competence.

Acquire the habit of reading specs and talking to your peers.

CVS does not have change control

Change control refers to a number of things. First of all it can mean *bug-tracking*, that is being able to keep a database of reported bugs and the sta-

tus of each one (is it fixed? in what release? has the bug submitter agreed that it is fixed?). For interfacing CVS to an external bug-tracking system, see the `'rcsinfo'` and `'verifymsg'` files (see Appendix C [Administrative files], page 125).

Another aspect of change control is keeping track of the fact that changes to several files were in fact changed together as one logical change. If you check in several files in a single `cvs commit` operation, CVS then forgets that those files were checked in together, and the fact that they have the same log message is the only thing tying them together. Keeping a GNU style `'ChangeLog'` can help somewhat.

Another aspect of change control, in some systems, is the ability to keep track of the status of each change. Some changes have been written by a developer, others have been reviewed by a second developer, and so on. Generally, the way to do this with CVS is to generate a diff (using `cvs diff` or `diff`) and email it to someone who can then apply it using the `patch` utility. This is very flexible, but depends on mechanisms outside CVS to make sure nothing falls through the cracks.

CVS is not an automated testing program

It should be possible to enforce mandatory use of a testsuite using the `commitinfo` file. I haven't heard a lot about projects trying to do that or whether there are subtle gotchas, however.

CVS does not have a builtin process model

Some systems provide ways to ensure that changes or releases go through various steps, with various approvals as needed. Generally, one can accomplish this with CVS but it might be a little more work. In some cases you'll want to use the `'commitinfo'`, `'loginfo'`, `'rcsinfo'`, or `'verifymsg'` files, to require that certain steps be performed before CVS will allow a checkin. Also consider whether features such as branches and tags can be used to perform tasks such as doing work in a development tree and then merging certain changes over to a stable tree only once they have been proven.

2 A sample session

As a way of introducing CVS, we'll go through a typical work-session using CVS. The first thing to understand is that CVS stores all files in a centralized *repository* (see Chapter 3 [Repository], page 13); this section assumes that a repository is set up.

Suppose you are working on a simple compiler. The source consists of a handful of C files and a `Makefile`. The compiler is called `tc` (Trivial Compiler), and the repository is set up so that there is a module called `tc`.

2.1 Getting the source

The first thing you must do is to get your own working copy of the source for `tc`. For this, you use the `checkout` command:

```
$ cvs checkout tc
```

This will create a new directory called `tc` and populate it with the source files.

```
$ cd tc
$ ls
CVS          Makefile    backend.c   driver.c    frontend.c  parser.c
```

The `CVS` directory is used internally by CVS. Normally, you should not modify or remove any of the files in it.

You start your favorite editor, hack away at `backend.c`, and a couple of hours later you have added an optimization pass to the compiler. A note to RCS and SCCS users: There is no need to lock the files that you want to edit. See Chapter 5 [Multiple developers], page 33, for an explanation.

2.2 Committing your changes

When you have checked that the compiler is still compilable you decide to make a new version of `backend.c`. This will store your new `backend.c` in the repository and make it available to anyone else who is using that same repository.

```
$ cvs commit backend.c
```

CVS starts an editor, to allow you to enter a log message. You type in “Added an optimization pass.”, save the temporary file, and exit the editor.

The environment variable `CVSEEDITOR` determines which editor is started. If `CVSEEDITOR` is not set, then if the environment variable `EDITOR` is set, it will be used. If both `CVSEEDITOR` and `EDITOR` are not set then there is a default which will vary with your operating system, for example `vi` for unix or `notepad` for Windows NT/95.

When CVS starts the editor, it includes a list of files which are modified. For the CVS client, this list is based on comparing the modification time of the file against the modification time that the file had when it was last gotten or updated. Therefore, if a file's modification time has changed but its contents have not, it will show up as modified. The simplest way to handle this is simply not to worry about it—if you proceed with the commit CVS will detect that the contents are not modified and treat it as an unmodified file. The

next `update` will clue CVS in to the fact that the file is unmodified, and it will reset its stored timestamp so that the file will not show up in future editor sessions.

If you want to avoid starting an editor you can specify the log message on the command line using the `-m` flag instead, like this:

```
$ cvs commit -m "Added an optimization pass" backend.c
```

2.3 Cleaning up

Before you turn to other tasks you decide to remove your working copy of `tc`. One acceptable way to do that is of course

```
$ cd ..
$ rm -r tc
```

but a better way is to use the `release` command (see Section A.15 [release], page 108):

```
$ cd ..
$ cvs release -d tc
M driver.c
? tc
You have [1] altered files in this repository.
Are you sure you want to release (and delete) module 'tc': n
** 'release' aborted by user choice.
```

The `release` command checks that all your modifications have been committed. If history logging is enabled it also makes a note in the history file. See Section C.10 [history file], page 134.

When you use the `-d` flag with `release`, it also removes your working copy.

In the example above, the `release` command wrote a couple of lines of output. `? tc` means that the file `tc` is unknown to CVS. That is nothing to worry about: `tc` is the executable compiler, and it should not be stored in the repository. See Section C.9 [cvsignore], page 133, for information about how to make that warning go away. See Section A.15.2 [release output], page 108, for a complete explanation of all possible output from `release`.

`M driver.c` is more serious. It means that the file `driver.c` has been modified since it was checked out.

The `release` command always finishes by telling you how many modified files you have in your working copy of the sources, and then asks you for confirmation before deleting any files or making any note in the history file.

You decide to play it safe and answer `n` (RET) when `release` asks for confirmation.

2.4 Viewing differences

You do not remember modifying `driver.c`, so you want to see what has happened to that file.

```
$ cd tc
$ cvs diff driver.c
```

This command runs `diff` to compare the version of `driver.c` that you checked out with your working copy. When you see the output you remember that you added a command line option that enabled the optimization pass. You check it in, and release the module.

```
$ cvs commit -m "Added an optimization pass" driver.c
Checking in driver.c;
/usr/local/cvsroot/tc/driver.c,v <-- driver.c
new revision: 1.2; previous revision: 1.1
done
$ cd ..
$ cvs release -d tc
? tc
You have [0] altered files in this repository.
Are you sure you want to release (and delete) module 'tc': y
```


3 The Repository

The CVS *repository* stores a complete copy of all the files and directories which are under version control.

Normally, you never access any of the files in the repository directly. Instead, you use CVS commands to get your own copy of the files into a *working directory*, and then work on that copy. When you've finished a set of changes, you check (or *commit*) them back into the repository. The repository then contains the changes which you have made, as well as recording exactly what you changed, when you changed it, and other such information. Note that the repository is not a subdirectory of the working directory, or vice versa; they should be in separate locations.

CVS can access a repository by a variety of means. It might be on the local computer, or it might be on a computer across the room or across the world. To distinguish various ways to access a repository, the repository name can start with an *access method*. For example, the access method `:local:` means to access a repository directory, so the repository `:local:/usr/local/cvsroot` means that the repository is in `/usr/local/cvsroot` on the computer running CVS. For information on other access methods, see Section 3.9 [Remote repositories], page 21.

If the access method is omitted, then if the repository does not contain `:`, then `:local:` is assumed. If it does contain `:` than either `:ext:` or `:server:` is assumed. For example, if you have a local repository in `/usr/local/cvsroot`, you can use `/usr/local/cvsroot` instead of `:local:/usr/local/cvsroot`. But if (under Windows NT, for example) your local repository is `c:\src\cvsroot`, then you must specify the access method, as in `:local:c:\src\cvsroot`.

The repository is split in two parts. `CVSROOT` contains administrative files for CVS. The other directories contain the actual user-defined modules.

3.1 Telling CVS where your repository is

There are a couple of different ways to tell CVS where to find the repository. You can name the repository on the command line explicitly, with the `-d` (for "directory") option:

```
cvs -d /usr/local/cvsroot checkout yoyodyne/tc
```

Or you can set the `CVSROOT` environment variable to an absolute path to the root of the repository, `/usr/local/cvsroot` in this example. To set `CVSROOT`, all `csh` and `tcsh` users should have this line in their `.cshrc` or `.tcshrc` files:

```
setenv CVSROOT /usr/local/cvsroot
```

`sh` and `bash` users should instead have these lines in their `.profile` or `.bashrc`:

```
CVSROOT=/usr/local/cvsroot
export CVSROOT
```

A repository specified with `-d` will override the `CVSROOT` environment variable. Once you've checked a working copy out from the repository, it will remember where its repository is (the information is recorded in the `CVS/Root` file in the working copy).

The `-d` option and the `CVS/Root` file both override the `CVSROOT` environment variable. If `-d` option differs from `CVS/Root`, the former is used (and specifying `-d` will cause

‘CVS/Root’ to be updated). Of course, for proper operation they should be two ways of referring to the same repository.

3.2 How data is stored in the repository

For most purposes it isn’t important *how* CVS stores information in the repository. In fact, the format has changed in the past, and is likely to change in the future. Since in almost all cases one accesses the repository via CVS commands; such changes need not be disruptive.

However, in some cases it may be necessary to understand how CVS stores data in the repository, for example you might need to track down CVS locks (see Section 5.5 [Concurrency], page 37) or you might need to deal with the file permissions appropriate for the repository.

3.2.1 Where files are stored within the repository

The overall structure of the repository is a directory tree corresponding to the directories in the working directory. For example, supposing the repository is in

```
/usr/local/cvsroot
```

here is a possible directory tree (showing only the directories):

```
/usr
|
+--local
|  |
|  +--cvsroot
|  |   |
|  |   +--CVSROOT
|  |       | (administrative files)
|  |       |
|  |       +--gnu
|  |           |
|  |           +--diff
|  |               | (source code to GNU diff)
|  |               |
|  |               +--rcs
|  |                   | (source code to RCS)
|  |                   |
|  |                   +--cvs
|  |                       | (source code to CVS)
|  |                       |
|  |                       +--yoyodyne
|  |                           |
|  |                           +--tc
|  |                               |
|  |                               +--man
|  |                               |
|  |                               +--testing
```



```

|
+--(other Yoyodyne software)

```

With the directories are *history files* for each file under version control. The name of the history file is the name of the corresponding file with ‘,v’ appended to the end. Here is what the repository for the ‘yoyodyne/tc’ directory might look like:

```

$CVSRROOT
|
+--yoyodyne
| |
|  +--tc
|  | |
|  |  +--Makefile,v
|  |  +--backend.c,v
|  |  +--driver.c,v
|  |  +--frontend.c,v
|  |  +--parser.c,v
|  |  +--man
|  |  |
|  |  +--tc.1,v
|  |  |
|  |  +--testing
|  |  |
|  |  +--testpgm.t,v
|  |  +--test2.t,v

```

The history files contain, among other things, enough information to recreate any revision of the file, a log of all commit messages and the user-name of the person who committed the revision. The history files are known as *RCS files*, because the first program to store files in that format was a version control system known as RCS. For a full description of the file format, see the `man` page `rcsfile(5)`, distributed with RCS. This file format has become very common—many systems other than CVS or RCS can at least import history files in this format.

The RCS files used in CVS differ in a few ways from the standard format. The biggest difference is magic branches; for more information see Section 6.8 [Magic branch numbers], page 50. Also in CVS the valid tag names are a subset of what RCS accepts; for CVS’s rules see Section 6.4 [Tags], page 45.

3.2.2 File permissions

All ‘,v’ files are created read-only, and you should not change the permission of those files. The directories inside the repository should be writable by the persons that have permission to modify the files in each directory. This normally means that you must create a UNIX group (see `group(5)`) consisting of the persons that are to edit the files in a project, and set up the repository so that it is that group that owns the directory.

This means that you can only control access to files on a per-directory basis.

Note that users must also have write access to check out files, because CVS needs to create lock files (see Section 5.5 [Concurrency], page 37).

Also note that users must have write access to the ‘CVSROOT/val-tags’ file. CVS uses it to keep track of what tags are valid tag names (it is sometimes updated when tags are used, as well as when they are created, though).

CVS tries to set up reasonable file permissions for new directories that are added inside the tree, but you must fix the permissions manually when a new directory should have different permissions than its parent directory. If you set the `CVSUMASK` environment variable that will control the file permissions which CVS uses in creating directories and/or files in the repository. `CVSUMASK` does not affect the file permissions in the working directory; such files have the permissions which are typical for newly created files, except that sometimes CVS creates them read-only (see the sections on watches, Section 5.6.1 [Setting a watch], page 38; `-r`, Section A.4 [Global options], page 86; or `CVSREAD`, Appendix D [Environment variables], page 137).

Note that using the client/server CVS (see Section 3.9 [Remote repositories], page 21), there is no good way to set `CVSUMASK`; the setting on the client machine has no effect. If you are connecting with `rsh`, you can set `CVSUMASK` in ‘`.bashrc`’ or ‘`.cshrc`’, as described in the documentation for your operating system. This behavior might change in future versions of CVS; do not rely on the setting of `CVSUMASK` on the client having no effect.

Since CVS was not written to be run `setuid`, it is unsafe to try to run it `setuid`. You cannot use the `setuid` features of RCS together with CVS.

3.2.3 The attic

You will notice that sometimes CVS stores an RCS file in the `Attic`. For example, if the `CVSROOT` is ‘`/usr/local/cvsroot`’ and we are talking about the file ‘`backend.c`’ in the directory ‘`yoyodyne/tc`’, then the file normally would be in

```
/usr/local/cvsroot/yoyodyne/tc/backend.c,v
```

but if it goes in the attic, it would be in

```
/usr/local/cvsroot/yoyodyne/tc/Attic/backend.c,v
```

instead. It should not matter from a user point of view whether a file is in the attic; CVS keeps track of this and looks in the attic when it needs to. But in case you want to know, the rule is that the RCS file is stored in the attic if and only if the head revision on the trunk has state `dead`. A `dead` state means that file has been removed, or never added, for that revision. For example, if you add a file on a branch, it will have a trunk revision in `dead` state, and a branch revision in a non-`dead` state.

3.3 How data is stored in the working directory

While we are discussing CVS internals which may become visible from time to time, we might as well talk about what CVS puts in the ‘CVS’ directories in the working directories. As with the repository, CVS handles this information and one can usually access it via CVS commands. But in some cases it may be useful to look at it, and other programs, such as the `jCVS` graphical user interface or the `VC` package for `emacs`, may need to look at it. Such programs should follow the recommendations in this section if they hope to be able to work with other programs which use those files, including future versions of the programs just mentioned and the command-line CVS client.

The ‘CVS’ directory contains several files. Programs which are reading this directory should silently ignore files which are in the directory but which are not documented here, to allow for future expansion.

‘Root’ This file contains the current CVS root, as described in Section 3.1 [Specifying a repository], page 13.

‘Repository’

This file contains the directory within the repository which the current directory corresponds with. For historical reasons it is an absolute pathname, although it would make more sense for it to be relative to the root. For example, after the command

```
cvcs -d :local:/usr/local/cvsroot checkout yoyodyne/tc
```

‘Root’ will contain

```
:local:/usr/local/cvsroot
```

and ‘Repository’ will contain

```
/usr/local/cvsroot/yoyodyne/tc
```

‘Entries’ This file lists the files and directories in the working directory. It is a text file according to the conventions appropriate for the operating system in question. The first character of each line indicates what sort of line it is. If the character is unrecognized, programs reading the file should silently skip that line, to allow for future expansion.

If the first character is ‘/’, then the format is:

```
/name/revision/timestamp[+conflict]/options/tagdate
```

where ‘[’ and ‘]’ are not part of the entry, but instead indicate that the ‘+’ and conflict marker are optional. *name* is the name of the file within the directory. *revision* is the revision that the file in the working derives from, or ‘0’ for an added file, or ‘-’ followed by a revision for a removed file. *timestamp* is the timestamp of the file at the time that CVS created it; if the timestamp differs with the actual modification time of the file it means the file has been modified. It is in Universal Time (UT), stored in the format used by the ISO C `asctime()` function (for example, ‘Sun Apr 7 01:29:26 1996’). One may write a string which is not in that format, for example, ‘Result of merge’, to indicate that the file should always be considered to be modified. This is not a special case; to see whether a file is modified a program should take the timestamp of the file and simply do a string compare with *timestamp*. *conflict* indicates that there was a conflict; if it is the same as the actual modification time of the file it means that the user has obviously not resolved the conflict. *options* contains sticky options (for example ‘-kb’ for a binary file). *tagdate* contains ‘T’ followed by a tag name, or ‘D’ for a date, followed by a sticky tag or date. Note that if *timestamp* contains a pair of timestamps separated by a space, rather than a single timestamp, you are dealing with a version of CVS earlier than CVS 1.5 (not documented here).

If the first character of a line in ‘Entries’ is ‘D’, then it indicates a subdirectory. ‘D’ on a line all by itself indicates that the program which wrote the ‘Entries’

file does record subdirectories (therefore, if there is such a line and no other lines beginning with ‘D’, one knows there are no subdirectories). Otherwise, the line looks like:

```
D/name/filler1/filler2/filler3/filler4
```

where *name* is the name of the subdirectory, and all the *filler* fields should be silently ignored, for future expansion. Programs which modify `Entries` files should preserve these fields.

‘`Entries.Log`’

This file does not record any information beyond that in ‘`Entries`’, but it does provide a way to update the information without having to rewrite the entire ‘`Entries`’ file, including the ability to preserve the information even if the program writing ‘`Entries`’ and ‘`Entries.Log`’ abruptly aborts. Programs which are reading the ‘`Entries`’ file should also check for ‘`Entries.Log`’. If the latter exists, they should read ‘`Entries`’ and then apply the changes mentioned in ‘`Entries.Log`’. After applying the changes, the recommended practice is to rewrite ‘`Entries`’ and then delete ‘`Entries.Log`’. The format of a line in ‘`Entries.Log`’ is a single character command followed by a space followed by a line in the format specified for a line in ‘`Entries`’. The single character command is ‘A’ to indicate that the entry is being added, ‘R’ to indicate that the entry is being removed, or any other character to indicate that the entire line in ‘`Entries.Log`’ should be silently ignored (for future expansion). If the second character of the line in ‘`Entries.Log`’ is not a space, then it was written by an older version of CVS (not documented here).

‘`Entries.Backup`’

This is a temporary file. Recommended usage is to write a new entries file to ‘`Entries.Backup`’, and then to rename it (atomically, where possible) to ‘`Entries`’.

‘`Entries.Static`’

The only relevant thing about this file is whether it exists or not. If it exists, then it means that only part of a directory was gotten and CVS will not create additional files in that directory. To clear it, use the `update` command with the ‘-d’ option, which will get the additional files and remove ‘`Entries.Static`’.

‘`Tag`’

This file contains per-directory sticky tags or dates. The first character is ‘T’ for a branch tag, ‘N’ for a non-branch tag, or ‘D’ for a date, or another character to mean the file should be silently ignored, for future expansion. This character is followed by the tag or date. Note that per-directory sticky tags or dates are used for things like applying to files which are newly added; they might not be the same as the sticky tags or dates on individual files. For general information on sticky tags and dates, see Section 6.7 [Sticky tags], page 48.

‘`Checkin.prog`’

‘`Update.prog`’

These files store the programs specified by the ‘-i’ and ‘-u’ options in the modules file, respectively.

- 'Notify'** This file stores notifications (for example, for `edit` or `unedit`) which have not yet been sent to the server. Its format is not yet documented here.
- 'Notify.tmp'** This file is to **'Notify'** as **'Entries.Backup'** is to **'Entries'**. That is, to write **'Notify'**, first write the new contents to **'Notify.tmp'** and then (atomically where possible), rename it to **'Notify'**.
- 'Base'** If watches are in use, then an `edit` command stores the original copy of the file in the **'Base'** directory. This allows the `unedit` command to operate even if it is unable to communicate with the server.
- 'Template'** This file contains the template specified by the **'rcsinfo'** file (see Section C.8 [rcsinfo], page 133). It is only used by the client; the non-client/server CVS consults **'rcsinfo'** directly.

3.4 The administrative files

The directory **'\$CVSROOT/CVSROOT'** contains some *administrative files*. See Appendix C [Administrative files], page 125, for a complete description. You can use CVS without any of these files, but some commands work better when at least the **'modules'** file is properly set up.

The most important of these files is the **'modules'** file. It defines all modules in the repository. This is a sample **'modules'** file.

```
CVSROOT      CVSROOT
modules      CVSROOT modules
cvs          gnu/cvs
rcs          gnu/rcs
diff         gnu/diff
tc          yoyodyne/tc
```

The **'modules'** file is line oriented. In its simplest form each line contains the name of the module, whitespace, and the directory where the module resides. The directory is a path relative to **'\$CVSROOT'**. The last four lines in the example above are examples of such lines.

The line that defines the module called **'modules'** uses features that are not explained here. See Section C.1 [modules], page 125, for a full explanation of all the available features.

3.4.1 Editing administrative files

You edit the administrative files in the same way that you would edit any other module. Use `cvs checkout CVSROOT` to get a working copy, edit it, and commit your changes in the normal way.

It is possible to commit an erroneous administrative file. You can often fix the error and check in a new revision, but sometimes a particularly bad error in the administrative file makes it impossible to commit new revisions.

3.5 Multiple repositories

In some situations it is a good idea to have more than one repository, for instance if you have two development groups that work on separate projects without sharing any code. All you have to do to have several repositories is to specify the appropriate repository, using the `CVSROOT` environment variable, the `-d` option to `CVS`, or (once you have checked out a working directory) by simply allowing `CVS` to use the repository that was used to check out the working directory (see Section 3.1 [Specifying a repository], page 13).

The big advantage of having multiple repositories is that they can reside on different servers. The big disadvantage is that you cannot have a single `CVS` command recurse into directories which comes from different repositories. Generally speaking, if you are thinking of setting up several repositories on the same machine, you might want to consider using several directories within the same repository.

None of the examples in this manual show multiple repositories.

3.6 Creating a repository

To set up a `CVS` repository, first choose the machine and disk on which you want to store the revision history of the source files. CPU and memory requirements are modest—a server with 32M of memory or even less can handle a fairly large source tree with a fair amount of activity. To estimate disk space requirements, if you are importing RCS files from another system, the size of those files is the approximate initial size of your repository, or if you are starting without any version history, a rule of thumb is to allow for the server approximately three times the size of the code to be under `CVS` for the repository (you will eventually outgrow this, but not for a while). On the machines on which the developers will be working, you'll want disk space for approximately one working directory for each developer (either the entire tree or a portion of it, depending on what each developer uses). Don't worry about CPU and memory requirements for the clients—any machine with enough capacity to run the operating system in question should have little trouble.

The repository should be accessible (directly or via a networked file system) from all machines which want to use `CVS` in server or local mode; the client machines need not have any access to it other than via the `CVS` protocol. It is not possible to use `CVS` to read from a repository which one only has read access to; `CVS` needs to be able to create lock files (see Section 5.5 [Concurrency], page 37).

To create a repository, run the `cvs init` command. It will set up an empty repository in the `CVS` root specified in the usual way (see Chapter 3 [Repository], page 13). For example,

```
cvs -d /usr/local/cvsroot init
```

`cvs init` is careful to never overwrite any existing files in the repository, so no harm is done if you run `cvs init` on an already set-up repository.

`cvs init` will enable history logging; if you don't want that, remove the history file after running `cvs init`. See Section C.10 [history file], page 134.

3.7 Backing up a repository

There is nothing particularly magical about the files in the repository; for the most part it is possible to back them up just like any other files. However, there are a few issues to consider.

The first is that to be paranoid, one should either not use CVS during the backup, or have the backup program lock CVS while doing the backup. To not use CVS, you might forbid logins to machines which can access the repository, turn off your CVS server, or similar mechanisms. The details would depend on your operating system and how you have CVS set up. To lock CVS, you would create `#cvs.rfl` locks in each repository directory. See Section 5.5 [Concurrency], page 37, for more on CVS locks. Having said all this, if you just back up without any of these precautions, the results are unlikely to be particularly dire. Restoring from backup, the repository might be in an inconsistent state, but this would not be particularly hard to fix manually.

When you restore a repository from backup, assuming that changes in the repository were made after the time of the backup, working directories which were not affected by the failure may refer to revisions which no longer exist in the repository. Trying to run CVS in such directories will typically produce an error message. One way to get those changes back into the repository is as follows:

- Get a new working directory.
- Copy the files from the working directory from before the failure over to the new working directory (do not copy the contents of the ‘CVS’ directories, of course).
- Working in the new working directory, use commands such as `cvs update` and `cvs diff` to figure out what has changed, and then when you are ready, commit the changes into the repository.

3.8 Moving a repository

Just as backing up the files in the repository is pretty much like backing up any other files, if you need to move a repository from one place to another it is also pretty much like just moving any other collection of files.

The main thing to consider is that working directories point to the repository. The simplest way to deal with a moved repository is to just get a fresh working directory after the move. Of course, you’ll want to make sure that the old working directory had been checked in before the move, or you figured out some other way to make sure that you don’t lose any changes. If you really do want to reuse the existing working directory, it should be possible with manual surgery on the ‘CVS/Repository’ files. You can see Section 3.3 [Working directory storage], page 16, for information on the ‘CVS/Repository’ and ‘CVS/Root’ files, but unless you are sure you want to bother, it probably isn’t worth it.

3.9 Remote repositories

Your working copy of the sources can be on a different machine than the repository. Using CVS in this manner is known as *client/server* operation. You run CVS on a machine which can mount your working directory, known as the *client*, and tell it to communicate to

a machine which can mount the repository, known as the *server*. Generally, using a remote repository is just like using a local one, except that the format of the repository name is:

```
:method:user@hostname:/path/to/repository
```

The details of exactly what needs to be set up depend on how you are connecting to the server.

If *method* is not specified, and the repository name contains ':', then the default is **ext** or **server**, depending on your platform; both are described in Section 3.9.2 [Connecting via rsh], page 22.

3.9.1 Server requirements

The quick answer to what sort of machine is suitable as a server is that requirements are modest—a server with 32M of memory or even less can handle a fairly large source tree with a fair amount of activity.

The real answer, of course, is more complicated. The CVS server consists of two processes for each client that it is serving. Memory consumption on the child process should remain fairly small. Memory consumption on the parent process, particularly if the network connection to the client is slow, can be expected to grow to slightly more than the size of the sources in a single directory, or two megabytes, whichever is larger.

Multiplying the size of each CVS server by the number of servers which you expect to have active at one time should give an idea of memory requirements for the server. For the most part, the memory consumed by the parent process probably can be swap space rather than physical memory.

Resource consumption for the client or the non-client/server CVS is even more modest—any machine with enough capacity to run the operating system in question should have little trouble.

3.9.2 Connecting with rsh

CVS uses the 'rsh' protocol to perform these operations, so the remote user host needs to have a '.rhosts' file which grants access to the local user.

For example, suppose you are the user 'mozart' on the local machine 'toe.grunge.com', and the server machine is 'chainsaw.yard.com'. On chainsaw, put the following line into the file '.rhosts' in 'bach's home directory:

```
toe.grunge.com  mozart
```

Then test that rsh is working with

```
rsh -l bach chainsaw.yard.com 'echo $PATH'
```

Next you have to make sure that rsh will be able to find the server. Make sure that the path which rsh printed in the above example includes the directory containing a program named cvs which is the server. You need to set the path in '.bashrc', '.cshrc', etc., not '.login' or '.profile'. Alternately, you can set the environment variable CVS_SERVER on the client machine to the filename of the server you want to use, for example '/usr/local/bin/cvs-1.6'.

There is no need to edit `inetd.conf` or start a CVS server daemon.

There are two access methods that you use in CVSROOT for rsh. `:server:` specifies an internal rsh client, which is supported only by some CVS ports. `:ext:` specifies an external rsh program. By default this is `rsh` but you may set the `CVS_RSH` environment variable to invoke another program which can access the remote server (for example, `remsh` on HP-UX 9 because `rsh` is something different). It must be a program which can transmit data to and from the server without modifying it; for example the Windows NT `rsh` is not suitable since it by default translates between CRLF and LF. The OS/2 CVS port has a hack to pass `-b` to `rsh` to get around this, but since this could potentially cause problems for programs other than the standard `rsh`, it may change in the future. If you set `CVS_RSH` to `SSH` or some other rsh replacement, the instructions in the rest of this section concerning `.rhosts` and so on are likely to be inapplicable; consult the documentation for your rsh replacement.

Continuing our example, supposing you want to access the module `'foo'` in the repository `'/usr/local/cvsroot/'`, on machine `'chainsaw.yard.com'`, you are ready to go:

```
cvs -d :ext:bach@chainsaw.yard.com:/usr/local/cvsroot checkout foo
```

(The `'bach@'` can be omitted if the username is the same on both the local and remote hosts.)

3.9.3 Direct connection with password authentication

The cvs client can also connect to the server using a password protocol. This is particularly useful if using `rsh` is not feasible (for example, the server is behind a firewall), and Kerberos also is not available.

To use this method, it is necessary to make some adjustments on both the server and client sides.

3.9.3.1 Setting up the server for password authentication

On the server side, the file `'/etc/inetd.conf'` needs to be edited so `inetd` knows to run the command `cvs pserver` when it receives a connection on the right port. By default, the port number is 2401; it would be different if your client were compiled with `CVS_AUTH_PORT` defined to something else, though.

If your `inetd` allows raw port numbers in `'/etc/inetd.conf'`, then the following (all on a single line in `'inetd.conf'`) should be sufficient:

```
2401 stream tcp nowait root /usr/local/bin/cvs
cvs -b /usr/local/bin --allow-root=/usr/cvsroot pserver
```

The `'-b'` option specifies the directory which contains the RCS binaries on the server. You could also use the `'-T'` option to specify a temporary directory.

The `'--allow-root'` option specifies the allowable CVSROOT directory. Clients which attempt to use a different CVSROOT directory will not be allowed to connect. If there is more than one CVSROOT directory which you want to allow, repeat the option.

If your `inetd` wants a symbolic service name instead of a raw port number, then put this in `'/etc/services'`:

```
cvspserver      2401/tcp
```

and put `cvspserver` instead of 2401 in `inetd.conf`.

Once the above is taken care of, restart your `inetd`, or do whatever is necessary to force it to reread its initialization files.

Because the client stores and transmits passwords in cleartext (almost—see Section 3.9.3.3 [Password authentication security], page 25, for details), a separate CVS password file may be used, so people don't compromise their regular passwords when they access the repository. This file is `$CVSROOT/CVSROOT/passwd` (see Section 3.4 [Intro administrative files], page 19). Its format is similar to `/etc/passwd`, except that it only has two fields, username and password. For example:

```
bach:ULtgRLXo7NRxs
cwang:1s0p854gDF3DY
```

The password is encrypted according to the standard Unix `crypt()` function, so it is possible to paste in passwords directly from regular Unix `'passwd'` files.

When authenticating a password, the server first checks for the user in the CVS `'passwd'` file. If it finds the user, it compares against that password. If it does not find the user, or if the CVS `'passwd'` file does not exist, then the server tries to match the password using the system's user-lookup routine. When using the CVS `'passwd'` file, the server runs under as the username specified in the the third argument in the entry, or as the first argument if there is no third argument (in this way CVS allows imaginary usernames provided the CVS `'passwd'` file indicates corresponding valid system usernames). In any case, CVS will have no privileges which the (valid) user would not have.

It is possible to “map” cvs-specific usernames onto system usernames (i.e., onto system login names) in the `$CVSROOT/CVSROOT/passwd` file by appending a colon and the system username after the password. For example:

```
cvs:ULtgRLXo7NRxs:kfogel
generic:1s0p854gDF3DY:spwang
anyone:1s0p854gDF3DY:spwang
```

Thus, someone remotely accessing the repository on `'chainsaw.yard.com'` with the following command:

```
cvs -d :pserver:cvs@chainsaw.yard.com:/usr/local/cvsroot checkout foo
```

would end up running the server under the system identity `kfogel`, assuming successful authentication. However, the remote user would not necessarily need to know `kfogel`'s system password, as the `$CVSROOT/CVSROOT/passwd` file might contain a different password, used only for CVS. And as the example above indicates, it is permissible to map multiple cvs usernames onto a single system username.

This feature is designed to allow people repository access without full system access (in particular, see See Section 3.10 [Read-only access], page 26); however, also See Section 3.9.3.3 [Password authentication security], page 25. Any sort of repository access very likely implies a degree of general system access as well.

Right now, the only way to put a password in the CVS `'passwd'` file is to paste it there from somewhere else. Someday, there may be a `cvs passwd` command.

3.9.3.2 Using the client with password authentication

Before connecting to the server, the client must *log in* with the command `cvs login`. Logging in verifies a password with the server, and also records the password for later transactions with the server. The `cvs login` command needs to know the username, server hostname, and full repository path, and it gets this information from the repository argument or the `CVSROOT` environment variable.

`cvs login` is interactive — it prompts for a password:

```
cvs -d :pserver:bach@chainsaw.yard.com:/usr/local/cvsroot login
CVS password:
```

The password is checked with the server; if it is correct, the `login` succeeds, else it fails, complaining that the password was incorrect.

Once you have logged in, you can force `CVS` to connect directly to the server and authenticate with the stored password:

```
cvs -d :pserver:bach@chainsaw.yard.com:/usr/local/cvsroot checkout foo
```

The `:pserver:` is necessary because without it, `CVS` will assume it should use `rsh` to connect with the server (see Section 3.9.2 [Connecting via `rsh`], page 22). (Once you have a working copy checked out and are running `CVS` commands from within it, there is no longer any need to specify the repository explicitly, because `CVS` records it in the working copy's `'CVS'` subdirectory.)

Passwords are stored by default in the file `'$HOME/.cvspass'`. Its format is human-readable, but don't edit it unless you know what you are doing. The passwords are not stored in cleartext, but are trivially encoded to protect them from "innocent" compromise (i.e., inadvertently being seen by a system administrator who happens to look at that file).

The password for the currently chosen remote repository can be removed from the `CVS_PASSFILE` by using the `cvs logout` command.

The `CVS_PASSFILE` environment variable overrides this default. If you use this variable, make sure you set it *before* `cvs login` is run. If you were to set it after running `cvs login`, then later `CVS` commands would be unable to look up the password for transmission to the server.

3.9.3.3 Security considerations with password authentication

The passwords are stored on the client side in a trivial encoding of the cleartext, and transmitted in the same encoding. The encoding is done only to prevent inadvertent password compromises (i.e., a system administrator accidentally looking at the file), and will not prevent even a naive attacker from gaining the password.

The separate `CVS` password file (see Section 3.9.3.1 [Password authentication server], page 23) allows people to use a different password for repository access than for login access. On the other hand, once a user has access to the repository, she can execute programs on the server system through a variety of means. Thus, repository access implies fairly broad system access as well. It might be possible to modify `CVS` to prevent that, but no one has done so as of this writing. Furthermore, there may be other ways in which having access to `CVS` allows people to gain more general access to the system; no one has done a careful audit.

In summary, anyone who gets the password gets repository access, and some measure of general system access as well. The password is available to anyone who can sniff network packets or read a protected (i.e., user read-only) file. If you want real security, get Kerberos.

3.9.4 Direct connection with kerberos

The main disadvantage of using rsh is that all the data needs to pass through additional programs, so it may be slower. So if you have kerberos installed you can connect via a direct TCP connection, authenticating with kerberos.

To do this, CVS needs to be compiled with kerberos support; when configuring CVS it tries to detect whether kerberos is present or you can use the ‘`--with-krb4`’ flag to configure.

The data transmitted is *not* encrypted by default. Encryption support must be compiled into both the client and server; use the ‘`--enable-encryption`’ configure option to turn it on. You must then use the `-x` global option to request encryption.

You need to edit `inetd.conf` on the server machine to run `cvserver`. The client uses port 1999 by default; if you want to use another port specify it in the `CVS_CLIENT_PORT` environment variable on the client.

When you want to use CVS, get a ticket in the usual way (generally `kinit`); it must be a ticket which allows you to log into the server machine. Then you are ready to go:

```
cvserver -d :kserver:chainsaw.yard.com:/user/local/cvsroot checkout foo
```

Previous versions of CVS would fall back to a connection via rsh; this version will not do so.

3.10 Read-only repository access

It is possible to grant read-only repository access to people using the password-authenticated server (see Section 3.9.3 [Password authenticated], page 23). (The other access methods do not have explicit support for read-only users because those methods all assume login access to the repository machine anyway, and therefore the user can do whatever local file permissions allow her to do.)

A user who has read-only access can do only those CVS operations which do not modify the repository, except for certain “administrative” files (such as lock files and the history file). It may be desirable to use this feature in conjunction with user-aliasing (see Section 3.9.3.1 [Password authentication server], page 23). However, note that read-only access does not repeal the existing security considerations in See Section 3.9.3.3 [Password authentication security], page 25.

There are two ways to specify read-only access for a user: by inclusion, and by exclusion.

"Inclusion" means listing that user specifically in the ‘`$CVSROOT/CVSROOT/readers`’ file, which is simply a newline-separated list of users. Here is a sample ‘`readers`’ file:

```
melissa
splotnik
jrandom
```

(Don’t forget the newline after the last user.)

"Exclusion" means explicitly listing everyone who has *write* access—if the file

`$CVSROOT/CVSROOT/writers`

exists, then only those users listed in it have write access, and everyone else has read-only access (of course, even the read-only users still need to be listed in the cvs `'passwd'` file). The `'writers'` file has the same format as the `'readers'` file.

Note: if your cvs `'passwd'` file maps cvs users onto system users (see Section 3.9.3.1 [Password authentication server], page 23), make sure you deny or grant read-only access using the cvs usernames, not the system usernames. That is, the `'readers'` and `'writers'` files contain cvs usernames, which may or may not be the same as system usernames.

Here is a complete description of the server's behavior in deciding whether to grant read-only or read-write access:

If `'readers'` exists, and this user is listed in it, then she gets read-only access. Or if `'writers'` exists, and this user is NOT listed in it, then she also gets read-only access (this is true even if `'readers'` exists but she is not listed there). Otherwise, she gets full read-write access.

Of course there is a conflict if the user is listed in both files. This is resolved in the more conservative way, it being better to protect the repository too much than too little: such a user gets read-only access.

3.11 Temporary directories for the server

While running, the cvs server creates temporary directories. They are named

`cvs-servpid`

where *pid* is the process identification number of the server. They are located in the directory specified by the `'TMPDIR'` environment variable (see Appendix D [Environment variables], page 137), the `'-T'` global option (see Section A.4 [Global options], page 86), or failing that `'/tmp'`.

In most cases the server will remove the temporary directory when it is done, whether it finishes normally or abnormally. However, there are a few cases in which the server does not or cannot remove the temporary directory, for example:

- If the server aborts due to an internal server error, it may preserve the directory to aid in debugging
- If the server is killed in a way that it has no way of cleaning up (most notably, `'kill -KILL'` on unix).
- If the system shuts down without an orderly shutdown, which tells the server to clean up.

In cases such as this, you will need to manually remove the `'cvs-servpid'` directories. As long as there is no server running with process identification number *pid*, it is safe to do so.

4 Starting a project with CVS

Because renaming files and moving them between directories is somewhat inconvenient, the first thing you do when you start a new project should be to think through your file organization. It is not impossible to rename or move files, but it does increase the potential for confusion and CVS does have some quirks particularly in the area of renaming directories. See Chapter 13 [Moving files], page 67.

What to do next depends on the situation at hand.

4.1 Setting up the files

The first step is to create the files inside the repository. This can be done in a couple of different ways.

4.1.1 Creating a directory tree from a number of files

When you begin using CVS, you will probably already have several projects that can be put under CVS control. In these cases the easiest way is to use the `import` command. An example is probably the easiest way to explain how to use it. If the files you want to install in CVS reside in `wdir`, and you want them to appear in the repository as `CVSROOT/yoyodyne/rdir`, you can do this:

```
$ cd wdir
$ cvs import -m "Imported sources" yoyodyne/rdir yoyo start
```

Unless you supply a log message with the `-m` flag, CVS starts an editor and prompts for a message. The string `yoyo` is a *vendor tag*, and `start` is a *release tag*. They may fill no purpose in this context, but since CVS requires them they must be present. See Chapter 12 [Tracking sources], page 63, for more information about them.

You can now verify that it worked, and remove your original source directory.

```
$ cd ..
$ mv dir dir.orig
$ cvs checkout yoyodyne/dir          # Explanation below
$ diff -r dir.orig yoyodyne/dir
$ rm -r dir.orig
```

Erasing the original sources is a good idea, to make sure that you do not accidentally edit them in `dir`, bypassing CVS. Of course, it would be wise to make sure that you have a backup of the sources before you remove them.

The `checkout` command can either take a module name as argument (as it has done in all previous examples) or a path name relative to `CVSROOT`, as it did in the example above.

It is a good idea to check that the permissions CVS sets on the directories inside `CVSROOT` are reasonable, and that they belong to the proper groups. See Section 3.2.2 [File permissions], page 15.

If some of the files you want to import are binary, you may want to use the wrappers features to specify which files are binary and which are not. See Section C.2 [Wrappers], page 126.

4.1.2 Creating Files From Other Version Control Systems

If you have a project which you are maintaining with another version control system, such as RCS, you may wish to put the files from that project into CVS, and preserve the revision history of the files.

From RCS If you have been using RCS, find the RCS files—usually a file named ‘foo.c’ will have its RCS file in ‘RCS/foo.c,v’ (but it could be other places; consult the RCS documentation for details). Then create the appropriate directories in CVS if they do not already exist. Then copy the files into the appropriate directories in the CVS repository (the name in the repository must be the name of the source file with ‘,v’ added; the files go directly in the appropriate directory of the repository, not in an ‘RCS’ subdirectory). This is one of the few times when it is a good idea to access the CVS repository directly, rather than using CVS commands. Then you are ready to check out a new working directory.

The RCS file should not be locked when you move it into CVS; if it is, CVS will have trouble letting you operate on it.

From another version control system

Many version control systems have the ability to export RCS files in the standard format. If yours does, export the RCS files and then follow the above instructions.

From SCCS

There is a script in the ‘contrib’ directory of the CVS source distribution called ‘sccs2rcs’ which converts SCCS files to RCS files. Note: you must run it on a machine which has both SCCS and RCS installed, and like everything else in contrib it is unsupported (your mileage may vary).

4.1.3 Creating a directory tree from scratch

For a new project, the easiest thing to do is probably to create an empty directory structure, like this:

```
$ mkdir tc
$ mkdir tc/man
$ mkdir tc/testing
```

After that, you use the `import` command to create the corresponding (empty) directory structure inside the repository:

```
$ cd tc
$ cvs import -m "Created directory structure" yoyodyne/dir yoyo start
```

Then, use `add` to add files (and new directories) as they appear.

Check that the permissions CVS sets on the directories inside ‘\$CVSROOT’ are reasonable.

4.2 Defining the module

The next step is to define the module in the ‘modules’ file. This is not strictly necessary, but modules can be convenient in grouping together related files and directories.

In simple cases these steps are sufficient to define a module.

1. Get a working copy of the modules file.

```
$ cvs checkout CVSROOT/modules
$ cd CVSROOT
```

2. Edit the file and insert a line that defines the module. See Section 3.4 [Intro administrative files], page 19, for an introduction. See Section C.1 [modules], page 125, for a full description of the modules file. You can use the following line to define the module 'tc':

```
tc    yoyodyne/tc
```

3. Commit your changes to the modules file.

```
$ cvs commit -m "Added the tc module." modules
```

4. Release the modules module.

```
$ cd ..
$ cvs release -d CVSROOT
```


5 Multiple developers

When more than one person works on a software project things often get complicated. Often, two people try to edit the same file simultaneously. One solution, known as *file locking* or *reserved checkouts*, is to allow only one person to edit each file at a time. This is the only solution with some version control systems, including RCS and SCCS. Currently the usual way to get reserved checkouts with CVS is the `cv`s `admin -l` command (see Section A.6.1 [admin options], page 91). This is not as nicely integrated into CVS as the watch features, described below, but it seems that most people with a need for reserved checkouts find it adequate. It also may be possible to use the watches features described below, together with suitable procedures (not enforced by software), to avoid having two people edit at the same time.

The default model with CVS is known as *unreserved checkouts*. In this model, developers can edit their own *working copy* of a file simultaneously. The first person that commits his changes has no automatic way of knowing that another has started to edit it. Others will get an error message when they try to commit the file. They must then use CVS commands to bring their working copy up to date with the repository revision. This process is almost automatic.

CVS also supports mechanisms which facilitate various kinds of communication, without actually enforcing rules like reserved checkouts do.

The rest of this chapter describes how these various models work, and some of the issues involved in choosing between them.

5.1 File status

Based on what operations you have performed on a checked out file, and what operations others have performed to that file in the repository, one can classify a file in a number of states. The states, as reported by the `status` command, are:

Up-to-date

The file is identical with the latest revision in the repository for the branch in use.

Locally Modified

You have edited the file, and not yet committed your changes.

Locally Added

You have added the file with `add`, and not yet committed your changes.

Locally Removed

You have removed the file with `remove`, and not yet committed your changes.

Needs Checkout

Someone else has committed a newer revision to the repository. The name is slightly misleading; you will ordinarily use `update` rather than `checkout` to get that newer revision.

Needs Patch

Like Needs Checkout, but the CVS server will send a patch rather than the entire file. Sending a patch or sending an entire file accomplishes the same thing.

Needs Merge

Someone else has committed a newer revision to the repository, and you have also made modifications to the file.

File had conflicts on merge

This is like `Locally Modified`, except that a previous `update` command gave a conflict. If you have not already done so, you need to resolve the conflict as described in Section 5.3 [Conflicts example], page 35.

Unknown CVS doesn't know anything about this file. For example, you have created a new file and have not run `add`.

To help clarify the file status, `status` also reports the **Working revision** which is the revision that the file in the working directory derives from, and the **Repository revision** which is the latest revision in the repository for the branch in use.

The options to `status` are listed in Appendix B [Invoking CVS], page 115. For information on its **Sticky tag** and **Sticky date** output, see Section 6.7 [Sticky tags], page 48. For information on its **Sticky options** output, see the `'-k'` option in Section A.18.1 [update options], page 111.

You can think of the `status` and `update` commands as somewhat complementary. You use `update` to bring your files up to date, and you can use `status` to give you some idea of what an `update` would do (of course, the state of the repository might change before you actually run `update`). In fact, if you want a command to display file status in a more brief format than is displayed by the `status` command, you can invoke

```
$ cvs -n -q update
```

The `'-n'` option means to not actually do the update, but merely to display statuses; the `'-q'` option avoids printing the name of each directory. For more information on the `update` command, and these options, see Appendix B [Invoking CVS], page 115.

5.2 Bringing a file up to date

When you want to update or merge a file, use the `update` command. For files that are not up to date this is roughly equivalent to a `checkout` command: the newest revision of the file is extracted from the repository and put in your working copy of the module.

Your modifications to a file are never lost when you use `update`. If no newer revision exists, running `update` has no effect. If you have edited the file, and a newer revision is available, CVS will merge all changes into your working copy.

For instance, imagine that you checked out revision 1.4 and started editing it. In the meantime someone else committed revision 1.5, and shortly after that revision 1.6. If you run `update` on the file now, CVS will incorporate all changes between revision 1.4 and 1.6 into your file.

If any of the changes between 1.4 and 1.6 were made too close to any of the changes you have made, an *overlap* occurs. In such cases a warning is printed, and the resulting file includes both versions of the lines that overlap, delimited by special markers. See Section A.18 [update], page 111, for a complete description of the `update` command.

5.3 Conflicts example

Suppose revision 1.4 of 'driver.c' contains this:

```
#include <stdio.h>

void main()
{
    parse();
    if (nerr == 0)
        gencode();
    else
        fprintf(stderr, "No code generated.\n");
    exit(nerr == 0 ? 0 : 1);
}
```

Revision 1.6 of 'driver.c' contains this:

```
#include <stdio.h>

int main(int argc,
         char **argv)
{
    parse();
    if (argc != 1)
    {
        fprintf(stderr, "tc: No args expected.\n");
        exit(1);
    }
    if (nerr == 0)
        gencode();
    else
        fprintf(stderr, "No code generated.\n");
    exit(!nerr);
}
```

Your working copy of 'driver.c', based on revision 1.4, contains this before you run 'cvs update':

```
#include <stdlib.h>
#include <stdio.h>

void main()
{
    init_scanner();
    parse();
    if (nerr == 0)
        gencode();
    else
        fprintf(stderr, "No code generated.\n");
    exit(nerr == 0 ? EXIT_SUCCESS : EXIT_FAILURE);
}
```

You run 'cvs update':

```
$ cvs update driver.c
RCS file: /usr/local/cvsroot/yoyodyne/tc/driver.c,v
retrieving revision 1.4
retrieving revision 1.6
Merging differences between 1.4 and 1.6 into driver.c
rcsmerge warning: overlaps during merge
cvs update: conflicts found in driver.c
C driver.c
```

CVS tells you that there were some conflicts. Your original working file is saved unmodified in '.#driver.c.1.4'. The new version of 'driver.c' contains this:

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc,
         char **argv)
{
    init_scanner();
    parse();
    if (argc != 1)
    {
        fprintf(stderr, "tc: No args expected.\n");
        exit(1);
    }
    if (nerr == 0)
        gencode();
    else
        fprintf(stderr, "No code generated.\n");
<<<<<<< driver.c
    exit(nerr == 0 ? EXIT_SUCCESS : EXIT_FAILURE);
=====
    exit(!nerr);
>>>>>>> 1.6
}
```

Note how all non-overlapping modifications are incorporated in your working copy, and that the overlapping section is clearly marked with '<<<<<<<', '=====' and '>>>>>>>'.

You resolve the conflict by editing the file, removing the markers and the erroneous line. Suppose you end up with this file:

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc,
         char **argv)
{
    init_scanner();
    parse();
```

```

if (argc != 1)
{
    fprintf(stderr, "tc: No args expected.\n");
    exit(1);
}
if (nerr == 0)
    gencode();
else
    fprintf(stderr, "No code generated.\n");
exit(nerr == 0 ? EXIT_SUCCESS : EXIT_FAILURE);
}

```

You can now go ahead and commit this as revision 1.7.

```

$ cvs commit -m "Initialize scanner. Use symbolic exit values." driver.c
Checking in driver.c;
/usr/local/cvsroot/yoyodyne/tc/driver.c,v <-- driver.c
new revision: 1.7; previous revision: 1.6
done

```

For your protection, CVS will refuse to check in a file if a conflict occurred and you have not resolved the conflict. Currently to resolve a conflict, you must change the timestamp on the file, and must also insure that the file contains no conflict markers. If your file legitimately contains conflict markers (that is, occurrences of '>>>>>>' at the start of a line that don't mark a conflict), then CVS has trouble handling this and you need to start hacking on the CVS/Entries file or other such workarounds.

If you use release 1.04 or later of pcl-cvs (a GNU Emacs front-end for CVS) you can use an Emacs package called emerge to help you resolve conflicts. See the documentation for pcl-cvs.

5.4 Informing others about commits

It is often useful to inform others when you commit a new revision of a file. The '-i' option of the 'modules' file, or the 'logininfo' file, can be used to automate this process. See Section C.1 [modules], page 125. See Section C.7 [logininfo], page 131. You can use these features of CVS to, for instance, instruct CVS to mail a message to all developers, or post a message to a local newsgroup.

5.5 Several developers simultaneously attempting to run CVS

If several developers try to run CVS at the same time, one may get the following message:

```
[11:43:23] waiting for bach's lock in /usr/local/cvsroot/foo
```

CVS will try again every 30 seconds, and either continue with the operation or print the message again, if it still needs to wait. If a lock seems to stick around for an undue amount of time, find the person holding the lock and ask them about the cvs command they are running. If they aren't running a cvs command, look in the repository directory mentioned in the message and remove files which they own whose names start with '#cvs.tf1', '#cvs.rf1', or '#cvs.wf1'.

Note that these locks are to protect CVS's internal data structures and have no relationship to the word *lock* in the sense used by RCS—which refers to reserved checkouts (see Chapter 5 [Multiple developers], page 33).

Any number of people can be reading from a given repository at a time; only when someone is writing do the locks prevent other people from reading or writing.

One might hope for the following property

```
If someone commits some changes in one cvs command,
then an update by someone else will either get all the
changes, or none of them.
```

but CVS does *not* have this property. For example, given the files

```
a/one.c
a/two.c
b/three.c
b/four.c
```

if someone runs

```
cvs ci a/two.c b/three.c
```

and someone else runs `cvs update` at the same time, the person running `update` might get only the change to 'b/three.c' and not the change to 'a/two.c'.

5.6 Mechanisms to track who is editing files

For many groups, use of CVS in its default mode is perfectly satisfactory. Users may sometimes go to check in a modification only to find that another modification has intervened, but they deal with it and proceed with their check in. Other groups prefer to be able to know who is editing what files, so that if two people try to edit the same file they can choose to talk about who is doing what when rather than be surprised at check in time. The features in this section allow such coordination, while retaining the ability of two developers to edit the same file at the same time.

For maximum benefit developers should use `cvs edit` (not `chmod`) to make files read-write to edit them, and `cvs release` (not `rm`) to discard a working directory which is no longer in use, but CVS is not able to enforce this behavior.

5.6.1 Telling CVS to watch certain files

To enable the watch features, you first specify that certain files are to be watched.

cvs watch on [-1R] files ... Command

Specify that developers should run `cvs edit` before editing *files*. CVS will create working copies of *files* read-only, to remind developers to run the `cvs edit` command before working on them.

If *files* includes the name of a directory, CVS arranges to watch all files added to the corresponding repository directory, and sets a default for files added in the future; this allows the user to set notification policies on a per-directory basis. The contents of the directory are processed recursively, unless the `-1`

option is given. The `-R` option can be used to force recursion if the `-l` option is set in `'~/ .cvsrc'` (see Section A.3 [`~/ .cvsrc`], page 86).

If *files* is omitted, it defaults to the current directory.

cvs **watch off** `[-lR] files ...` Command

Do not provide notification about work on *files*. CVS will create working copies of *files* read-write.

The *files* and options are processed as for `cv`s `watch on`.

5.6.2 Telling CVS to notify you

You can tell CVS that you want to receive notifications about various actions taken on a file. You can do this without using `cv`s `watch on` for the file, but generally you will want to use `cv`s `watch on`, so that developers use the `cv`s `edit` command.

cvs **watch add** `[-a action] [-lR] files ...` Command

Add the current user to the list of people to receive notification of work done on *files*.

The `-a` option specifies what kinds of events CVS should notify the user about. *action* is one of the following:

edit Another user has applied the `cv`s `edit` command (described below) to a file.

unedit Another user has applied the `cv`s `unedit` command (described below) or the `cv`s `release` command to a file, or has deleted the file and allowed `cv`s `update` to recreate it.

commit Another user has committed changes to a file.

all All of the above.

none None of the above. (This is useful with `cv`s `edit`, described below.)

The `-a` option may appear more than once, or not at all. If omitted, the action defaults to `all`.

The *files* and options are processed as for the `cv`s `watch` commands.

cvs **watch remove** `[-a action] [-lR] files ...` Command

Remove a notification request established using `cv`s `watch add`; the arguments are the same. If the `-a` option is present, only watches for the specified actions are removed.

When the conditions exist for notification, CVS calls the `'notify'` administrative file. Edit `'notify'` as one edits the other administrative files (see Section 3.4 [Intro administrative files], page 19). This file follows the usual conventions for administrative files (see Section C.3.1 [syntax], page 128), where each line is a regular expression followed by a command to execute. The command should contain a single occurrence of `'%s'` which will be replaced by the user to notify; the rest of the information regarding the notification will be supplied to the command on standard input. The standard thing to put in the `notify` file is the single line:

```
ALL mail %s -s \"CVS notification\"
```

This causes users to be notified by electronic mail.

Note that if you set this up in the straightforward way, users receive notifications on the server machine. One could of course write a `notify` script which directed notifications elsewhere, but to make this easy, CVS allows you to associate a notification address for each user. To do so create a file `users` in `CVSROOT` with a line for each user in the format `user:value`. Then instead of passing the name of the user to be notified to `notify`, CVS will pass the *value* (normally an email address on some other machine).

CVS does not notify you for your own changes. Currently this check is done based on whether the user name of the person taking the action which triggers notification matches the user name of the person getting notification. In fact, in general, the watches features only track one edit by each user. It probably would be more useful if watches tracked each working directory separately, so this behavior might be worth changing.

5.6.3 How to edit a file which is being watched

Since a file which is being watched is checked out read-only, you cannot simply edit it. To make it read-write, and inform others that you are planning to edit it, use the `cvsexit` command. Some systems call this a *checkout*, but CVS uses that term for obtaining a copy of the sources (see Section 2.1 [Getting the source], page 9), an operation which those systems call a *get* or a *fetch*.

cvsexit [*options*] *files* ... Command

Prepare to edit the working files *files*. CVS makes the *files* read-write, and notifies users who have requested `edit` notification for any of *files*.

The `cvsexit` command accepts the same *options* as the `cvsexit add` command, and establishes a temporary watch for the user on *files*; CVS will remove the watch when *files* are `unedited` or `committed`. If the user does not wish to receive notifications, she should specify `-a none`.

The *files* and options are processed as for the `cvsexit watch` commands.

Normally when you are done with a set of changes, you use the `cvsexit commit` command, which checks in your changes and returns the watched files to their usual read-only state. But if you instead decide to abandon your changes, or not to make any changes, you can use the `cvsexit unedit` command.

cvsexit unedit [-1R] *files* ... Command

Abandon work on the working files *files*, and revert them to the repository versions on which they are based. CVS makes those *files* read-only for which users have requested notification using `cvsexit watch on`. CVS notifies users who have requested `unedited` notification for any of *files*.

The *files* and options are processed as for the `cvsexit watch` commands.

If watches are not in use, the `unedited` command probably does not work, and the way to revert to the repository version is to remove the file and then use `cvsexit update` to get a new copy. The meaning is not precisely the same; removing and updating may also bring in some changes which have been made in the repository since the last time you updated.

When using client/server CVS, you can use the `cvsexit` and `cvsexit` commands even if CVS is unable to successfully communicate with the server; the notifications will be sent upon the next successful CVS command.

5.6.4 Information about who is watching and editing

cvsexit [-1R] *files* ... Command

List the users currently watching changes to *files*. The report includes the files being watched, and the mail address of each watcher.

The *files* and options are processed as for the `cvsexit` commands.

cvsexitors [-1R] *files* ... Command

List the users currently working on *files*. The report includes the mail address of each user, the time when the user began working with the file, and the host and path of the working directory containing the file.

The *files* and options are processed as for the `cvsexit` commands.

5.6.5 Using watches with old versions of CVS

If you use the watch features on a repository, it creates 'CVS' directories in the repository and stores the information about watches in that directory. If you attempt to use CVS 1.6 or earlier with the repository, you get an error message such as the following (all on one line):

```
cvsexit: cannot open CVS/Entries for reading:  
No such file or directory
```

and your operation will likely be aborted. To use the watch features, you must upgrade all copies of CVS which use that repository in local or server mode. If you cannot upgrade, use the `watch off` and `watch remove` commands to remove all watches, and that will restore the repository to a state which CVS 1.6 can cope with.

5.7 Choosing between reserved or unreserved checkouts

Reserved and unreserved checkouts each have pros and cons. Let it be said that a lot of this is a matter of opinion or what works given different groups' working styles, but here is a brief description of some of the issues. There are many ways to organize a team of developers. CVS does not try to enforce a certain organization. It is a tool that can be used in several ways.

Reserved checkouts can be very counter-productive. If two persons want to edit different parts of a file, there may be no reason to prevent either of them from doing so. Also, it is common for someone to take out a lock on a file, because they are planning to edit it, but then forget to release the lock.

People, especially people who are familiar with reserved checkouts, often wonder how often conflicts occur if unreserved checkouts are used, and how difficult they are to resolve. The experience with many groups is that they occur rarely and usually are relatively straightforward to resolve.

The rarity of serious conflicts may be surprising, until one realizes that they occur only when two developers disagree on the proper design for a given section of code; such a disagreement suggests that the team has not been communicating properly in the first place. In order to collaborate under *any* source management regimen, developers must agree on the general design of the system; given this agreement, overlapping changes are usually straightforward to merge.

In some cases unreserved checkouts are clearly inappropriate. If no merge tool exists for the kind of file you are managing (for example word processor files or files edited by Computer Aided Design programs), and it is not desirable to change to a program which uses a mergeable data format, then resolving conflicts is going to be unpleasant enough that you generally will be better off to simply avoid the conflicts instead, by using reserved checkouts.

The watches features described above in Section 5.6 [Watches], page 38 can be considered to be an intermediate model between reserved checkouts and unreserved checkouts. When you go to edit a file, it is possible to find out who else is editing it. And rather than having the system simply forbid both people editing the file, it can tell you what the situation is and let you figure out whether it is a problem in that particular case or not. Therefore, for some groups it can be considered the best of both the reserved checkout and unreserved checkout worlds.

6 Revisions and branches

For many uses of CVS, one doesn't need to worry too much about revision numbers; CVS assigns numbers such as 1.1, 1.2, and so on, and that is all one needs to know. However, some people prefer to have more knowledge and control concerning how CVS assigns revision numbers.

If one wants to keep track of a set of revisions involving more than one file, such as which revisions went into a particular release, one uses a *tag*, which is a symbolic revision which can be assigned to a numeric revision in each file.

Another useful feature, especially when maintaining several releases of a software product at once, is the ability to make branches on the revision tree.

6.1 Revision numbers

Each version of a file has a unique *revision number*. Revision numbers look like '1.1', '1.2', '1.3.2.2' or even '1.3.2.2.4.5'. A revision number always has an even number of period-separated decimal integers. By default revision 1.1 is the first revision of a file. Each successive revision is given a new number by increasing the rightmost number by one. The following figure displays a few revisions, with newer revisions to the right.

```
+-----+   +-----+   +-----+   +-----+   +-----+
! 1.1 !----! 1.2 !----! 1.3 !----! 1.4 !----! 1.5 !
+-----+   +-----+   +-----+   +-----+   +-----+
```

CVS is not limited to linear development. The *revision tree* can be split into *branches*, where each branch is a self-maintained line of development. Changes made on one branch can easily be moved back to the main trunk.

Each branch has a *branch number*, consisting of an odd number of period-separated decimal integers. The branch number is created by appending an integer to the revision number where the corresponding branch forked off. Having branch numbers allows more than one branch to be forked off from a certain revision.

between things like release 1 versus release 2 of your product (see Section 6.4 [Tags], page 45). However, if you want to set the numeric revisions, the `-r` option to `cvs commit` can do that. The `-r` option implies the `-f` option, in the sense that it causes the files to be committed even if they are not modified.

For example, to bring all your files up to revision 3.0 (including those that haven't changed), you might invoke:

```
$ cvs commit -r 3.0
```

Note that the number you specify with `-r` must be larger than any existing revision number. That is, if revision 3.0 exists, you cannot `cvs commit -r 1.3`. If you want to maintain several releases in parallel, you need to use a branch (see Chapter 6 [Revisions and branches], page 43).

6.4 Tags—Symbolic revisions

The revision numbers live a life of their own. They need not have anything at all to do with the release numbers of your software product. Depending on how you use CVS the revision numbers might change several times between two releases. As an example, some of the source files that make up RCS 5.6 have the following revision numbers:

```
ci.c          5.21
co.c          5.9
ident.c       5.3
rcs.c         5.12
rcsbase.h     5.11
rcsdiff.c     5.10
rcsedit.c     5.11
rcsfcmp.c     5.9
rcsgen.c      5.10
rcslex.c      5.11
rcsmap.c      5.2
rcsutil.c     5.10
```

You can use the `tag` command to give a symbolic name to a certain revision of a file. You can use the `-v` flag to the `status` command to see all tags that a file has, and which revision numbers they represent. Tag names must start with an uppercase or lowercase letter and can contain uppercase and lowercase letters, digits, `-`, and `_`. The two tag names `BASE` and `HEAD` are reserved for use by CVS. It is expected that future names which are special to CVS will be specially named, for example by starting with `.`, rather than being named analogously to `BASE` and `HEAD`, to avoid conflicts with actual tag names.

You'll want to choose some convention for naming tags, based on information such as the name of the program and the version number of the release. For example, one might take the name of the program, immediately followed by the version number with `.` changed to `-`, so that CVS 1.9 would be tagged with the name `cvs1-9`. If you choose a consistent convention, then you won't constantly be guessing whether a tag is `cvs-1-9` or `cvs1_9` or what. You might even want to consider enforcing your convention in the `taginfo` file (see Section 15.3 [user-defined logging], page 71).

The following example shows how you can add a tag to a file. The commands must be issued inside your working copy of the module. That is, you should issue the command in the directory where 'backend.c' resides.

```
$ cvs tag release-0-4 backend.c
T backend.c
$ cvs status -v backend.c
=====
File: backend.c          Status: Up-to-date

Version:                 1.4      Tue Dec  1 14:39:01 1992
RCS Version:             1.4      /u/cvsroot/yoyodyne/tc/backend.c,v
Sticky Tag:              (none)
Sticky Date:             (none)
Sticky Options:          (none)

Existing Tags:
    release-0-4          (revision: 1.4)
```

There is seldom reason to tag a file in isolation. A more common use is to tag all the files that constitute a module with the same tag at strategic points in the development life-cycle, such as when a release is made.

```
$ cvs tag release-1-0 .
cvs tag: Tagging .
T Makefile
T backend.c
T driver.c
T frontend.c
T parser.c
```

(When you give CVS a directory as argument, it generally applies the operation to all the files in that directory, and (recursively), to any subdirectories that it may contain. See Chapter 8 [Recursive behavior], page 55.)

The `checkout` command has a flag, '-r', that lets you check out a certain revision of a module. This flag makes it easy to retrieve the sources that make up release 1.0 of the module 'tc' at any time in the future:

```
$ cvs checkout -r release-1-0 tc
```

This is useful, for instance, if someone claims that there is a bug in that release, but you cannot find the bug in the current working copy.

You can also check out a module as it was at any given date. See Section A.7.1 [checkout options], page 95.

When you tag more than one file with the same tag you can think about the tag as "a curve drawn through a matrix of filename vs. revision number." Say we have 5 files with the following revisions:


```

file1  file2  file3  file4  file5

1.1    1.1    1.1    1.1  /--1.1*    <-*-- TAG
1.2*-  1.2    1.2    -1.2*-
1.3   \- 1.3*-  1.3    / 1.3
1.4           \ 1.4  / 1.4
           \-1.5*-  1.5
                1.6

```

At some time in the past, the `*` versions were tagged. You can think of the tag as a handle attached to the curve drawn through the tagged revisions. When you pull on the handle, you get all the tagged revisions. Another way to look at it is that you "sight" through a set of revisions that is "flat" along the tagged revisions, like this:

```

file1  file2  file3  file4  file5

                1.1
                1.2
            1.1  1.3
1.1    1.2    1.4    1.1    /
1.2*---1.3*---1.5*---1.2*---1.1  (--- <--- Look here
1.3           1.6    1.3    \_
1.4           1.4
                1.5

```

6.5 What branches are good for

Suppose that release 1.0 of `tc` has been made. You are continuing to develop `tc`, planning to create release 1.1 in a couple of months. After a while your customers start to complain about a fatal bug. You check out release 1.0 (see Section 6.4 [Tags], page 45) and find the bug (which turns out to have a trivial fix). However, the current revision of the sources are in a state of flux and are not expected to be stable for at least another month. There is no way to make a bugfix release based on the newest sources.

The thing to do in a situation like this is to create a *branch* on the revision trees for all the files that make up release 1.0 of `tc`. You can then make modifications to the branch without disturbing the main trunk. When the modifications are finished you can select to either incorporate them on the main trunk, or leave them on the branch.

6.6 Creating a branch

The `rtag` command can be used to create a branch. The `rtag` command is much like `tag`, but it does not require that you have a working copy of the module. See Section A.16 [rtag], page 109. (You can also use the `tag` command; see Section A.17 [tag], page 110).

```
$ cvs rtag -b -r release-1-0 release-1-0-patches tc
```

The `'-b'` flag makes `rtag` create a branch (rather than just a symbolic revision name). `'-r release-1-0'` says that this branch should be rooted at the node (in the revision tree) that corresponds to the tag `'release-1-0'`. Note that the numeric revision number that

matches ‘release-1-0’ will probably be different from file to file. The name of the new branch is ‘release-1-0-patches’, and the module affected is ‘tc’.

To fix the problem in release 1.0, you need a working copy of the branch you just created.

```
$ cvs checkout -r release-1-0-patches tc
$ cvs status -v driver.c backend.c
=====
File: driver.c          Status: Up-to-date

    Version:            1.7      Sat Dec  5 18:25:54 1992
    RCS Version:        1.7      /u/cvsroot/yoyodyne/tc/driver.c,v
    Sticky Tag:         release-1-0-patches (branch: 1.7.2)
    Sticky Date:        (none)
    Sticky Options:     (none)

    Existing Tags:
        release-1-0-patches          (branch: 1.7.2)
        release-1-0                  (revision: 1.7)

=====
File: backend.c        Status: Up-to-date

    Version:            1.4      Tue Dec  1 14:39:01 1992
    RCS Version:        1.4      /u/cvsroot/yoyodyne/tc/backend.c,v
    Sticky Tag:         release-1-0-patches (branch: 1.4.2)
    Sticky Date:        (none)
    Sticky Options:     (none)

    Existing Tags:
        release-1-0-patches          (branch: 1.4.2)
        release-1-0                  (revision: 1.4)
        release-0-4                  (revision: 1.4)
```

As the output from the `status` command shows the branch number is created by adding a digit at the tail of the revision number it is based on. (If ‘release-1-0’ corresponds to revision 1.4, the branch’s revision number will be 1.4.2. For obscure reasons CVS always gives branches even numbers, starting at 2. See Section 6.1 [Revision numbers], page 43.).

6.7 Sticky tags

The ‘-r release-1-0-patches’ flag that was given to `checkout` in the previous example is *sticky*, that is, it will apply to subsequent commands in this directory. If you commit any modifications, they are committed on the branch. You can later merge the modifications into the main trunk. See Chapter 7 [Merging], page 51.

You can use the `status` command to see what sticky tags or dates are set:

```
$ vi driver.c # Fix the bugs
$ cvs commit -m "Fixed initialization bug" driver.c
Checking in driver.c;
```

```

/usr/local/cvsroot/yoyodyne/tc/driver.c,v <-- driver.c
new revision: 1.7.2.1; previous revision: 1.7
done
$ cvs status -v driver.c
=====
File: driver.c          Status: Up-to-date

Version:                1.7.2.1 Sat Dec  5 19:35:03 1992
RCS Version:            1.7.2.1 /u/cvsroot/yoyodyne/tc/driver.c,v
Sticky Tag:             release-1-0-patches (branch: 1.7.2)
Sticky Date:            (none)
Sticky Options:         (none)

Existing Tags:
    release-1-0-patches      (branch: 1.7.2)
    release-1-0              (revision: 1.7)

```

The sticky tags will remain on your working files until you delete them with ‘`cvs update -A`’. The ‘-A’ option retrieves the version of the file from the head of the trunk, and forgets any sticky tags, dates, or options.

Sticky tags are not just for branches. For example, suppose that you want to avoid updating your working directory, to isolate yourself from possibly destabilizing changes other people are making. You can, of course, just refrain from running `cvs update`. But if you want to avoid updating only a portion of a larger tree, then sticky tags can help. If you check out a certain revision (such as 1.4) it will become sticky. Subsequent `cvs update` will not retrieve the latest revision until you reset the tag with `cvs update -A`. Likewise, use of the ‘-D’ option to `update` or `checkout` sets a *sticky date*, which, similarly, causes that date to be used for future retrievals.

Many times you will want to retrieve an old version of a file without setting a sticky tag. The way to do that is with the ‘-p’ option to `checkout` or `update`, which sends the contents of the file to standard output. For example, suppose you have a file named ‘file1’ which existed as revision 1.1, and you then removed it (thus adding a dead revision 1.2). Now suppose you want to add it again, with the same contents it had previously. Here is how to do it:

```

$ cvs update -p -r 1.1 file1 >file1
=====
Checking out file1
RCS:  /tmp/cvs-sanity/cvsroot/first-dir/Attic/file1,v
VERS: 1.1
*****
$ cvs add file1
cvs add: re-adding file file1 (in place of dead revision 1.2)
cvs add: use 'cvs commit' to add this file permanently
$ cvs commit -m test
Checking in file1;
/tmp/cvs-sanity/cvsroot/first-dir/file1,v <-- file1
new revision: 1.3; previous revision: 1.2

```

```
done  
$
```

6.8 Magic branch numbers

This section describes a CVS feature called *magic branches*. For most purposes, you need not worry about magic branches; CVS handles them for you. However, they are visible to you in certain circumstances, so it may be useful to have some idea of how it works.

Externally, branch numbers consist of an odd number of dot-separated decimal integers. See Section 6.1 [Revision numbers], page 43. That is not the whole truth, however. For efficiency reasons CVS sometimes inserts an extra 0 in the second rightmost position (1.2.3 becomes 1.2.0.3, 8.9.10.11.12 becomes 8.9.10.11.0.12 and so on).

CVS does a pretty good job at hiding these so called magic branches, but in a few places the hiding is incomplete:

- The magic branch number appears in the output from `cv`s `log`.
- You cannot specify a symbolic branch name to `cv`s `admin`.

You can use the `admin` command to reassign a symbolic name to a branch the way RCS expects it to be. If `NR4patches` is assigned to the branch 1.4.2 (magic branch number 1.4.0.2) in file `numbers.c` you can do this:

```
$ cvs admin -NR4patches:1.4.2 numbers.c
```

It only works if at least one revision is already committed on the branch. Be very careful so that you do not assign the tag to the wrong number. (There is no way to see how the tag was assigned yesterday).

7 Merging

You can include the changes made between any two revisions into your working copy, by *merging*. You can then commit that revision, and thus effectively copy the changes onto another branch.

7.1 Merging an entire branch

You can merge changes made on a branch into your working copy by giving the ‘-j *branch*’ flag to the `update` command. With one ‘-j *branch*’ option it merges the changes made between the point where the branch forked and newest revision on that branch (into your working copy).

The ‘-j’ stands for “join”.

Consider this revision tree:

```

+-----+   +-----+   +-----+   +-----+
! 1.1 !----! 1.2 !----! 1.3 !----! 1.4 !      <- The main trunk
+-----+   +-----+   +-----+   +-----+
                |
                |
                | +-----+   +-----+
Branch R1fix -> +----! 1.2.2.1 !----! 1.2.2.2 !
                +-----+   +-----+

```

The branch 1.2.2 has been given the tag (symbolic name) ‘R1fix’. The following example assumes that the module ‘mod’ contains only one file, ‘m.c’.

```

$ cvs checkout mod           # Retrieve the latest revision, 1.4

$ cvs update -j R1fix m.c    # Merge all changes made on the branch,
                             # i.e. the changes between revision 1.2
                             # and 1.2.2.2, into your working copy
                             # of the file.

$ cvs commit -m "Included R1fix" # Create revision 1.5.

```

A conflict can result from a merge operation. If that happens, you should resolve it before committing the new revision. See Section 5.3 [Conflicts example], page 35.

The `checkout` command also supports the ‘-j *branch*’ flag. The same effect as above could be achieved with this:

```

$ cvs checkout -j R1fix mod
$ cvs commit -m "Included R1fix"

```

7.2 Merging from a branch several times

Continuing our example, the revision tree now looks like this:

```

+-----+   +-----+   +-----+   +-----+   +-----+
! 1.1 !----! 1.2 !----! 1.3 !----! 1.4 !----! 1.5 !      <- The main trunk
+-----+   +-----+   +-----+   +-----+   +-----+

```

```

          !
          !
          ! +-----+ +-----+
Branch R1fix -> +---! 1.2.2.1 !---! 1.2.2.2 !
                +-----+ +-----+

```

where the starred line represents the merge from the ‘R1fix’ branch to the main trunk, as just discussed.

Now suppose that development continues on the ‘R1fix’ branch:

```

+-----+ +-----+ +-----+ +-----+ +-----+
! 1.1 !---! 1.2 !---! 1.3 !---! 1.4 !---! 1.5 !   <- The main trunk
+-----+ +-----+ +-----+ +-----+ +-----+
          !
          !
          ! +-----+ +-----+ +-----+
Branch R1fix -> +---! 1.2.2.1 !---! 1.2.2.2 !---! 1.2.2.3 !
                +-----+ +-----+ +-----+

```

and then you want to merge those new changes onto the main trunk. If you just use the `cv`s `update -j R1fix m.c` command again, CVS will attempt to merge again the changes which you have already merged, which can have undesirable side effects.

So instead you need to specify that you only want to merge the changes on the branch which have not yet been merged into the trunk. To do that you specify two ‘-j’ options, and CVS merges the changes from the first revision to the second revision. For example, in this case the simplest way would be

```

cv
```

s update -j 1.2.2.2 -j R1fix m.c # Merge changes from 1.2.2.2 to the
 # head of the R1fix branch

The problem with this is that you need to specify the 1.2.2.2 revision manually. A slightly better approach might be to use the date the last merge was done:

```

cv
```

s update -j R1fix:yesterday -j R1fix m.c

Better yet, tag the R1fix branch after every merge into the trunk, and then use that tag for subsequent merges:

```

cv
```

s update -j merged_from_R1fix_to_trunk -j R1fix m.c

7.3 Merging differences between any two revisions

With two ‘-j *revision*’ flags, the `update` (and `checkout`) command can merge the differences between any two revisions into your working file.

```

$ cv
```

s update -j 1.5 -j 1.3 backend.c

will *remove* all changes made between revision 1.3 and 1.5. Note the order of the revisions!

If you try to use this option when operating on multiple files, remember that the numeric revisions will probably be very different between the various files that make up a module. You almost always use symbolic tags rather than revision numbers when operating on multiple files.

7.4 Merging can add or remove files

If the changes which you are merging involve removing or adding some files, `update -j` will reflect such additions or removals.

For example:

```
cv$ update -A
touch a b c
cv$ add a b c ; cv$ ci -m "added" a b c
cv$ tag -b branchtag
cv$ update -r branchtag
touch d ; cv$ add d
rm a ; cv$ rm a
cv$ ci -m "added d, removed a"
cv$ update -A
cv$ update -jbranchtag
```


8 Recursive behavior

Almost all of the subcommands of CVS work recursively when you specify a directory as an argument. For instance, consider this directory structure:

```

$HOME
|
+--tc
|  |
|  +--CVS
|  |   (internal CVS files)
|  +--Makefile
|  +--backend.c
|  +--driver.c
|  +--frontend.c
|  +--parser.c
|  +--man
|  |  |
|  |  +--CVS
|  |  |   (internal CVS files)
|  |  +--tc.1
|  |
|  +--testing
|  |  |
|  |  +--CVS
|  |  |   (internal CVS files)
|  |  +--testpgm.t
|  |  +--test2.t

```

If 'tc' is the current working directory, the following is true:

- 'cvs update testing' is equivalent to


```
cvs update testing/testpgm.t testing/test2.t
```
- 'cvs update testing man' updates all files in the subdirectories
- 'cvs update .' or just 'cvs update' updates all files in the tc module

If no arguments are given to `update` it will update all files in the current working directory and all its subdirectories. In other words, '.' is a default argument to `update`. This is also true for most of the CVS subcommands, not only the `update` command.

The recursive behavior of the CVS subcommands can be turned off with the '-l' option. Conversely, the '-R' option can be used to force recursion if '-l' is specified in '~/.cvsrc' (see Section A.3 [~/cvsrc], page 86).

```
$ cvs update -l          # Don't update files in subdirectories
```


9 Adding files to a directory

To add a new file to a directory, follow these steps.

- You must have a working copy of the directory. See Section 2.1 [Getting the source], page 9.
- Create the new file inside your working copy of the directory.
- Use `'cvs add filename'` to tell cvs that you want to version control the file. If the file contains binary data, specify `'-kb'` (see Chapter 17 [Binary files], page 77).
- Use `'cvs commit filename'` to actually check in the file into the repository. Other developers cannot see the file until you perform this step.

You can also use the `add` command to add a new directory.

Unlike most other commands, the `add` command is not recursive. You cannot even type `'cvs add foo/bar'`! Instead, you have to

```
$ cd foo
$ cvs add bar
```

cvs add `[-k kflag] [-m message] files ...` Command

Schedule *files* to be added to the repository. The files or directories specified with `add` must already exist in the current directory. To add a whole new directory hierarchy to the source repository (for example, files received from a third-party vendor), use the `import` command instead. See Section A.12 [import], page 103.

The added files are not placed in the source repository until you use `commit` to make the change permanent. Doing an `add` on a file that was removed with the `remove` command will undo the effect of the `remove`, unless a `commit` command intervened. See Chapter 10 [Removing files], page 59, for an example.

The `'-k'` option specifies the default way that this file will be checked out; for more information see Section 16.4 [Substitution modes], page 75.

The `'-m'` option specifies a description for the file. This description appears in the history log (if it is enabled, see Section C.10 [history file], page 134). It will also be saved in the version history inside the repository when the file is committed. The `log` command displays this description. The description can be changed using `'admin -t'`. See Section A.6 [admin], page 91. If you omit the `'-m description'` flag, an empty string will be used. You will not be prompted for a description.

For example, the following commands add the file `'backend.c'` to the repository:

```
$ cvs add backend.c
$ cvs commit -m "Early version. Not yet compilable." backend.c
```

When you add a file it is added only on the branch which you are working on (see Chapter 6 [Revisions and branches], page 43). You can later merge the additions to another branch if you want (see Section 7.4 [Merging adds and removals], page 53).

10 Removing files

Modules change. New files are added, and old files disappear. Still, you want to be able to retrieve an exact copy of old releases.

Here is what you can do to remove a file, but remain able to retrieve old revisions:

- Make sure that you have not made any uncommitted modifications to the file. See Section 2.4 [Viewing differences], page 10, for one way to do that. You can also use the `status` or `update` command. If you remove the file without committing your changes, you will of course not be able to retrieve the file as it was immediately before you deleted it.
- Remove the file from your working copy of the directory. You can for instance use `rm`.
- Use `'cvs remove filename'` to tell cvs that you really want to delete the file.
- Use `'cvs commit filename'` to actually perform the removal of the file from the repository.

When you commit the removal of the file, cvs records the fact that the file no longer exists. It is possible for a file to exist on only some branches and not on others, or to re-add another file with the same name later. CVS will correctly create or not create the file, based on the `'-r'` and `'-D'` options specified to `checkout` or `update`.

cvs remove [*options*] *files* ... Command

Schedule file(s) to be removed from the repository (files which have not already been removed from the working directory are not processed). This command does not actually remove the file from the repository until you commit the removal. For a full list of options, see Appendix B [Invoking CVS], page 115.

Here is an example of removing several files:

```
$ cd test
$ rm *.c
$ cvs remove
cvs remove: Removing .
cvs remove: scheduling a.c for removal
cvs remove: scheduling b.c for removal
cvs remove: use 'cvs commit' to remove these files permanently
$ cvs ci -m "Removed unneeded files"
cvs commit: Examining .
cvs commit: Committing .
```

As a convenience you can remove the file and `cvs remove` it in one step, by specifying the `'-f'` option. For example, the above example could also be done like this:

```
$ cd test
$ cvs remove -f *.c
cvs remove: scheduling a.c for removal
cvs remove: scheduling b.c for removal
cvs remove: use 'cvs commit' to remove these files permanently
$ cvs ci -m "Removed unneeded files"
cvs commit: Examining .
```

```
cvcs commit: Committing .
```

If you execute `remove` for a file, and then change your mind before you commit, you can undo the `remove` with an `add` command.

```
$ ls
CVS  ja.h  oj.c
$ rm oj.c
$ cvs remove oj.c
cvs remove: scheduling oj.c for removal
cvs remove: use 'cvs commit' to remove this file permanently
$ cvs add oj.c
U oj.c
cvs add: oj.c, version 1.1.1.1, resurrected
```

If you realize your mistake before you run the `remove` command you can use `update` to resurrect the file:

```
$ rm oj.c
$ cvs update oj.c
cvs update: warning: oj.c was lost
U oj.c
```

When you remove a file it is removed only on the branch which you are working on (see Chapter 6 [Revisions and branches], page 43). You can later merge the removals to another branch if you want (see Section 7.4 [Merging adds and removals], page 53).

11 Removing directories

In concept removing directories is somewhat similar to removing files—you want the directory to not exist in your current working directories, but you also want to be able to retrieve old releases in which the directory existed.

The way that you remove a directory is to remove all the files in it. Then specify the `-P` option to `cvs update`, `cvs checkout`, or `cvs export`, which will cause cvs to remove empty directories from working directories. Probably the best way to do this is to always specify `-P`; if you want an empty directory then put a dummy file (for example `.keepme`) in it to prevent `-P` from removing it.

Note that `-P` is implied by the `-r` or `-D` options of `checkout` and `export`. This way cvs will be able to correctly create the directory or not depending on whether the particular version you are checking out contains any files in that directory.

12 Tracking third-party sources

If you modify a program to better fit your site, you probably want to include your modifications when the next release of the program arrives. CVS can help you with this task.

In the terminology used in CVS, the supplier of the program is called a *vendor*. The unmodified distribution from the vendor is checked in on its own branch, the *vendor branch*. CVS reserves branch 1.1.1 for this use.

When you modify the source and commit it, your revision will end up on the main trunk. When a new release is made by the vendor, you commit it on the vendor branch and copy the modifications onto the main trunk.

Use the `import` command to create and update the vendor branch. After a successful `import` the vendor branch is made the ‘head’ revision, so anyone that checks out a copy of the file gets that revision. When a local modification is committed it is placed on the main trunk, and made the ‘head’ revision.

12.1 Importing a module for the first time

Use the `import` command to check in the sources for the first time. When you use the `import` command to track third-party sources, the *vendor tag* and *release tags* are useful. The *vendor tag* is a symbolic name for the branch (which is always 1.1.1, unless you use the `-b branch` flag—See Section 12.6 [Multiple vendor branches], page 65.). The *release tags* are symbolic names for a particular release, such as ‘FSF_0_04’.

Note that `import` does *not* change the directory in which you invoke it. In particular, it does not set up that directory as a CVS working directory; if you want to work with the sources import them first and then check them out into a different directory (see Section 2.1 [Getting the source], page 9).

Suppose you have the sources to a program called `wdiff` in a directory ‘`wdiff-0.04`’, and are going to make private modifications that you want to be able to use even when new releases are made in the future. You start by importing the source to your repository:

```
$ cd wdiff-0.04
$ cvs import -m "Import of FSF v. 0.04" fsf/wdiff FSF_DIST WDIFF_0_04
```

The vendor tag is named ‘`FSF_DIST`’ in the above example, and the only release tag assigned is ‘`WDIFF_0_04`’.

12.2 Updating a module with the import command

When a new release of the source arrives, you import it into the repository with the same `import` command that you used to set up the repository in the first place. The only difference is that you specify a different release tag this time.

```
$ tar xzf wdiff-0.05.tar.gz
$ cd wdiff-0.05
$ cvs import -m "Import of FSF v. 0.05" fsf/wdiff FSF_DIST WDIFF_0_05
```

For files that have not been modified locally, the newly created revision becomes the head revision. If you have made local changes, `import` will warn you that you must merge the changes into the main trunk, and tell you to use `'checkout -j'` to do so.

```
$ cvs checkout -jFSF_DIST:yesterday -jFSF_DIST wdiff
```

The above command will check out the latest revision of `'wdiff'`, merging the changes made on the vendor branch `'FSF_DIST'` since yesterday into the working copy. If any conflicts arise during the merge they should be resolved in the normal way (see Section 5.3 [Conflicts example], page 35). Then, the modified files may be committed.

Using a date, as suggested above, assumes that you do not import more than one release of a product per day. If you do, you can always use something like this instead:

```
$ cvs checkout -jWDIFF_0_04 -jWDIFF_0_05 wdiff
```

In this case, the two above commands are equivalent.

12.3 Reverting to the latest vendor release

You can also revert local changes completely and return to the latest vendor release by changing the `'head'` revision back to the vendor branch on all files. For example, if you have a checked-out copy of the sources in `'~/work.d/wdiff'`, and you want to revert to the vendor's version for all the files in that directory, you would type:

```
$ cd ~/work.d/wdiff
$ cvs admin -bWDIFF .
```

You must specify the `'-bWDIFF'` without any space after the `'-b'`. See Section A.6.1 [admin options], page 91.

12.4 How to handle binary files with cvs import

Use the `'-k'` wrapper option to tell import which files are binary. See Section C.2 [Wrappers], page 126.

12.5 How to handle keyword substitution with cvs import

The sources which you are importing may contain keywords (see Chapter 16 [Keyword substitution], page 73). For example, the vendor may use CVS or some other system which uses similar keyword expansion syntax. If you just import the files in the default fashion, then the keyword expansions supplied by the vendor will be replaced by keyword expansions supplied by your own copy of CVS. It may be more convenient to maintain the expansions supplied by the vendor, so that this information can supply information about the sources that you imported from the vendor.

To maintain the keyword expansions supplied by the vendor, supply the `'-ko'` option to `cvs import` the first time you import the file. This will turn off keyword expansion for that file entirely, so if you want to be more selective you'll have to think about what you want and use the `'-k'` option to `cvs update` or `cvs admin` as appropriate.

12.6 Multiple vendor branches

All the examples so far assume that there is only one vendor from which you are getting sources. In some situations you might get sources from a variety of places. For example, suppose that you are dealing with a project where many different people and teams are modifying the software. There are a variety of ways to handle this, but in some cases you have a bunch of source trees lying around and what you want to do more than anything else is just to all put them in CVS so that you at least have them in one place.

For handling situations in which there may be more than one vendor, you may specify the `-b` option to `cvs import`. It takes as an argument the vendor branch to import to. The default is `-b 1.1.1`.

For example, suppose that there are two teams, the red team and the blue team, that are sending you sources. You want to import the red team's efforts to branch 1.1.1 and use the vendor tag RED. You want to import the blue team's efforts to branch 1.1.3 and use the vendor tag BLUE. So the commands you might use are:

```
$ cvs import dir RED RED_1-0
$ cvs import -b 1.1.3 dir BLUE BLUE_1-5
```

Note that if your vendor tag does not match your `-b` option, CVS will not detect this case! For example,

```
$ cvs import -b 1.1.3 dir RED RED_1-0
```

Be careful; this kind of mismatch is sure to sow confusion or worse. I can't think of a useful purpose for the ability to specify a mismatch here, but if you discover such a use, don't. CVS is likely to make this an error in some future release.

13 Moving and renaming files

Moving files to a different directory or renaming them is not difficult, but some of the ways in which this works may be non-obvious. (Moving or renaming a directory is even harder. See Chapter 14 [Moving directories], page 69.).

The examples below assume that the file *old* is renamed to *new*.

13.1 The Normal way to Rename

The normal way to move a file is to copy *old* to *new*, and then issue the normal CVS commands to remove *old* from the repository, and add *new* to it.

```
$ mv old new
$ cvs remove old
$ cvs add new
$ cvs commit -m "Renamed old to new" old new
```

This is the simplest way to move a file, it is not error-prone, and it preserves the history of what was done. Note that to access the history of the file you must specify the old or the new name, depending on what portion of the history you are accessing. For example, `cvs log old` will give the log up until the time of the rename.

When *new* is committed its revision numbers will start again, usually at 1.1, so if that bothers you, use the `-r rev` option to commit. For more information see Section 6.3 [Assigning revisions], page 44.

13.2 Moving the history file

This method is more dangerous, since it involves moving files inside the repository. Read this entire section before trying it out!

```
$ cd $CVSR00T/module
$ mv old,v new,v
```

Advantages:

- The log of changes is maintained intact.
- The revision numbers are not affected.

Disadvantages:

- Old releases of the module cannot easily be fetched from the repository. (The file will show up as *new* even in revisions from the time before it was renamed).
- There is no log information of when the file was renamed.
- Nasty things might happen if someone accesses the history file while you are moving it. Make sure no one else runs any of the CVS commands while you move it.

13.3 Copying the history file

This way also involves direct modifications to the repository. It is safe, but not without drawbacks.

```
# Copy the RCS file inside the repository
$ cd $CVSR00T/module
$ cp old,v new,v
# Remove the old file
$ cd ~/module
$ rm old
$ cvs remove old
$ cvs commit old
# Remove all tags from new
$ cvs update new
$ cvs log new           # Remember the non-branch tag names
$ cvs tag -d tag1 new
$ cvs tag -d tag2 new
...
```

By removing the tags you will be able to check out old revisions of the module.

Advantages:

- Checking out old revisions works correctly, as long as you use ‘-rtag’ and not ‘-Ddate’ to retrieve the revisions.
- The log of changes is maintained intact.
- The revision numbers are not affected.

Disadvantages:

- You cannot easily see the history of the file across the rename.

14 Moving and renaming directories

The normal way to rename or move a directory is to rename or move each file within it as described in Section 13.1 [Outside], page 67. Then check out with the ‘-P’ option, as described in Chapter 11 [Removing directories], page 61.

If you really want to hack the repository to rename or delete a directory in the repository, you can do it like this:

1. Inform everyone who has a copy of the module that the directory will be renamed. They should commit all their changes, and remove their working copies of the module, before you take the steps below.
2. Rename the directory inside the repository.

```
$ cd $CVSROOT/module
$ mv old-dir new-dir
```

3. Fix the cvs administrative files, if necessary (for instance if you renamed an entire module).
4. Tell everyone that they can check out the module and continue working.

If someone had a working copy of the module the cvs commands will cease to work for him, until he removes the directory that disappeared inside the repository.

It is almost always better to move the files in the directory instead of moving the directory. If you move the directory you are unlikely to be able to retrieve old releases correctly, since they probably depend on the name of the directories.

15 History browsing

Once you have used CVS to store a version control history—what files have changed when, how, and by whom, there are a variety of mechanisms for looking through the history.

15.1 Log messages

Whenever you commit a file you specify a log message.

To look through the log messages which have been specified for every revision which has been committed, use the `cv`s `log` command (see Section A.13 [log], page 104).

15.2 The history database

You can use the history file (see Section C.10 [history file], page 134) to log various CVS actions. To retrieve the information from the history file, use the `cv`s `history` command (see Section A.11 [history], page 101).

15.3 User-defined logging

You can customize CVS to log various kinds of actions, in whatever manner you choose. These mechanisms operate by executing a script at various times. The script might append a message to a file listing the information and the programmer who created it, or send mail to a group of developers, or, perhaps, post a message to a particular newsgroup. To log commits, use the ‘`loginfo`’ file (see Section C.7 [loginfo], page 131). To log commits, checkouts, exports, and tags, respectively, you can also use the ‘`-i`’, ‘`-o`’, ‘`-e`’, and ‘`-t`’ options in the modules file. For a more flexible way of giving notifications to various users, which requires less in the way of keeping centralized scripts up to date, use the `cv`s `watch add` command (see Section 5.6.2 [Getting Notified], page 39); this command is useful even if you are not using `cv`s `watch on`.

The ‘`taginfo`’ file defines programs to execute when someone executes a `tag` or `rtag` command. The ‘`taginfo`’ file has the standard form for administrative files (see Appendix C [Administrative files], page 125), where each line is a regular expression followed by a command to execute. The arguments passed to the command are, in order, the *tagname*, *operation* (add for `tag`, `mov` for `tag -F`, and `del` for `tag -d`), *repository*, and any remaining are pairs of *filename revision*. A non-zero exit of the filter program will cause the tag to be aborted.

15.4 Annotate command

`cv`s `annotate` [-f1R] [-r rev| -D date] *files* ... Command

For each file in *files*, print the head revision of the trunk, together with information on the last modification for each line. For example:

```
$ cvs annotate ssfile
Annotations for ssfile
*****
```

```
1.1          (mary      27-Mar-96): ssfile line 1
1.2          (joe       28-Mar-96): ssfile line 2
```

The file ‘`ssfile`’ currently contains two lines. The `ssfile line 1` line was checked in by `mary` on March 27. Then, on March 28, `joe` added a line `ssfile line 2`, without modifying the `ssfile line 1` line. This report doesn’t tell you anything about lines which have been deleted or replaced; you need to use `cvs diff` for that (see Section A.9 [diff], page 99).

The options to `cvs annotate` are listed in Appendix B [Invoking CVS], page 115, and can be used to select the files and revisions to annotate. The options are described in more detail in Section A.5 [Common options], page 88.

16 Keyword substitution

As long as you edit source files inside your working copy of a module you can always find out the state of your files via ‘`cvstatus`’ and ‘`cvlog`’. But as soon as you export the files from your development environment it becomes harder to identify which revisions they are.

CVS can use a mechanism known as *keyword substitution* (or *keyword expansion*) to help identifying the files. Embedded strings of the form `$keyword$` and `$keyword:..$` in a file are replaced with strings of the form `$keyword:value$` whenever you obtain a new revision of the file.

16.1 RCS Keywords

This is a list of the keywords:

- `$Author$` The login name of the user who checked in the revision.
- `$Date$` The date and time (UTC) the revision was checked in.
- `$Header$` A standard header containing the full pathname of the RCS file, the revision number, the date (UTC), the author, the state, and the locker (if locked). Files will normally never be locked when you use CVS.
- `Id` Same as `$Header$`, except that the RCS filename is without a path.
- `$Name$` Tag name used to check out this file.
- `$Locker$` The login name of the user who locked the revision (empty if not locked, and thus almost always useless when you are using CVS).
- `Log` The log message supplied during commit, preceded by a header containing the RCS filename, the revision number, the author, and the date (UTC). Existing log messages are *not* replaced. Instead, the new log message is inserted after `$Log:..$`. Each new line is prefixed with a *comment leader* which RCS guesses from the file name extension. It can be changed with `cvadmin -c`. See Section A.6.1 [admin options], page 91. This keyword is useful for accumulating a complete change log in a source file, but for several reasons it can be problematic. See Section 16.5 [Log keyword], page 75.
- `$RCSfile$`
The name of the RCS file without a path.
- `$Revision$`
The revision number assigned to the revision.
- `$Source$` The full pathname of the RCS file.
- `$State$` The state assigned to the revision. States can be assigned with `cvadmin -s`—See Section A.6.1 [admin options], page 91.

16.2 Using keywords

To include a keyword string you simply include the relevant text string, such as `Id`, inside the file, and commit the file. CVS will automatically expand the string as part of the commit operation.

It is common to embed `Id` string in the C source code. This example shows the first few lines of a typical file, after keyword substitution has been performed:

```
static char *rcsid="$Id: samp.c,v 1.5 1993/10/19 14:57:32 ceder Exp $";
/* The following lines will prevent gcc version 2.x
   from issuing an "unused variable" warning. */
#if __GNUC__ == 2
#define USE(var) static void * use_##var = (&use_##var, (void *) &var)
USE (rcsid);
#endif
```

Even though a clever optimizing compiler could remove the unused variable `rcsid`, most compilers tend to include the string in the binary. Some compilers have a `#pragma` directive to include literal text in the binary.

The `ident` command (which is part of the RCS package) can be used to extract keywords and their values from a file. This can be handy for text files, but it is even more useful for extracting keywords from binary files.

```
$ ident samp.c
samp.c:
    $Id: samp.c,v 1.5 1993/10/19 14:57:32 ceder Exp $
$ gcc samp.c
$ ident a.out
a.out:
    $Id: samp.c,v 1.5 1993/10/19 14:57:32 ceder Exp $
```

SCCS is another popular revision control system. It has a command, `what`, which is very similar to `ident` and used for the same purpose. Many sites without RCS have SCCS. Since `what` looks for the character sequence `@(#)` it is easy to include keywords that are detected by either command. Simply prefix the RCS keyword with the magic SCCS phrase, like this:

```
static char *id="@(#) $Id: ab.c,v 1.5 1993/10/19 14:57:32 ceder Exp $";
```

16.3 Avoiding substitution

Keyword substitution has its disadvantages. Sometimes you might want the literal text string `‘$Author$’` to appear inside a file without RCS interpreting it as a keyword and expanding it into something like `‘$Author: ceder $’`.

There is unfortunately no way to selectively turn off keyword substitution. You can use `‘-ko’` (see Section 16.4 [Substitution modes], page 75) to turn off keyword substitution entirely.

In many cases you can avoid using keywords in the source, even though they appear in the final product. For example, the source for this manual contains `‘$@asis{}Author$’` whenever the text `‘$Author$’` should appear. In `nroff` and `troff` you can embed the null-character `\&` inside the keyword for a similar effect.

16.4 Substitution modes

Each file has a stored default substitution mode, and each working directory copy of a file also has a substitution mode. The former is set by the `-k` option to `cvs add` and `cvs admin`; the latter is set by the `-k` or `-A` options to `cvs checkout` or `cvs update`. `cvs diff` also has a `-k` option. For some examples, See Chapter 17 [Binary files], page 77.

The modes available are:

- `'-kkv'` Generate keyword strings using the default form, e.g. `$Revision: 5.7 $` for the `Revision` keyword.
- `'-kkv1'` Like `'-kkv'`, except that a locker's name is always inserted if the given revision is currently locked. This option is normally not useful when `cvs` is used.
- `'-kk'` Generate only keyword names in keyword strings; omit their values. For example, for the `Revision` keyword, generate the string `$Revision$` instead of `$Revision: 5.7 $`. This option is useful to ignore differences due to keyword substitution when comparing different revisions of a file.
- `'-ko'` Generate the old keyword string, present in the working file just before it was checked in. For example, for the `Revision` keyword, generate the string `$Revision: 1.1 $` instead of `$Revision: 5.7 $` if that is how the string appeared when the file was checked in.
- `'-kb'` Like `'-ko'`, but also inhibit conversion of line endings between the canonical form in which they are stored in the repository (linefeed only), and the form appropriate to the operating system in use on the client. For systems, like `unix`, which use linefeed only to terminate lines, this is the same as `'-ko'`. For more information on binary files, see Chapter 17 [Binary files], page 77.
- `'-kv'` Generate only keyword values for keyword strings. For example, for the `Revision` keyword, generate the string `5.7` instead of `$Revision: 5.7 $`. This can help generate files in programming languages where it is hard to strip keyword delimiters like `$Revision: $` from a string. However, further keyword substitution cannot be performed once the keyword names are removed, so this option should be used with care.

One often would like to use `'-kv'` with `cvs export`—see Section A.10 [export], page 100. But be aware that doesn't handle an export containing binary files correctly.

16.5 Problems with the `Log` keyword.

The `Log` keyword is somewhat controversial. As long as you are working on your development system the information is easily accessible even if you do not use the `Log` keyword—just do a `cvs log`. Once you export the file the history information might be useless anyhow.

A more serious concern is that `cvs` is not good at handling `Log` entries when a branch is merged onto the main trunk. Conflicts often result from the merging operation.

People also tend to "fix" the log entries in the file (correcting spelling mistakes and maybe even factual errors). If that is done the information from `cvs log` will not be consistent with the information inside the file. This may or may not be a problem in real life.

It has been suggested that the `Log` keyword should be inserted *last* in the file, and not in the files header, if it is to be used at all. That way the long list of change messages will not interfere with everyday source file browsing.

17 Handling binary files

There are two issues with using CVS to store binary files. The first is that CVS by default convert line endings between the canonical form in which they are stored in the repository (linefeed only), and the form appropriate to the operating system in use on the client (for example, carriage return followed by line feed for Windows NT).

The second is that a binary file might happen to contain data which looks like a keyword (see Chapter 16 [Keyword substitution], page 73), so keyword expansion must be turned off.

The `-kb` option available with some CVS commands insures that neither line ending conversion nor keyword expansion will be done. If you are using an old version of RCS without this option, and you are using an operating system, such as UNIX, which terminates lines with linefeeds only, you can use `-ko` instead; if you are on another operating system, upgrade to a version of RCS, such as 5.7 or later, which supports `-kb`.

Here is an example of how you can create a new file using the `-kb` flag:

```
$ echo '$Id$' > kotest
$ cvs add -kb -m"A test file" kotest
$ cvs ci -m"First checkin; contains a keyword" kotest
```

If a file accidentally gets added without `-kb`, one can use the `cvs admin` command to recover. For example:

```
$ echo '$Id$' > kotest
$ cvs add -m"A test file" kotest
$ cvs ci -m"First checkin; contains a keyword" kotest
$ cvs admin -kb kotest
$ cvs update -A kotest
$ cvs commit -m "make it binary" kotest # For non-unix systems
```

When you check in the file `'kotest'` the keywords are expanded. (Try the above example, and do a `cat kotest` after every command). The `cvs admin -kb` command sets the default keyword substitution method for this file, but it does not alter the working copy of the file that you have. The easiest way to get the unexpanded version of `'kotest'` is `cvs update -A`. If you need to cope with line endings (that is, you are using a CVS client on a non-UNIX system), then you need to check in a new copy of the file, as shown by the `cvs commit` command above.

However, in using `cvs admin -k` to change the keyword expansion, be aware that the keyword expansion mode is not version controlled. This means that, for example, that if you have a text file in old releases, and a binary file with the same name in new releases, CVS provides no way to check out the file in text or binary mode depending on what version you are checking out. There is no good workaround for this problem.

You can also set a default for whether `cvs add` and `cvs import` treat a file as binary based on its name; for example you could say that files whose names end in `'.exe'` are binary. See Section C.2 [Wrappers], page 126. There is currently no way to have CVS detect whether a file is binary based on its contents. The main difficulty with designing such a feature is that it is not clear how to distinguish between binary and non-binary files, and the rules to apply would vary considerably with the operating system.

18 How your build system interacts with CVS

As mentioned in the introduction, CVS does not contain software for building your software from source code. This section describes how various aspects of your build system might interact with CVS.

One common question, especially from people who are accustomed to RCS, is how to make their build get an up to date copy of the sources. The answer to this with CVS is two-fold. First of all, since CVS itself can recurse through directories, there is no need to modify your `Makefile` (or whatever configuration file your build tool uses) to make sure each file is up to date. Instead, just use two commands, first `cvsv -q update` and then `make` or whatever the command is to invoke your build tool. Secondly, you do not necessarily *want* to get a copy of a change someone else made until you have finished your own work. One suggested approach is to first update your sources, then implement, build and test the change you were thinking of, and then commit your sources (updating first if necessary). By periodically (in between changes, using the approach just described) updating your entire tree, you ensure that your sources are sufficiently up to date.

One common need is to record which versions of which source files went into a particular build. This kind of functionality is sometimes called *bill of materials* or something similar. The best way to do this with CVS is to use the `tag` command to record which versions went into a given build (see Section 6.4 [Tags], page 45).

Using CVS in the most straightforward manner possible, each developer will have a copy of the entire source tree which is used in a particular build. If the source tree is small, or if developers are geographically dispersed, this is the preferred solution. In fact one approach for larger projects is to break a project down into smaller separately-compiled subsystems, and arrange a way of releasing them internally so that each developer need check out only those subsystems which are they are actively working on.

Another approach is to set up a structure which allows developers to have their own copies of some files, and for other files to access source files from a central location. Many people have come up with some such a system using features such as the symbolic link feature found in many operating systems, or the `VPATH` feature found in many versions of `make`. One build tool which is designed to help with this kind of thing is Odin (see <ftp://ftp.cs.colorado.edu/pub/distrib/odin>).

19 Compatibility between CVS Versions

The repository format is compatible going back to cvs 1.3. But see Section 5.6.5 [Watches Compatibility], page 41, if you have copies of cvs 1.6 or older and you want to use the optional developer communication features.

The working directory format is compatible going back to cvs 1.5. It did change between cvs 1.3 and cvs 1.5. If you run cvs 1.5 or newer on a working directory checked out with cvs 1.3, cvs will convert it, but to go back to cvs 1.3 you need to check out a new working directory with cvs 1.3.

The remote protocol is interoperable going back to cvs 1.5, but no further (1.5 was the first official release with the remote protocol, but some older versions might still be floating around). In many cases you need to upgrade both the client and the server to take advantage of new features and bugfixes, however.

20 Revision management

If you have read this far, you probably have a pretty good grasp on what CVS can do for you. This chapter talks a little about things that you still have to decide.

If you are doing development on your own using CVS you could probably skip this chapter. The questions this chapter takes up become more important when more than one person is working in a repository.

20.1 When to commit?

Your group should decide which policy to use regarding commits. Several policies are possible, and as your experience with CVS grows you will probably find out what works for you.

If you commit files too quickly you might commit files that do not even compile. If your partner updates his working sources to include your buggy file, he will be unable to compile the code. On the other hand, other persons will not be able to benefit from the improvements you make to the code if you commit very seldom, and conflicts will probably be more common.

It is common to only commit files after making sure that they can be compiled. Some sites require that the files pass a test suite. Policies like this can be enforced using the `commitinfo` file (see Section C.4 [commitinfo], page 128), but you should think twice before you enforce such a convention. By making the development environment too controlled it might become too regimented and thus counter-productive to the real goal, which is to get software written.

Appendix A Guide to CVS commands

This appendix describes the overall structure of CVS commands, and describes some commands in detail (others are described elsewhere; for a quick reference to CVS commands, see Appendix B [Invoking CVS], page 115).

A.1 Overall structure of CVS commands

The overall format of all CVS commands is:

```
cvs [ cvs_options ] cvs_command [ command_options ] [ command_args ]
```

`cvs` The name of the CVS program.

`cvs_options`

Some options that affect all sub-commands of CVS. These are described below.

`cvs_command`

One of several different sub-commands. Some of the commands have aliases that can be used instead; those aliases are noted in the reference manual for that command. There are only two situations where you may omit '`cvs_command`': '`cvs -H`' elicits a list of available commands, and '`cvs -v`' displays version information on CVS itself.

`command_options`

Options that are specific for the command.

`command_args`

Arguments to the commands.

There is unfortunately some confusion between `cvs_options` and `command_options`. '`-l`', when given as a `cvs_option`, only affects some of the commands. When it is given as a `command_option` it has a different meaning, and is accepted by more commands. In other words, do not take the above categorization too seriously. Look at the documentation instead.

A.2 CVS's exit status

CVS can indicate to the calling environment whether it succeeded or failed by setting its *exit status*. The exact way of testing the exit status will vary from one operating system to another. For example in a unix shell script the ' `$?` ' variable will be 0 if the last command returned a successful exit status, or greater than 0 if the exit status indicated failure.

If CVS is successful, it returns a successful status; if there is an error, it prints an error message and returns a failure status. The one exception to this is the `cvs diff` command. It will return a successful status if it found no differences, or a failure status if there were differences or if there was an error. Because this behavior provides no good way to detect errors, in the future it is possible that `cvs diff` will be changed to behave like the other CVS commands.

A.3 Default options and the `~/.cvsrc` file

There are some `command_options` that are used so often that you might have set up an alias or some other means to make sure you always specify that option. One example (the one that drove the implementation of the `.cvsrc` support, actually) is that many people find the default output of the `diff` command to be very hard to read, and that either context diffs or unidiffs are much easier to understand.

The `~/.cvsrc` file is a way that you can add default options to `cvs_commands` within `cvs`, instead of relying on aliases or other shell scripts.

The format of the `~/.cvsrc` file is simple. The file is searched for a line that begins with the same name as the `cvs_command` being executed. If a match is found, then the remainder of the line is split up (at whitespace characters) into separate options and added to the command arguments *before* any options from the command line.

If a command has two names (e.g., `checkout` and `co`), the official name, not necessarily the one used on the command line, will be used to match against the file. So if this is the contents of the user's `~/.cvsrc` file:

```
log -N
diff -u
update -P
co -P
```

the command `cvs checkout foo` would have the `-P` option added to the arguments, as well as `cvs co foo`.

With the example file above, the output from `cvs diff foobar` will be in unidiff format. `cvs diff -c foobar` will provide context diffs, as usual. Getting "old" format diffs would be slightly more complicated, because `diff` doesn't have an option to specify use of the "old" format, so you would need `cvs -f diff foobar`.

In place of the command name you can use `cvs` to specify global options (see Section A.4 [Global options], page 86). For example the following line in `.cvsrc`

```
cvs -z6
```

causes CVS to use compression level 6

A.4 Global options

The available `cvs_options` (that are given to the left of `cvs_command`) are:

`--allow-root=rootdir`

Specify legal CVSROOT directory. See Section 3.9.3.1 [Password authentication server], page 23.

`-b bindir` Use `bindir` as the directory where RCS programs are located. Overrides the setting of the `$RCSBIN` environment variable and any precompiled directory. This parameter should be specified as an absolute pathname.

`-T tempdir`

Use `tempdir` as the directory where temporary files are located. Overrides the setting of the `$TMPDIR` environment variable and any precompiled directory. This parameter should be specified as an absolute pathname.

- d** *cvsrc* *root* *directory*
Use *cvsrc* *root* *directory* as the root directory pathname of the repository. Overrides the setting of the `$CVSRROOT` environment variable. See Chapter 3 [Repository], page 13.
- e** *editor*
Use *editor* to enter revision log information. Overrides the setting of the `$CVSEEDITOR` and `$EDITOR` environment variables. For more information, see Section 2.2 [Committing your changes], page 9.
- f**
Do not read the `~/ .cvsrc` file. This option is most often used because of the non-orthogonality of the `cvsrc` option set. For example, the `cvsrc log` option `-N` (turn off display of tag names) does not have a corresponding option to turn the display on. So if you have `-N` in the `~/ .cvsrc` entry for `log`, you may need to use `-f` to show the tag names.
- H**
- help**
Display usage information about the specified `cvsrc` *command* (but do not actually execute the command). If you don't specify a command name, `cvsrc -H` displays overall help for `cvsrc`, including a list of other help options.
- l**
Do not log the `cvsrc` *command* in the command history (but execute it anyway). See Section A.11 [history], page 101, for information on command history.
- n**
Do not change any files. Attempt to execute the `cvsrc` *command*, but only to issue reports; do not remove, update, or merge any existing files, or create any new files.

Note that `cvsrc` will not necessarily produce exactly the same output as without `-n`. In some cases the output will be the same, but in other cases `cvsrc` will skip some of the processing that would have been required to produce the exact same output.
- Q**
Cause the command to be really quiet; the command will only generate output for serious problems.
- q**
Cause the command to be somewhat quiet; informational messages, such as reports of recursion through subdirectories, are suppressed.
- r**
Make new working files files read-only. Same effect as if the `$CVSREAD` environment variable is set (see Appendix D [Environment variables], page 137). The default is to make working files writable, unless watches are on (see Section 5.6 [Watches], page 38).
- s** *variable*=*value*
Set a user variable (see Section C.11 [Variables], page 134).
- t**
Trace program execution; display messages showing the steps of `cvsrc` activity. Particularly useful with `-n` to explore the potential impact of an unfamiliar command.
- v**
- version**
Display version and copyright information for `cvsrc`.

- `-w` Make new working files read-write. Overrides the setting of the `$CVSREAD` environment variable. Files are created read-write by default, unless `$CVSREAD` is set or `'-r'` is given.
- `-x` Encrypt all communication between the client and the server. Only has an effect on the cvs client. As of this writing, this is only implemented when using a Kerberos connection (see Section 3.9.4 [Kerberos authenticated], page 26). Encryption support is not available by default; it must be enabled using a special configure option, `'--enable-encryption'`, when you build cvs.
- `-z gzip-level`
Set the compression level. Only has an effect on the cvs client.

A.5 Common command options

This section describes the `'command_options'` that are available across several cvs commands. These options are always given to the right of `'cvs_command'`. Not all commands support all of these options; each option is only supported for commands where it makes sense. However, when a command has one of these options you can almost always count on the same behavior of the option as in other commands. (Other command options, which are listed with the individual commands, may have different behavior from one cvs command to the other).

Warning: the `'history'` command is an exception; it supports many options that conflict even with these standard options.

`-D date_spec`

Use the most recent revision no later than *date_spec*. *date_spec* is a single argument, a date description specifying a date in the past.

The specification is *sticky* when you use it to make a private copy of a source file; that is, when you get a working file using `'-D'`, cvs records the date you specified, so that further updates in the same directory will use the same date (for more information on sticky tags/dates, see Section 6.7 [Sticky tags], page 48).

`'-D'` is available with the `checkout`, `diff`, `export`, `history`, `rdiff`, `rtag`, and `update` commands. (The `history` command uses this option in a slightly different way; see Section A.11.1 [history options], page 101).

A wide variety of date formats are supported by cvs. The most standard ones are ISO8601 (from the International Standards Organization) and the Internet e-mail standard (specified in RFC822 as amended by RFC1123).

ISO8601 dates have many variants but a few examples are:

```
1972-09-24
1972-09-24 20:05
```

For more details about ISO8601 dates, see:

```
http://www.ft.uni-erlangen.de/~mskuhn/iso-time.html
```

In addition to the dates allowed in Internet e-mail itself, cvs also allows some of the fields to be omitted. For example:

```
24 Sep 1972 20:05
24 Sep
```

The date is interpreted as being in the local timezone, unless a specific timezone is specified.

These two date formats are preferred. However, CVS currently accepts a wide variety of other date formats. They are intentionally not documented here in any detail, and future versions of CVS might not accept all of them.

One such format is *month/day/year*. This may confuse people who are accustomed to having the month and day in the other order; '1/4/96' is January 4, not April 1.

Remember to quote the argument to the '-D' flag so that your shell doesn't interpret spaces as argument separators. A command using the '-D' flag can look like this:

```
$ cvs diff -D "1 hour ago" cvs.texinfo
```

-f When you specify a particular date or tag to CVS commands, they normally ignore files that do not contain the tag (or did not exist prior to the date) that you specified. Use the '-f' option if you want files retrieved even when there is no match for the tag or date. (The most recent revision of the file will be used).

'-f' is available with these commands: `annotate`, `checkout`, `export`, `rdiff`, `rtag`, and `update`.

Warning: The `commit` command also has a '-f' option, but it has a different behavior for that command. See Section A.8.1 [commit options], page 97.

-k *kflag* Alter the default RCS processing of keywords. See Chapter 16 [Keyword substitution], page 73, for the meaning of *kflag*. Your *kflag* specification is *sticky* when you use it to create a private copy of a source file; that is, when you use this option with the `checkout` or `update` commands, CVS associates your selected *kflag* with the file, and continues to use it with future update commands on the same file until you specify otherwise.

The '-k' option is available with the `add`, `checkout`, `diff`, `import` and `update` commands.

-l Local; run only in current working directory, rather than recursing through subdirectories.

Warning: this is not the same as the overall 'cvs -l' option, which you can specify to the left of a cvs command!

Available with the following commands: `annotate`, `checkout`, `commit`, `diff`, `edit`, `editors`, `export`, `log`, `rdiff`, `remove`, `rtag`, `status`, `tag`, `unedit`, `update`, `watch`, and `watchers`.

-m *message*

Use *message* as log information, instead of invoking an editor.

Available with the following commands: `add`, `commit` and `import`.

- n** Do not run any checkout/commit/tag program. (A program can be specified to run on each of these activities, in the modules database (see Section C.1 [modules], page 125); this option bypasses it).
Warning: this is not the same as the overall ‘cvs -n’ option, which you can specify to the left of a cvs command!
 Available with the **checkout**, **commit**, **export**, and **rtag** commands.
- P** Prune empty directories. See See Chapter 11 [Removing directories], page 61.
- p** Pipe the files retrieved from the repository to standard output, rather than writing them in the current directory. Available with the **checkout** and **update** commands.
- R** Process directories recursively. This is on by default.
 Available with the following commands: **annotate**, **checkout**, **commit**, **diff**, **edit**, **editors**, **export**, **rdiff**, **remove**, **rtag**, **status**, **tag**, **unedit**, **update**, **watch**, and **watchers**.
- r tag** Use the revision specified by the *tag* argument instead of the default *head* revision. As well as arbitrary tags defined with the **tag** or **rtag** command, two special tags are always available: ‘HEAD’ refers to the most recent version available in the repository, and ‘BASE’ refers to the revision you last checked out into the current working directory.
 The tag specification is sticky when you use this with **checkout** or **update** to make your own copy of a file: CVS remembers the tag and continues to use it on future update commands, until you specify otherwise (for more information on sticky tags/dates, see Section 6.7 [Sticky tags], page 48). The tag can be either a symbolic or numeric tag. See Section 6.4 [Tags], page 45.
 Specifying the ‘-q’ global option along with the ‘-r’ command option is often useful, to suppress the warning messages when the RCS history file does not contain the specified tag.
Warning: this is not the same as the overall ‘cvs -r’ option, which you can specify to the left of a cvs command!
 ‘-r’ is available with the **checkout**, **commit**, **diff**, **history**, **export**, **rdiff**, **rtag**, and **update** commands.
- W** Specify file names that should be filtered. You can use this option repeatedly. The spec can be a file name pattern of the same type that you can specify in the ‘.cvswrappers’ file. Available with the following commands: **import**, and **update**.

A.6 admin—Administration front end for rcs

- Requires: repository, working directory.
- Changes: repository.
- Synonym: rcs

This is the CVS interface to assorted administrative RCS facilities, documented in `rsc(1)`. `admin` simply passes all its options and arguments to the `rsc` command; it does no filtering or other processing. This command *does* work recursively, however, so extreme care should be used.

If there is a group whose name matches a compiled in value which defaults to `cvsgroup`, only members of that group can use `cvsgroup admin`. To disallow `cvsgroup admin` for all users, create a group with no users in it.

A.6.1 admin options

Not all valid `rsc` options are useful together with CVS. Some even makes it impossible to use CVS until you undo the effect!

This description of the available options is based on the ‘`rsc(1)`’ man page, but modified to suit readers that are more interested in CVS than RCS.

- Aoldfile* Might not work together with CVS. Append the access list of *oldfile* to the access list of the RCS file.
- alogins* Might not work together with CVS. Append the login names appearing in the comma-separated list *logins* to the access list of the RCS file.
- b[rev]* When used with bare RCS, this option sets the default branch to *rev*; in CVS sticky tags (see Section 6.7 [Sticky tags], page 48) are a better way to decide which branch you want to work on. There is one use with CVS: to revert to the vendor’s version when using vendor branches (see Section 12.3 [Reverting local changes], page 64).
- cstring* Useful with CVS. Sets the comment leader to *string*. The comment leader is printed before every log message line generated by the keyword `Log` (see Chapter 16 [Keyword substitution], page 73). This is useful for programming languages without multi-line comments. RCS initially guesses the value of the comment leader from the file name extension when the file is first committed.
- e [logins]* Might not work together with CVS. Erase the login names appearing in the comma-separated list *logins* from the access list of the RCS file. If *logins* is omitted, erase the entire access list.
- I* Run interactively, even if the standard input is not a terminal.
- i* Useless with CVS. When using bare RCS, this is used to create and initialize a new RCS file, without depositing a revision.
- ksubst* Useful with CVS. Set the default keyword substitution to *subst*. See Chapter 16 [Keyword substitution], page 73. Giving an explicit ‘-k’ option to `cvsgroup update`, `cvsgroup export`, or `cvsgroup checkout` overrides this default.
- l[rev]* Lock the revision with number *rev*. If a branch is given, lock the latest revision on that branch. If *rev* is omitted, lock the latest revision on the default branch. This can be used in conjunction with the ‘`rsclock.pl`’ script in the ‘`contrib`’ directory of the CVS source distribution to provide reserved checkouts (where

only one user can be editing a given file at a time). See the comments in that file for details (and see the ‘README’ file in that directory for disclaimers about the unsupported nature of contrib). According to comments in that file, locking must set to strict (which is the default).

-L Set locking to strict. Strict locking means that the owner of an RCS file is not exempt from locking for checkin. For use with cvs, strict locking must be set; see the discussion under the ‘-l’ option above.

-mrev:msg

Replace the log message of revision *rev* with *msg*.

-Nname[:[rev]]

Act like ‘-n’, except override any previous assignment of *name*.

-nname[:[rev]]

Associate the symbolic name *name* with the branch or revision *rev*. It is normally better to use ‘cvs tag’ or ‘cvs rtag’ instead. Delete the symbolic name if both ‘:’ and *rev* are omitted; otherwise, print an error message if *name* is already associated with another number. If *rev* is symbolic, it is expanded before association. A *rev* consisting of a branch number followed by a ‘.’ stands for the current latest revision in the branch. A ‘:’ with an empty *rev* stands for the current latest revision on the default branch, normally the trunk. For example, ‘rcs -nname: RCS/*’ associates *name* with the current latest revision of all the named RCS files; this contrasts with ‘rcs -nname:\$ RCS/*’ which associates *name* with the revision numbers extracted from keyword strings in the corresponding working files.

-orange

Potentially useful, but dangerous, with cvs (see below). Deletes (*outdates*) the revisions given by *range*. A range consisting of a single revision number means that revision. A range consisting of a branch number means the latest revision on that branch. A range of the form ‘rev1:rev2’ means revisions *rev1* to *rev2* on the same branch, ‘:rev’ means from the beginning of the branch containing *rev* up to and including *rev*, and ‘rev:’ means from revision *rev* to the end of the branch containing *rev*. None of the outdated revisions may have branches or locks.

Due to the way cvs handles branches *rev* cannot be specified symbolically if it is a branch. See Section 6.8 [Magic branch numbers], page 50, for an explanation.

Make sure that no-one has checked out a copy of the revision you outdate. Strange things will happen if he starts to edit it and tries to check it back in. For this reason, this option is not a good way to take back a bogus commit; commit a new revision undoing the bogus change instead (see Section 7.3 [Merging two revisions], page 52).

-q

Run quietly; do not print diagnostics.

-sstate[:rev]

Useful with cvs. Set the state attribute of the revision *rev* to *state*. If *rev* is a branch number, assume the latest revision on that branch. If *rev* is omitted, assume the latest revision on the default branch. Any identifier is acceptable for

state. A useful set of states is ‘Exp’ (for experimental), ‘Stab’ (for stable), and ‘Rel’ (for released). By default, the state of a new revision is set to ‘Exp’ when it is created. The state is visible in the output from *cvs log* (see Section A.13 [log], page 104), and in the ‘\$Log\$’ and ‘\$State\$’ keywords (see Chapter 16 [Keyword substitution], page 73). Note that *CVS* uses the **dead** state for its own purposes; to take a file to or from the **dead** state use commands like *cvs remove* and *cvs add*, not *cvs admin -s*.

- t [*file*] Useful with *CVS*. Write descriptive text from the contents of the named *file* into the RCS file, deleting the existing text. The *file* pathname may not begin with ‘-’. If *file* is omitted, obtain the text from standard input, terminated by end-of-file or by a line containing ‘.’ by itself. Prompt for the text if interaction is possible; see ‘-I’. The descriptive text can be seen in the output from ‘*cvs log*’ (see Section A.13 [log], page 104).
- t-*string* Similar to ‘-t*file*’. Write descriptive text from the *string* into the RCS file, deleting the existing text.
- U Set locking to non-strict. Non-strict locking means that the owner of a file need not lock a revision for checkin. For use with *CVS*, strict locking must be set; see the discussion under the ‘-l’ option above.
- u [*rev*] See the option ‘-l’ above, for a discussion of using this option with *CVS*. Unlock the revision with number *rev*. If a branch is given, unlock the latest revision on that branch. If *rev* is omitted, remove the latest lock held by the caller. Normally, only the locker of a revision may unlock it. Somebody else unlocking a revision breaks the lock. This causes a mail message to be sent to the original locker. The message contains a commentary solicited from the breaker. The commentary is terminated by end-of-file or by a line containing . by itself.
- V*n* Emulate RCS version *n*. Use -V*n* to make an RCS file acceptable to RCS version *n* by discarding information that would confuse version *n*.
- x*suffixes* Useless with *CVS*. Use *suffixes* to characterize RCS files.

A.6.2 admin examples

A.6.2.1 Outdating is dangerous

First, an example of how *not* to use the **admin** command. It is included to stress the fact that this command can be quite dangerous unless you know *exactly* what you are doing.

The ‘-o’ option can be used to *outdate* old revisions from the history file. If you are short on disc this option might help you. But think twice before using it—there is no way short of restoring the latest backup to undo this command!

The next line is an example of a command that you would *not* like to execute.

```
$ cvs admin -o:R_1_02 .
```

The above command will delete all revisions up to, and including, the revision that corresponds to the tag R_1_02. But beware! If there are files that have not changed between

R_1.02 and R_1.03 the file will have *the same* numerical revision number assigned to the tags R_1.02 and R_1.03. So not only will it be impossible to retrieve R_1.02; R_1.03 will also have to be restored from the tapes!

A.6.2.2 Comment leaders

If you use the `Log` keyword and you do not agree with the guess for comment leader that CVS has done, you can enforce your will with `cvs admin -c`. This might be suitable for `nroff` source:

```
$ cvs admin -c'.\" ' *.man
$ rm *.man
$ cvs update
```

The two last steps are to make sure that you get the versions with correct comment leaders in your working files.

A.7 checkout—Check out sources for editing

- Synopsis: `checkout [options] modules...`
- Requires: repository.
- Changes: working directory.
- Synonyms: `co`, `get`

Make a working directory containing copies of the source files specified by *modules*. You must execute `checkout` before using most of the other CVS commands, since most of them operate on your working directory.

The *modules* part of the command are either symbolic names for some collection of source directories and files, or paths to directories or files in the repository. The symbolic names are defined in the ‘`modules`’ file. See Section C.1 [modules], page 125.

Depending on the modules you specify, `checkout` may recursively create directories and populate them with the appropriate source files. You can then edit these source files at any time (regardless of whether other software developers are editing their own copies of the sources); update them to include new changes applied by others to the source repository; or commit your work as a permanent change to the source repository.

Note that `checkout` is used to create directories. The top-level directory created is always added to the directory where `checkout` is invoked, and usually has the same name as the specified module. In the case of a module alias, the created sub-directory may have a different name, but you can be sure that it will be a sub-directory, and that `checkout` will show the relative path leading to each file as it is extracted into your private work area (unless you specify the ‘`-Q`’ global option).

The files created by `checkout` are created read-write, unless the ‘`-r`’ option to CVS (see Section A.4 [Global options], page 86) is specified, the `CVSREAD` environment variable is specified (see Appendix D [Environment variables], page 137), or a watch is in effect for that file (see Section 5.6 [Watches], page 38).

Running `checkout` on a directory that was already built by a prior `checkout` is also permitted, and has the same effect as specifying the ‘`-d`’ option to the `update` command,

that is, any new directories that have been created in the repository will appear in your work area. See Section A.18 [update], page 111.

For the output produced by the `checkout` command see Section A.18.2 [update output], page 113.

A.7.1 checkout options

These standard options are supported by `checkout` (see Section A.5 [Common options], page 88, for a complete description of them):

- `-D date` Use the most recent revision no later than *date*. This option is sticky, and implies `'-P'`. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.
- `-f` Only useful with the `'-D date'` or `'-r tag'` flags. If no matching revision is found, retrieve the most recent revision (instead of ignoring the file).
- `-k kflag` Process RCS keywords according to *kflag*. See `co(1)`. This option is sticky; future updates of this file in this working directory will use the same *kflag*. The `status` command can be viewed to see the sticky options. See Appendix B [Invoking CVS], page 115, for more information on the `status` command.
- `-l` Local; run only in current working directory.
- `-n` Do not run any checkout program (as specified with the `'-o'` option in the modules file; see Section C.1 [modules], page 125).
- `-P` Prune empty directories. See Chapter 14 [Moving directories], page 69.
- `-p` Pipe files to the standard output.
- `-R` Checkout directories recursively. This option is on by default.
- `-r tag` Use revision *tag*. This option is sticky, and implies `'-P'`. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.

In addition to those, you can use these special command options with `checkout`:

- `-A` Reset any sticky tags, dates, or `'-k'` options. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.
- `-c` Copy the module file, sorted, to the standard output, instead of creating or modifying any files or directories in your working directory.
- `-d dir` Create a directory called *dir* for the working files, instead of using the module name. Unless you also use `'-N'`, the paths created under *dir* will be as short as possible.
- `-j tag` With two `'-j'` options, merge changes from the revision specified with the first `'-j'` option to the revision specified with the second `'j'` option, into the working directory.
With one `'-j'` option, merge changes from the ancestor revision to the revision specified with the `'-j'` option, into the working directory. The ancestor revision

is the common ancestor of the revision which the working directory is based on, and the revision specified in the ‘-j’ option.

In addition, each -j option can contain an optional date specification which, when used with branches, can limit the chosen revision to one within a specific date. An optional date is specified by adding a colon (:) to the tag: ‘-j*Symbolic_Tag:Date_Specifier*’.

See Chapter 7 [Merging], page 51.

- N Only useful together with ‘-d *dir*’. With this option, cvs will not shorten module paths in your working directory. (Normally, cvs shortens paths as much as possible when you specify an explicit target directory).
- s Like ‘-c’, but include the status of all modules, and sort it by the status string. See Section C.1 [modules], page 125, for info about the ‘-s’ option that is used inside the modules file to set the module status.

A.7.2 checkout examples

Get a copy of the module ‘tc’:

```
$ cvs checkout tc
```

Get a copy of the module ‘tc’ as it looked one day ago:

```
$ cvs checkout -D yesterday tc
```

A.8 commit—Check files into the repository

- Synopsis: `commit [-lnRf] [-m 'log_message' | -F file] [-r revision] [files...]`
- Requires: working directory, repository.
- Changes: repository.
- Synonym: `ci`

Use `commit` when you want to incorporate changes from your working source files into the source repository.

If you don’t specify particular files to commit, all of the files in your working current directory are examined. `commit` is careful to change in the repository only those files that you have really changed. By default (or if you explicitly specify the ‘-R’ option), files in subdirectories are also examined and committed if they have changed; you can use the ‘-1’ option to limit `commit` to the current directory only.

`commit` verifies that the selected files are up to date with the current revisions in the source repository; it will notify you, and exit without committing, if any of the specified files must be made current first with `update` (see Section A.18 [update], page 111). `commit` does not call the `update` command for you, but rather leaves that for you to do when the time is right.

When all is well, an editor is invoked to allow you to enter a log message that will be written to one or more logging programs (see Section C.1 [modules], page 125, and see Section C.7 [loginfo], page 131) and placed in the RCS history file inside the repository. This log message can be retrieved with the `log` command; See Section A.13 [log], page 104.

You can specify the log message on the command line with the ‘`-m message`’ option, and thus avoid the editor invocation, or use the ‘`-F file`’ option to specify that the argument file contains the log message.

A.8.1 commit options

These standard options are supported by `commit` (see Section A.5 [Common options], page 88, for a complete description of them):

- `-l` Local; run only in current working directory.
- `-n` Do not run any module program.
- `-R` Commit directories recursively. This is on by default.
- `-r revision`
Commit to *revision*. *revision* must be either a branch, or a revision on the main trunk that is higher than any existing revision number (see Section 6.3 [Assigning revisions], page 44). You cannot commit to a specific revision on a branch.

`commit` also supports these options:

- `-F file` Read the log message from *file*, instead of invoking an editor.
- `-f` Note that this is not the standard behavior of the ‘`-f`’ option as defined in See Section A.5 [Common options], page 88.
Force CVS to commit a new revision even if you haven’t made any changes to the file. If the current revision of *file* is 1.7, then the following two commands are equivalent:

```
$ cvs commit -f file
$ cvs commit -r 1.8 file
```

The ‘`-f`’ option disables recursion (i.e., it implies ‘`-1`’). To force CVS to commit a new revision for all files in all subdirectories, you must use ‘`-f -R`’.

- `-m message`
Use *message* as the log message, instead of invoking an editor.

A.8.2 commit examples

A.8.2.1 Committing to a branch

You can commit to a branch revision (one that has an even number of dots) with the ‘`-r`’ option. To create a branch revision, use the ‘`-b`’ option of the `rtag` or `tag` commands (see Section A.17 [tag], page 110 or see Section A.16 [rtag], page 109). Then, either `checkout` or `update` can be used to base your sources on the newly created branch. From that point on, all `commit` changes made within these working sources will be automatically added to a branch revision, thereby not disturbing main-line development in any way. For example, if you had to create a patch to the 1.2 version of the product, even though the 2.0 version is already under development, you might do:

```

$ cvs rtag -b -r FCS1_2 FCS1_2_Patch product_module
$ cvs checkout -r FCS1_2_Patch product_module
$ cd product_module
[[ hack away ]]
$ cvs commit

```

This works automatically since the ‘-r’ option is sticky.

A.8.2.2 Creating the branch after editing

Say you have been working on some extremely experimental software, based on whatever revision you happened to checkout last week. If others in your group would like to work on this software with you, but without disturbing main-line development, you could commit your change to a new branch. Others can then checkout your experimental stuff and utilize the full benefit of cvs conflict resolution. The scenario might look like:

```

[[ hacked sources are present ]]
$ cvs tag -b EXPR1
$ cvs update -r EXPR1
$ cvs commit

```

The `update` command will make the ‘-r EXPR1’ option sticky on all files. Note that your changes to the files will never be removed by the `update` command. The `commit` will automatically commit to the correct branch, because the ‘-r’ is sticky. You could also do like this:

```

[[ hacked sources are present ]]
$ cvs tag -b EXPR1
$ cvs commit -r EXPR1

```

but then, only those files that were changed by you will have the ‘-r EXPR1’ sticky flag. If you hack away, and commit without specifying the ‘-r EXPR1’ flag, some files may accidentally end up on the main trunk.

To work with you on the experimental change, others would simply do

```

$ cvs checkout -r EXPR1 whatever_module

```

A.9 diff—Show differences between revisions

- Synopsis: `diff [-lR] [rcsdiff_options] [[-r rev1 | -D date1] [-r rev2 | -D date2]] [files...]`
- Requires: working directory, repository.
- Changes: nothing.

The `diff` command is used to compare different revisions of files. The default action is to compare your working files with the revisions they were based on, and report any differences that are found.

If any file names are given, only those files are compared. If any directories are given, all files under them will be compared.

The exit status for `diff` is different than for other cvs commands; for details Section A.2 [Exit status], page 85.

A.9.1 diff options

These standard options are supported by `diff` (see Section A.5 [Common options], page 88, for a complete description of them):

- `-D date` Use the most recent revision no later than *date*. See ‘-r’ for how this affects the comparison.
- `-k kflag` Process RCS keywords according to *kflag*. See `co(1)`.
- `-l` Local; run only in current working directory.
- `-R` Examine directories recursively. This option is on by default.
- `-r tag` Compare with revision *tag*. Zero, one or two ‘-r’ options can be present. With no ‘-r’ option, the working file will be compared with the revision it was based on. With one ‘-r’, that revision will be compared to your current working file. With two ‘-r’ options those two revisions will be compared (and your working file will not affect the outcome in any way).
One or both ‘-r’ options can be replaced by a ‘-D *date*’ option, described above.
- `--ifdef=arg`
Output in `ifdef` format. Consult the documentation of your underlying `diff` program concerning the ‘-D’ option to `diff`, for more information on this format.

The following options are passed through to `rcsdiff`, which in turn passes them to `diff`. The exact meaning of the options depends on which `diff` you are using. See the documentation for your `diff` for details.

```
-a -b -B -c -C nlines -d -e -f -h -H -i -n -N -p -s -t -u -U nlines -F regexp -I regexp
-L label -T -V arg -W columns -w -y -0 -1 -2 -3 -4 -5 -6 -7 -8 -9 --binary --brief
--changed-group-format=arg --context[=lines] --ed --expand-tabs --forward-ed
--horizon-lines=arg --ignore-all-space --ignore-blank-lines --ignore-case --
ignore-matching-lines=regexp --ignore-space-change --initial-tab --label=label
--left-column --minimal --new-file --new-line-format=arg --old-line-format=arg
--paginate --rcs --report-identical-files --code-c-function --side-by-side --
show-function-line=regexp --speed-large-files --suppress-common-lines --text
--unchanged-group-format=arg --unified[=lines] --width=columns
```

A.9.2 diff examples

The following line produces a Unidiff (‘-u’ flag) between revision 1.14 and 1.19 of ‘`backend.c`’. Due to the ‘-kk’ flag no keywords are substituted, so differences that only depend on keyword substitution are ignored.

```
$ cvs diff -kk -u -r 1.14 -r 1.19 backend.c
```

Suppose the experimental branch `EXPR1` was based on a set of files tagged `RELEASE_1_0`. To see what has happened on that branch, the following can be used:

```
$ cvs diff -r RELEASE_1_0 -r EXPR1
```

A command like this can be used to produce a context diff between two releases:

```
$ cvs diff -c -r RELEASE_1_0 -r RELEASE_1_1 > diffs
```

If you are maintaining ChangeLogs, a command like the following just before you commit your changes may help you write the ChangeLog entry. All local modifications that have not yet been committed will be printed.

```
$ cvs diff -u | less
```

A.10 export—Export sources from CVS, similar to checkout

- Synopsis: `export [-fNnR] [-r rev|-D date] [-k subst] [-d dir] module...`
- Requires: repository.
- Changes: current directory.

This command is a variant of `checkout`; use it when you want a copy of the source for module without the CVS administrative directories. For example, you might use `export` to prepare source for shipment off-site. This command requires that you specify a date or tag (with `-D` or `-r`), so that you can count on reproducing the source you ship to others.

One often would like to use `-kv` with `cvs export`. This causes any RCS keywords to be expanded such that an import done at some other site will not lose the keyword revision information. But be aware that doesn't handle an export containing binary files correctly. Also be aware that after having used `-kv`, one can no longer use the `ident` command (which is part of the RCS suite—see `ident(1)`) which looks for RCS keyword strings. If you want to be able to use `ident` you must not use `-kv`.

A.10.1 export options

These standard options are supported by `export` (see Section A.5 [Common options], page 88, for a complete description of them):

- `-D date` Use the most recent revision no later than *date*.
- `-f` If no matching revision is found, retrieve the most recent revision (instead of ignoring the file).
- `-l` Local; run only in current working directory.
- `-n` Do not run any checkout program.
- `-R` Export directories recursively. This is on by default.
- `-r tag` Use revision *tag*.

In addition, these options (that are common to `checkout` and `export`) are also supported:

- `-d dir` Create a directory called *dir* for the working files, instead of using the module name. Unless you also use `-N`, the paths created under *dir* will be as short as possible.
- `-k subst` Set keyword expansion mode (see Section 16.4 [Substitution modes], page 75).
- `-N` Only useful together with `-d dir`. With this option, CVS will not shorten module paths in your working directory. (Normally, CVS shortens paths as much as possible when you specify an explicit target directory.)

A.11 history—Show status of files and users

- Synopsis: `history [-report] [-flags] [-options args] [files...]`
- Requires: the file `‘$CVSROOT/CVSROOT/history’`
- Changes: nothing.

CVS can keep a history file that tracks each use of the `checkout`, `commit`, `rtag`, `update`, and `release` commands. You can use `history` to display this information in various formats.

Logging must be enabled by creating the file `‘$CVSROOT/CVSROOT/history’`.

Warning: `history` uses `‘-f’`, `‘-l’`, `‘-n’`, and `‘-p’` in ways that conflict with the normal use inside CVS (see Section A.5 [Common options], page 88).

A.11.1 history options

Several options (shown above as `‘-report’`) control what kind of report is generated:

- `-c` Report on each time `commit` was used (i.e., each time the repository was modified).
- `-e` Everything (all record types); equivalent to specifying `‘-xMACFROGWUT’`.
- `-m module` Report on a particular module. (You can meaningfully use `‘-m’` more than once on the command line.)
- `-o` Report on checked-out modules.
- `-T` Report on all tags.
- `-x type` Extract a particular set of record types *type* from the CVS history. The types are indicated by single letters, which you may specify in combination.

Certain commands have a single record type:

F	release
O	checkout
E	export
T	rtag

One of four record types may result from an update:

C	A merge was necessary but collisions were detected (requiring manual merging).
G	A merge was necessary and it succeeded.
U	A working file was copied from the repository.
W	The working copy of a file was deleted during update (because it was gone from the repository).

One of three record types results from `commit`:

A	A file was added for the first time.
M	A file was modified.
R	A file was removed.

The options shown as ‘`-flags`’ constrain or expand the report without requiring option arguments:

- `-a` Show data for all users (the default is to show data only for the user executing `history`).
- `-l` Show last modification only.
- `-w` Show only the records for modifications done from the same working directory where `history` is executing.

The options shown as ‘`-options args`’ constrain the report based on an argument:

- `-b str` Show data back to a record containing the string `str` in either the module name, the file name, or the repository path.
- `-D date` Show data since `date`. This is slightly different from the normal use of ‘`-D date`’, which selects the newest revision older than `date`.
- `-p repository` Show data for a particular source repository (you can specify several ‘`-p`’ options on the same command line).
- `-r rev` Show records referring to revisions since the revision or tag named `rev` appears in individual RCS files. Each RCS file is searched for the revision or tag.
- `-t tag` Show records since tag `tag` was last added to the the history file. This differs from the ‘`-r`’ flag above in that it reads only the history file, not the RCS files, and is much faster.
- `-u name` Show records for user `name`.

A.12 `import`—Import sources into CVS, using vendor branches

- Synopsis: `import [-options] repository vendortag releasetag...`
- Requires: Repository, source distribution directory.
- Changes: repository.

Use `import` to incorporate an entire source distribution from an outside source (e.g., a source vendor) into your source repository directory. You can use this command both for initial creation of a repository, and for wholesale updates to the module from the outside source. See Chapter 12 [Tracking sources], page 63, for a discussion on this subject.

The `repository` argument gives a directory name (or a path to a directory) under the CVS root directory for repositories; if the directory did not exist, `import` creates it.

When you use `import` for updates to source that has been modified in your source repository (since a prior `import`), it will notify you of any files that conflict in the two

branches of development; use ‘`checkout -j`’ to reconcile the differences, as `import` instructs you to do.

If `cv`s decides a file should be ignored (see Section C.9 [cvsignore], page 133), it does not import it and prints ‘I ’ followed by the filename (see Section A.12.2 [import output], page 104, for a complete description of the output).

If the file ‘`$CVSROOT/CVSROOT/cvswrappers`’ exists, any file whose names match the specifications in that file will be treated as packages and the appropriate filtering will be performed on the file/directory before being imported, See Section C.2 [Wrappers], page 126.

The outside source is saved in a first-level RCS branch, by default 1.1.1. Updates are leaves of this branch; for example, files from the first imported collection of source will be revision 1.1.1.1, then files from the first imported update will be revision 1.1.1.2, and so on.

At least three arguments are required. *repository* is needed to identify the collection of source. *vendortag* is a tag for the entire branch (e.g., for 1.1.1). You must also specify at least one *releasetag* to identify the files at the leaves created each time you execute `import`.

Note that `import` does *not* change the directory in which you invoke it. In particular, it does not set up that directory as a `CVS` working directory; if you want to work with the sources import them first and then check them out into a different directory (see Section 2.1 [Getting the source], page 9).

A.12.1 import options

This standard option is supported by `import` (see Section A.5 [Common options], page 88, for a complete description):

`-m message`

Use *message* as log information, instead of invoking an editor.

There are three additional special options.

`-b branch` See Section 12.6 [Multiple vendor branches], page 65.

`-k subst` Indicate the RCS keyword expansion mode desired. This setting will apply to all files created during the import, but not to any files that previously existed in the repository. See Section 16.4 [Substitution modes], page 75, for a list of valid ‘-k’ settings.

`-I name` Specify file names that should be ignored during import. You can use this option repeatedly. To avoid ignoring any files at all (even those ignored by default), specify ‘-I !’.

name can be a file name pattern of the same type that you can specify in the ‘`.cvsignore`’ file. See Section C.9 [cvsignore], page 133.

`-W spec` Specify file names that should be filtered during import. You can use this option repeatedly.

spec can be a file name pattern of the same type that you can specify in the ‘`.cvswrappers`’ file. See Section C.2 [Wrappers], page 126.

A.12.2 import output

`import` keeps you informed of its progress by printing a line for each file, preceded by one character indicating the status of the file:

U <i>file</i>	The file already exists in the repository and has not been locally modified; a new revision has been created (if necessary).
N <i>file</i>	The file is a new file which has been added to the repository.
C <i>file</i>	The file already exists in the repository but has been locally modified; you will have to merge the changes.
I <i>file</i>	The file is being ignored (see Section C.9 [cvsignore], page 133).
L <i>file</i>	The file is a symbolic link; <code>cvs import</code> ignores symbolic links. People periodically suggest that this behavior should be changed, but if there is a consensus on what it should be changed to, it doesn't seem to be apparent. (Various options in the 'modules' file can be used to recreate symbolic links on checkout, update, etc.; see Section C.1 [modules], page 125.)

A.12.3 import examples

See Chapter 12 [Tracking sources], page 63, and See Section 4.1.1 [From files], page 29.

A.13 log—Print out log information for files

- Synopsis: `log [options] [files...]`
- Requires: repository, working directory.
- Changes: nothing.

Display log information for files. `log` used to call the RCS utility `rlog`. Although this is no longer true in the current sources, this history determines the format of the output and the options, which are not quite in the style of the other `CVS` commands.

The output includes the location of the RCS file, the *head* revision (the latest revision on the trunk), all symbolic names (tags) and some other things. For each revision, the revision number, the author, the number of lines added/deleted and the log message are printed. All times are displayed in Coordinated Universal Time (UTC). (Other parts of `CVS` print times in the local timezone).

Warning: `log` uses '-R' in a way that conflicts with the normal use inside `CVS` (see Section A.5 [Common options], page 88).

A.13.1 log options

By default, `log` prints all information that is available. All other options restrict the output.

- b** Print information about the revisions on the default branch, normally the highest branch on the trunk.

-d *dates* Print information about revisions with a checkin date/time in the range given by the semicolon-separated list of dates. The date formats accepted are those accepted by the ‘-D’ option to many other CVS commands (see Section A.5 [Common options], page 88). Dates can be combined into ranges as follows:

d1<*d2*

d2>*d1* Select the revisions that were deposited between *d1* and *d2*.

<*d*

d> Select all revisions dated *d* or earlier.

d<

>*d* Select all revisions dated *d* or later.

d Select the single, latest revision dated *d* or earlier.

The ‘>’ or ‘<’ characters may be followed by ‘=’ to indicate an inclusive range rather than an exclusive one.

Note that the separator is a semicolon (;).

-h Print only the RCS pathname, working pathname, head, default branch, access list, locks, symbolic names, and suffix.

-l Local; run only in current working directory. (Default is to run recursively).

-N Do not print the list of tags for this file. This option can be very useful when your site uses a lot of tags, so rather than “more”ing over 3 pages of tag information, the log information is presented without tags at all.

-R Print only the name of the RCS history file.

-r *revisions*

Print information about revisions given in the comma-separated list *revisions* of revisions and ranges. The following table explains the available range formats:

rev1:*rev2* Revisions *rev1* to *rev2* (which must be on the same branch).

:*rev* Revisions from the beginning of the branch up to and including *rev*.

rev: Revisions starting with *rev* to the end of the branch containing *rev*.

branch An argument that is a branch means all revisions on that branch.

branch1:*branch2*

A range of branches means all revisions on the branches in that range.

branch. The latest revision in *branch*.

A bare ‘-r’ with no revisions means the latest revision on the default branch, normally the trunk. There can be no space between the ‘-r’ option and its argument.

-s *states* Print information about revisions whose state attributes match one of the states given in the comma-separated list *states*.

-t Print the same as ‘-h’, plus the descriptive text.

-wlogins Print information about revisions checked in by users with login names appearing in the comma-separated list *logins*. If *logins* is omitted, the user's login is assumed. There can be no space between the **'-w'** option and its argument.

log prints the intersection of the revisions selected with the options **'-d'**, **'-s'**, and **'-w'**, intersected with the union of the revisions selected by **'-b'** and **'-r'**.

A.13.2 log examples

Contributed examples are gratefully accepted.

A.14 rdiff—'patch' format diffs between releases

- **rdiff** [-flags] [-V vn] [-r t|-D d [-r t2|-D d2]] modules...
- Requires: repository.
- Changes: nothing.
- Synonym: **patch**

Builds a Larry Wall format **patch(1)** file between two releases, that can be fed directly into the **patch** program to bring an old release up-to-date with the new release. (This is one of the few CVS commands that operates directly from the repository, and doesn't require a prior checkout.) The diff output is sent to the standard output device.

You can specify (using the standard **'-r'** and **'-D'** options) any combination of one or two revisions or dates. If only one revision or date is specified, the patch file reflects differences between that revision or date and the current head revisions in the RCS file.

Note that if the software release affected is contained in more than one directory, then it may be necessary to specify the **'-p'** option to the **patch** command when patching the old sources, so that **patch** is able to find the files that are located in other directories.

A.14.1 rdiff options

These standard options are supported by **rdiff** (see Section A.5 [Common options], page 88, for a complete description of them):

- D date** Use the most recent revision no later than *date*.
- f** If no matching revision is found, retrieve the most recent revision (instead of ignoring the file).
- l** Local; don't descend subdirectories.
- R** Examine directories recursively. This option is on by default.
- r tag** Use revision *tag*.

In addition to the above, these options are available:

- c** Use the context diff format. This is the default format.

- s Create a summary change report instead of a patch. The summary includes information about files that were changed or added between the releases. It is sent to the standard output device. This is useful for finding out, for example, which files have changed between two dates or revisions.
- t A diff of the top two revisions is sent to the standard output device. This is most useful for seeing what the last change to a file was.
- u Use the unidiff format for the context diffs. This option is not available if your diff does not support the unidiff format. Remember that old versions of the `patch` program can't handle the unidiff format, so if you plan to post this patch to the net you should probably not use `-u`.
- V *vn* Expand RCS keywords according to the rules current in RCS version *vn* (the expansion format changed with RCS version 5).

A.14.2 rdiff examples

Suppose you receive mail from `foo@bar.com` asking for an update from release 1.2 to 1.4 of the `tc` compiler. You have no such patches on hand, but with `CVS` that can easily be fixed with a command such as this:

```
$ cvs rdiff -c -r F001_2 -r F001_4 tc | \
$$ Mail -s 'The patches you asked for' foo@bar.com
```

Suppose you have made release 1.3, and forked a branch called `'R_1_3fix'` for bugfixes. `'R_1_3_1'` corresponds to release 1.3.1, which was made some time ago. Now, you want to see how much development has been done on the branch. This command can be used:

```
$ cvs patch -s -r R_1_3_1 -r R_1_3fix module-name
cvs rdiff: Diffing module-name
File ChangeLog,v changed from revision 1.52.2.5 to 1.52.2.6
File foo.c,v changed from revision 1.52.2.3 to 1.52.2.4
File bar.h,v changed from revision 1.29.2.1 to 1.2
```

A.15 release—Indicate that a Module is no longer in use

- `release [-d] directories...`
- Requires: Working directory.
- Changes: Working directory, history log.

This command is meant to safely cancel the effect of `'cvs checkout'`. Since `CVS` doesn't lock files, it isn't strictly necessary to use this command. You can always simply delete your working directory, if you like; but you risk losing changes you may have forgotten, and you leave no trace in the `CVS` history file (see Section C.10 [history file], page 134) that you've abandoned your checkout.

Use `'cvs release'` to avoid these problems. This command checks that no uncommitted changes are present; that you are executing it from immediately above a `CVS` working directory; and that the repository recorded for your files is the same as the repository defined in the module database.

If all these conditions are true, `'cvs release'` leaves a record of its execution (attesting to your intentionally abandoning your checkout) in the `CVS` history log.

A.15.1 release options

The `release` command supports one command option:

-d Delete your working copy of the file if the release succeeds. If this flag is not given your files will remain in your working directory.

Warning: The `release` command deletes all directories and files recursively. This has the very serious side-effect that any directory that you have created inside your checked-out sources, and not added to the repository (using the `add` command; see Chapter 9 [Adding files], page 57) will be silently deleted—even if it is non-empty!

A.15.2 release output

Before `release` releases your sources it will print a one-line message for any file that is not up-to-date.

Warning: Any new directories that you have created, but not added to the CVS directory hierarchy with the `add` command (see Chapter 9 [Adding files], page 57) will be silently ignored (and deleted, if `-d` is specified), even if they contain files.

U file

P file There exists a newer revision of this file in the repository, and you have not modified your local copy of the file (`'U'` and `'P'` mean the same thing).

A file The file has been added to your private copy of the sources, but has not yet been committed to the repository. If you delete your copy of the sources this file will be lost.

R file The file has been removed from your private copy of the sources, but has not yet been removed from the repository, since you have not yet committed the removal. See Section A.8 [commit], page 96.

M file The file is modified in your working directory. There might also be a newer revision inside the repository.

? file *file* is in your working directory, but does not correspond to anything in the source repository, and is not in the list of files for CVS to ignore (see the description of the `'-I'` option, and see Section C.9 [cvsignore], page 133). If you remove your working sources, this file will be lost.

A.15.3 release examples

Release the module, and delete your local working copy of the files.

```
$ cd ..           # You must stand immediately above the
                  # sources when you issue 'cvs release'.
$ cvs release -d tc
You have [0] altered files in this repository.
Are you sure you want to release (and delete) module 'tc': y
$
```

A.16 rtag—Add a symbolic tag to a module

- `rtag [-fahnR] [-b] [-d] [-r tag | -Ddate] symbolic_tag modules...`
- Requires: repository.
- Changes: repository.
- Synonym: `rfreeze`

You can use this command to assign symbolic tags to particular, explicitly specified source revisions in the repository. `rtag` works directly on the repository contents (and requires no prior checkout). Use `tag` instead (see Section A.17 [tag], page 110), to base the selection of revisions on the contents of your working directory.

If you attempt to use a tag name that already exists, CVS will complain and not overwrite that tag. Use the `-F` option to force the new tag value.

A.16.1 rtag options

These standard options are supported by `rtag` (see Section A.5 [Common options], page 88, for a complete description of them):

- `-D date` Tag the most recent revision no later than *date*.
- `-f` Only useful with the `-D date` or `-r tag` flags. If no matching revision is found, use the most recent revision (instead of ignoring the file).
- `-F` Overwrite an existing tag of the same name on a different revision.
- `-l` Local; run only in current working directory.
- `-n` Do not run any tag program that was specified with the `-t` flag inside the `'modules'` file. (see Section C.1 [modules], page 125).
- `-R` Tag directories recursively. This is on by default.
- `-r tag` Only tag those files that contain *tag*. This can be used to rename a tag: tag only the files identified by the old tag, then delete the old tag, leaving the new tag on exactly the same files as the old tag.

In addition to the above common options, these options are available:

- `-a` Use the `-a` option to have `rtag` look in the `'Attic'` (see Section 3.2.3 [Attic], page 16) for removed files that contain the specified tag. The tag is removed from these files, which makes it convenient to re-use a symbolic tag as development continues (and files get removed from the up-coming distribution).
- `-b` Make the tag a branch tag. See Chapter 6 [Revisions and branches], page 43.
- `-d` Delete the tag instead of creating it.

In general, tags (often the symbolic names of software distributions) should not be removed, but the `-d` option is available as a means to remove completely obsolete symbolic names if necessary (as might be the case for an Alpha release, or if you mistagged a module).

A.17 tag—Add a symbolic tag to checked out versions of files

- `tag [-lR] [-b] [-c] [-d] symbolic_tag [files...]`
- Requires: working directory, repository.
- Changes: repository.
- Synonym: freeze

Use this command to assign symbolic tags to the nearest repository versions to your working sources. The tags are applied immediately to the repository, as with `rtag`, but the versions are supplied implicitly by the CVS records of your working files' history rather than applied explicitly.

One use for tags is to record a snapshot of the current sources when the software freeze date of a project arrives. As bugs are fixed after the freeze date, only those changed sources that are to be part of the release need be re-tagged.

The symbolic tags are meant to permanently record which revisions of which files were used in creating a software distribution. The `checkout` and `update` commands allow you to extract an exact copy of a tagged release at any time in the future, regardless of whether files have been changed, added, or removed since the release was tagged.

This command can also be used to delete a symbolic tag, or to create a branch. See the options section below.

If you attempt to use a tag name that already exists, CVS will complain and not overwrite that tag. Use the `-F` option to force the new tag value.

A.17.1 tag options

These standard options are supported by `tag` (see Section A.5 [Common options], page 88, for a complete description of them):

- F Overwrite an existing tag of the same name on a different revision.
- l Local; run only in current working directory.
- R Tag directories recursively. This is on by default.

Two special options are available:

- b The `-b` option makes the tag a branch tag (see Chapter 6 [Revisions and branches], page 43), allowing concurrent, isolated development. This is most useful for creating a patch to a previously released software distribution.
- c The `-c` option checks that all files which are to be tagged are unmodified. This can be used to make sure that you can reconstruct the current file contents.
- d Delete a tag.
If you use `'cvs tag -d symbolic_tag'`, the symbolic tag you specify is deleted instead of being added. Warning: Be very certain of your ground before you delete a tag; doing this permanently discards some historical information, which may later turn out to be valuable.

A.18 update—Bring work tree in sync with repository

- update [-AdfPpR] [-d] [-r tag|-D date] files...
- Requires: repository, working directory.
- Changes: working directory.

After you've run `checkout` to create your private copy of source from the common repository, other developers will continue changing the central source. From time to time, when it is convenient in your development process, you can use the `update` command from within your working directory to reconcile your work with any revisions applied to the source repository since your last checkout or update.

A.18.1 update options

These standard options are available with `update` (see Section A.5 [Common options], page 88, for a complete description of them):

- D *date* Use the most recent revision no later than *date*. This option is sticky, and implies '-P'. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.
- f Only useful with the '-D *date*' or '-r *tag*' flags. If no matching revision is found, retrieve the most recent revision (instead of ignoring the file).
- k *kflag* Process RCS keywords according to *kflag*. See `co(1)`. This option is sticky; future updates of this file in this working directory will use the same *kflag*. The `status` command can be viewed to see the sticky options. See Appendix B [Invoking CVS], page 115, for more information on the `status` command.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- P Prune empty directories. See Chapter 14 [Moving directories], page 69.
- p Pipe files to the standard output.
- R Update directories recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *tag* Retrieve revision *tag*. This option is sticky, and implies '-P'. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.

These special options are also available with `update`.

- A Reset any sticky tags, dates, or '-k' options. See Section 6.7 [Sticky tags], page 48, for more information on sticky tags/dates.
- d Create any directories that exist in the repository if they're missing from the working directory. Normally, `update` acts only on directories and files that were already enrolled in your working directory.
This is useful for updating directories that were created in the repository since the initial checkout; but it has an unfortunate side effect. If you deliberately avoided certain directories in the repository when you created your working

directory (either through use of a module name or by listing explicitly the files and directories you wanted on the command line), then updating with ‘-d’ will create those directories, which may not be what you want.

- I *name* Ignore files whose names match *name* (in your working directory) during the update. You can specify ‘-I’ more than once on the command line to specify several files to ignore. Use ‘-I !’ to avoid ignoring any files at all. See Section C.9 [cvsignore], page 133, for other ways to make CVS ignore some files.
- W*spec* Specify file names that should be filtered during update. You can use this option repeatedly.
spec can be a file name pattern of the same type that you can specify in the ‘.cvswrappers’ file. See Section C.2 [Wrappers], page 126.
- j*revision* With two ‘-j’ options, merge changes from the revision specified with the first ‘-j’ option to the revision specified with the second ‘j’ option, into the working directory.
 With one ‘-j’ option, merge changes from the ancestor revision to the revision specified with the ‘-j’ option, into the working directory. The ancestor revision is the common ancestor of the revision which the working directory is based on, and the revision specified in the ‘-j’ option.
 In addition, each -j option can contain an optional date specification which, when used with branches, can limit the chosen revision to one within a specific date. An optional date is specified by adding a colon (:) to the tag: ‘-j*Symbolic_Tag:Date_Specifier*’.
 See Chapter 7 [Merging], page 51.

A.18.2 update output

update and **checkout** keep you informed of its progress by printing a line for each file, preceded by one character indicating the status of the file:

- U** *file* The file was brought up to date with respect to the repository. This is done for any file that exists in the repository but not in your source, and for files that you haven’t changed but are not the most recent versions available in the repository.
- P** *file* Like ‘U’, but the CVS server sends a patch instead of an entire file. These two things accomplish the same thing.
- A** *file* The file has been added to your private copy of the sources, and will be added to the source repository when you run **commit** on the file. This is a reminder to you that the file needs to be committed.
- R** *file* The file has been removed from your private copy of the sources, and will be removed from the source repository when you run **commit** on the file. This is a reminder to you that the file needs to be committed.
- M** *file* The file is modified in your working directory.

'M' can indicate one of two states for a file you're working on: either there were no modifications to the same file in the repository, so that your file remains as you last saw it; or there were modifications in the repository as well as in your copy, but they were merged successfully, without conflict, in your working directory.

CVS will print some messages if it merges your work, and a backup copy of your working file (as it looked before you ran **update**) will be made. The exact name of that file is printed while **update** runs.

- C** *file* A conflict was detected while trying to merge your changes to *file* with changes from the source repository. *file* (the copy in your working directory) is now the output of the `rcsmerge(1)` command on the two revisions; an unmodified copy of your file is also in your working directory, with the name `‘.#file.revision’` where *revision* is the RCS revision that your modified file started from. Resolve the conflict as described in Section 5.3 [Conflicts example], page 35 (Note that some systems automatically purge files that begin with `‘.#’` if they have not been accessed for a few days. If you intend to keep a copy of your original file, it is a very good idea to rename it.) Under VMS, the file name starts with `‘_’` rather than `‘.#’`.
- ?** *file* *file* is in your working directory, but does not correspond to anything in the source repository, and is not in the list of files for CVS to ignore (see the description of the `‘-I’` option, and see Section C.9 [cvsignore], page 133).

Appendix B Quick reference to CVS commands

This appendix describes how to invoke CVS, with references to where each command or feature is described in detail. Other relevant references are the ‘--help’/‘-H’ option to CVS (see Section A.4 [Global options], page 86) and [Index], page 153.

add [*options*] [*files...*]

Add a new file/directory. See Chapter 9 [Adding files], page 57.

-k *kflag* Set keyword expansion.

-m *msg* Set file description.

admin [*options*] [*files...*]

Administration of history files in the repository. See Section A.6 [admin], page 91.

-b[*rev*] Set default branch.

-c*string* Set comment leader.

-k*subst* Set keyword substitution. See Chapter 16 [Keyword substitution], page 73.

-l[*rev*] Lock revision *rev*, or latest revision.

-m*rev:msg*

Replace the log message of revision *rev* with *msg*.

-orange Delete revisions from the history files

-q Run quietly; do not print diagnostics.

-s*state[:rev]*

Set the state.

-t Set file description from standard input.

-t*file* Set file description from *file*.

-t-*string* Set file description to *string*.

-u[*rev*] Unlock revision *rev*, or latest revision.

annotate [*options*] [*files...*]

Show last revision where each line was modified. See Section 15.4 [annotate], page 71.

-D *date* Annotate the most recent revision no later than *date*. See Section A.5 [Common options], page 88.

-f Use head revision if tag/date not found. See Section A.5 [Common options], page 88.

-l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.

-R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

-r tag Annotate revision *tag*. See Section A.5 [Common options], page 88.

checkout [*options*] *modules...*

Get a copy of the sources. See Section A.7 [checkout], page 94.

-A Reset any sticky tags/date/options. See Section 6.7 [Sticky tags], page 48 and Chapter 16 [Keyword substitution], page 73.

-c Output the module database. See Section A.7.1 [checkout options], page 95.

-D date Check out revisions as of *date* (is sticky). See Section A.5 [Common options], page 88.

-d dir Check out into *dir*. See Section A.7.1 [checkout options], page 95.

-f Use head revision if tag/date not found. See Section A.5 [Common options], page 88.

-j rev Merge in changes. See Section A.7.1 [checkout options], page 95.

-k kflag Use *kflag* keyword expansion. See Section 16.4 [Substitution modes], page 75.

-l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.

-N Don't shorten module paths if **-d** specified. See Section A.7.1 [checkout options], page 95.

-n Do not run module program (if any). See Section A.7.1 [checkout options], page 95.

-P Prune empty directories. See Chapter 14 [Moving directories], page 69.

-p Check out files to standard output (avoids stickiness). See Section A.7.1 [checkout options], page 95.

-R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

-r tag Checkout revision *tag* (is sticky). See Section A.5 [Common options], page 88.

-s Like **-c**, but include module status. See Section A.7.1 [checkout options], page 95.

commit [*options*] [*files...*]

Check changes into the repository. See Section A.8 [commit], page 96.

-F file Read log message from *file*. See Section A.8.1 [commit options], page 97.

- f Force the file to be committed; disables recursion. See Section A.8.1 [commit options], page 97.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- m *msg* Use *msg* as log message. See Section A.8.1 [commit options], page 97.
- n Do not run module program (if any). See Section A.8.1 [commit options], page 97.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *rev* Commit to *rev*. See Section A.8.1 [commit options], page 97.

diff [*options*] [*files...*]

Show differences between revisions. See Section A.9 [diff], page 99. In addition to the options shown below, accepts a wide variety of options to control output style, for example ‘-c’ for context diffs.

- D *date1* Diff revision for *date* against working file. See Section A.9.1 [diff options], page 99.
- D *date2* Diff *rev1/date1* against *date2*. See Section A.9.1 [diff options], page 99.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- N Include diffs for added and removed files. See Section A.9.1 [diff options], page 99.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *rev1* Diff revision for *rev1* against working file. See Section A.9.1 [diff options], page 99.
- r *rev2* Diff *rev1/date1* against *rev2*. See Section A.9.1 [diff options], page 99.

edit [*options*] [*files...*]

Get ready to edit a watched file. See Section 5.6.3 [Editing files], page 40.

- a *actions* Specify actions for temporary watch, where *actions* is **edit**, **unedit**, **commit**, **all**, or **none**. See Section 5.6.3 [Editing files], page 40.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

editors [*options*] [*files...*]

See who is editing a watched file. See Section 5.6.4 [Watch information], page 41.

- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

export [*options*] *modules...*

Export files from CVS. See Section A.10 [export], page 100.

- D *date* Check out revisions as of *date*. See Section A.5 [Common options], page 88.
- d *dir* Check out into *dir*. See Section A.10.1 [export options], page 100.
- f Use head revision if tag/date not found. See Section A.5 [Common options], page 88.
- k *kflag* Use *kflag* keyword expansion. See Section 16.4 [Substitution modes], page 75.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- N Don't shorten module paths if -d specified. See Section A.10.1 [export options], page 100.
- n Do not run module program (if any). See Section A.10.1 [export options], page 100.
- P Prune empty directories. See Chapter 14 [Moving directories], page 69.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *tag* Checkout revision *tag* (is sticky). See Section A.5 [Common options], page 88.

history [*options*] [*files...*]

Show repository access history. See Section A.11 [history], page 101.

- a All users (default is self). See Section A.11.1 [history options], page 101.
- b *str* Back to record with *str* in module/file/repos field. See Section A.11.1 [history options], page 101.
- c Report on committed (modified) files. See Section A.11.1 [history options], page 101.
- D *date* Since *date*. See Section A.11.1 [history options], page 101.
- e Report on all record types. See Section A.11.1 [history options], page 101.
- l Last modified (committed or modified report). See Section A.11.1 [history options], page 101.

- m *module*** Report on *module* (repeatable). See Section A.11.1 [history options], page 101.
 - n *module*** In *module*. See Section A.11.1 [history options], page 101.
 - o** Report on checked out modules. See Section A.11.1 [history options], page 101.
 - r *rev*** Since revision *rev*. See Section A.11.1 [history options], page 101.
 - T** Produce report on all TAGs. See Section A.11.1 [history options], page 101.
 - t *tag*** Since tag record placed in history file (by anyone). See Section A.11.1 [history options], page 101.
 - u *user*** For user *user* (repeatable). See Section A.11.1 [history options], page 101.
 - w** Working directory must match. See Section A.11.1 [history options], page 101.
 - x *types*** Report on *types*, one or more of **TOEFWUCGMAR**. See Section A.11.1 [history options], page 101.
 - z *zone*** Output for time zone *zone*. See Section A.11.1 [history options], page 101.
- import** [*options*] *repository vendor-tag release-tags...*
 Import files into CVS, using vendor branches. See Section A.12 [import], page 103.
- b *bra*** Import to vendor branch *bra*. See Section 12.6 [Multiple vendor branches], page 65.
 - d** Use the file's modification time as the time of import. See Section A.12.1 [import options], page 103.
 - k *kflag*** Set default RCS keyword substitution mode. See Section A.12.1 [import options], page 103.
 - m *msg*** Use *msg* for log message. See Section A.12.1 [import options], page 103.
 - I *ign*** More files to ignore (! to reset). See Section A.12.1 [import options], page 103.
 - W *spec*** More wrappers. See Section A.12.1 [import options], page 103.
- init** Create a CVS repository if it doesn't exist. See Section 3.6 [Creating a repository], page 20.
- log** [*options*] [*files...*]
 Print out history information for files. See Section A.13 [log], page 104.

- b Only list revisions on the default branch. See Section A.13.1 [log options], page 105.
 - d *dates* Specify dates (*d1*<*d2* for range, *d* for latest before). See Section A.13.1 [log options], page 105.
 - h Only print header. See Section A.13.1 [log options], page 105.
 - l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
 - N Do not list tags. See Section A.13.1 [log options], page 105.
 - R Only print name of RCS file. See Section A.13.1 [log options], page 105.
 - r *revs* Only list revisions *revs*. See Section A.13.1 [log options], page 105.
 - s *states* Only list revisions with specified states. See Section A.13.1 [log options], page 105.
 - t Only print header and descriptive text. See Section A.13.1 [log options], page 105.
 - w *logins* Only list revisions checked in by specified logins. See Section A.13.1 [log options], page 105.
- login** Prompt for password for authenticating server. See Section 3.9.3.2 [Password authentication client], page 25.
- logout** Remove stored password for authenticating server. See Section 3.9.3.2 [Password authentication client], page 25.
- rdiff** [*options*] *modules...*
 Show differences between releases. See Section A.14 [rdiff], page 106.
- c Context diff output format (default). See Section A.14.1 [rdiff options], page 107.
 - D *date* Select revisions based on *date*. See Section A.5 [Common options], page 88.
 - f Use head revision if tag/date not found. See Section A.5 [Common options], page 88.
 - l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
 - R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
 - r *rev* Select revisions based on *rev*. See Section A.5 [Common options], page 88.
 - s Short patch - one liner per file. See Section A.14.1 [rdiff options], page 107.

- t Top two diffs - last change made to the file. See Section A.9.1 [diff options], page 99.
- u Unidiff output format. See Section A.14.1 [rdiff options], page 107.
- V *vers* Use RCS Version *vers* for keyword expansion. See Section A.14.1 [rdiff options], page 107.

release [*options*] *directory*

Indicate that a directory is no longer in use. See Section A.15 [release], page 108.

- d Delete the given directory. See Section A.15.1 [release options], page 108.

remove [*options*] [*files...*]

Remove an entry from the repository. See Chapter 10 [Removing files], page 59.

- f Delete the file before removing it. See Chapter 10 [Removing files], page 59.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

rtag [*options*] *tag modules...*

Add a symbolic tag to a module. See Section A.16 [rtag], page 109.

- a Clear tag from removed files that would not otherwise be tagged. See Section A.16.1 [rtag options], page 109.
- b Create a branch named *tag*. See Section A.16.1 [rtag options], page 109.
- D *date* Tag revisions as of *date*. See Section A.16.1 [rtag options], page 109.
- d Delete the given tag. See Section A.16.1 [rtag options], page 109.
- F Move tag if it already exists. See Section A.16.1 [rtag options], page 109.
- f Force a head revision match if tag/date not found. See Section A.16.1 [rtag options], page 109.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- n No execution of tag program. See Section A.16.1 [rtag options], page 109.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *tag* Tag existing tag *tag*. See Section A.16.1 [rtag options], page 109.

status [*options*] *files...*

Display status information in a working directory. See Section 5.1 [File status], page 33.

- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- v Include tag information for file. See Section 6.4 [Tags], page 45.

tag [*options*] *tag* [*files...*]

Add a symbolic tag to checked out version of files. See Section A.17 [tag], page 110.

- b Create a branch named *tag*. See Section A.17.1 [tag options], page 110.
- D *date* Tag revisions as of *date*. See Section A.17.1 [tag options], page 110.
- d Delete the given tag. See Section A.17.1 [tag options], page 110.
- F Move tag if it already exists. See Section A.17.1 [tag options], page 110.
- f Force a head revision match if tag/date not found. See Section A.17.1 [tag options], page 110.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- n No execution of tag program. See Section A.17.1 [tag options], page 110.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *tag* Tag existing tag *tag*. See Section A.17.1 [tag options], page 110.

unedit [*options*] [*files...*]

Undo an edit command. See Section 5.6.3 [Editing files], page 40.

- a *actions* Specify actions for temporary watch, where *actions* is **edit**, **unedit**, **commit**, **all**, or **none**. See Section 5.6.3 [Editing files], page 40.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

update [*options*] [*files...*]

Bring work tree in sync with repository. See Section A.18 [update], page 111.

- A Reset any sticky tags/date/options. See Section 6.7 [Sticky tags], page 48 and Chapter 16 [Keyword substitution], page 73.
- D *date* Check out revisions as of *date* (is sticky). See Section A.5 [Common options], page 88.
- d Create directories. See Section A.18.1 [update options], page 111.

- f Use head revision if tag/date not found. See Section A.5 [Common options], page 88.
- I *ign* More files to ignore (! to reset). See Section A.12.1 [import options], page 103.
- j *rev* Merge in changes. See Section A.18.1 [update options], page 111.
- k *kflag* Use *kflag* keyword expansion. See Section 16.4 [Substitution modes], page 75.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- P Prune empty directories. See Chapter 14 [Moving directories], page 69.
- p Check out files to standard output (avoids stickiness). See Section A.18.1 [update options], page 111.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.
- r *tag* Checkout revision *tag* (is sticky). See Section A.5 [Common options], page 88.
- W *spec* More wrappers. See Section A.12.1 [import options], page 103.

watch [on|off|add|remove] [*options*] [*files...*]

on/off: turn on/off read-only checkouts of files. See Section 5.6.1 [Setting a watch], page 38.

add/remove: add or remove notification on actions. See Section 5.6.2 [Getting Notified], page 39.

- a *actions* Specify actions for temporary watch, where *actions* is **edit**, **unedit**, **commit**, **all**, or **none**. See Section 5.6.3 [Editing files], page 40.
- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

watchers [*options*] [*files...*]

See who is watching a file. See Section 5.6.4 [Watch information], page 41.

- l Local; run only in current working directory. See Chapter 8 [Recursive behavior], page 55.
- R Operate recursively (default). See Chapter 8 [Recursive behavior], page 55.

Appendix C Reference manual for Administrative files

Inside the repository, in the directory ‘`$CVSROOT/CVSROOT`’, there are a number of supportive files for `CVS`. You can use `CVS` in a limited fashion without any of them, but if they are set up properly they can help make life easier. For a discussion of how to edit them, See Section 3.4 [Intro administrative files], page 19.

The most important of these files is the ‘`modules`’ file, which defines the modules inside the repository.

C.1 The modules file

The ‘`modules`’ file records your definitions of names for collections of source code. `CVS` will use these definitions if you use `CVS` to update the modules file (use normal commands like `add`, `commit`, etc).

The ‘`modules`’ file may contain blank lines and comments (lines beginning with ‘`#`’) as well as module definitions. Long lines can be continued on the next line by specifying a backslash (‘`\`’) as the last character on the line.

A module definition is a single line of the ‘`modules`’ file, in either of two formats. In both cases, *mname* represents the symbolic module name, and the remainder of the line is its definition.

mname -a *aliases*...

This represents the simplest way of defining a module *mname*. The ‘`-a`’ flags the definition as a simple alias: `CVS` will treat any use of *mname* (as a command argument) as if the list of names *aliases* had been specified instead. *aliases* may contain either other module names or paths. When you use paths in aliases, `checkout` creates all intermediate directories in the working directory, just as if the path had been specified explicitly in the `CVS` arguments.

mname [options] *dir* [*files...*] [&*module...*]

In the simplest case, this form of module definition reduces to ‘*mname dir*’. This defines all the files in directory *dir* as module *mname*. *dir* is a relative path (from `$CVSROOT`) to a directory of source in the source repository. In this case, on checkout, a single directory called *mname* is created as a working directory; no intermediate directory levels are used by default, even if *dir* was a path involving several directory levels.

By explicitly specifying files in the module definition after *dir*, you can select particular files from directory *dir*. The sample definition for ‘`modules`’ is an example of a module defined with a single file from a particular directory. Here is another example:

```
m4test unsupported/gnu/m4 foreach.m4 forloop.m4
```

With this definition, executing ‘`cvs checkout m4test`’ will create a single working directory ‘`m4test`’ containing the two files listed, which both come from a common directory several levels deep in the `CVS` source repository.

A module definition can refer to other modules by including ‘&*module*’ in its definition. `checkout` creates a subdirectory for each such module, in the directory containing the module. For example, if `modules` contains

m4test &unsupported

then a checkout will create an **m4test** directory which contains a directory called **unsupported**, which in turns contains all the directories and files which live there.

- d name** Name the working directory something other than the module name.
- e prog** Specify a program *prog* to run whenever files in a module are exported. *prog* runs with a single argument, the module name.
- i prog** Specify a program *prog* to run whenever files in a module are committed. *prog* runs with a single argument, the full pathname of the affected directory in a source repository. The **'commitinfo'**, **'loginfo'**, and **'verifymsg'** files provide other ways to call a program on commit.
- o prog** Specify a program *prog* to run whenever files in a module are checked out. *prog* runs with a single argument, the module name.
- s status** Assign a status to the module. When the module file is printed with **'cvs checkout -s'** the modules are sorted according to primarily module status, and secondarily according to the module name. This option has no other meaning. You can use this option for several things besides status: for instance, list the person that is responsible for this module.
- t prog** Specify a program *prog* to run whenever files in a module are tagged with **rtag**. *prog* runs with two arguments: the module name and the symbolic tag specified to **rtag**. There is no way to specify a program to run when **tag** is executed.
- u prog** Specify a program *prog* to run whenever **'cvs update'** is executed from the top-level directory of the checked-out module. *prog* runs with a single argument, the full path to the source repository for this module.

C.2 The **cvswrappers** file

Wrappers allow you to set a hook which transforms files on their way in and out of CVS.

The file **'cvswrappers'** defines the script that will be run on a file when its name matches a regular expression. There are two scripts that can be run on a file or directory. One script is executed on the file/directory before being checked into the repository (this is denoted with the **-t** flag) and the other when the file is checked out of the repository (this is denoted with the **-f** flag). The **'-t'/'-f'** feature does not work with client/server CVS.

The **'cvswrappers'** also has a **'-m'** option to specify the merge methodology that should be used when the file is updated. **MERGE** means the usual CVS behavior: try to merge the files (this generally will not work for binary files). **COPY** means that **cvs update** will merely copy one version over the other, and require the user using mechanisms outside CVS, to

insert any necessary changes. The `-m` wrapper option only affects behavior when merging is done on update; it does not affect how files are stored. See Chapter 17 [Binary files], page 77, for more on binary files.

The basic format of the file `'cvswrappers'` is:

```
wildcard      [option value][option value]...

where option is one of
-f            from cvs filter           value: path to filter
-t            to cvs filter             value: path to filter
-m            update methodology       value: MERGE or COPY
-k            keyword expansion        value: expansion mode

and value is a single-quote delimited value.

*.nib        -f 'unwrap %s' -t 'wrap %s %s' -m 'COPY'
*.c          -t 'indent %s %s'
```

The above example of a `'cvswrappers'` file states that all files/directories that end with a `.nib` should be filtered with the `'wrap'` program before checking the file into the repository. The file should be filtered though the `'unwrap'` program when the file is checked out of the repository. The `'cvswrappers'` file also states that a `COPY` methodology should be used when updating the files in the repository (that is no merging should be performed).

The last example line says that all files that end with a `*.c` should be filtered with `'indent'` before being checked into the repository. Unlike the previous example no filtering of the `*.c` file is done when it is checked out of the repository. The `-t` filter is called with two arguments, the first is the name of the file/directory to filter and the second is the pathname to where the resulting filtered file should be placed.

The `-f` filter is called with one argument, which is the name of the file to filter from. The end result of this filter will be a file in the users directory that they can work on as they normally would.

Note that the `'-t'/'-f'` features do not conveniently handle one portion of CVS's operation: determining when files are modified. CVS will still want a file (or directory) to exist, and it will use its modification time to determine whether a file is modified. If CVS erroneously thinks a file is unmodified (for example, a directory is unchanged but one of the files within it is changed), you can force it to check in the file anyway by specifying the `'-f'` option to `cvs commit` (see Section A.8.1 [commit options], page 97).

For another example, the following command imports a directory, treating files whose name ends in `.exe` as binary:

```
cvs import -I ! -W "*.exe -k 'b'" first-dir vendortag reltag
```

C.3 The commit support files

The `'-i'` flag in the `'modules'` file can be used to run a certain program whenever files are committed (see Section C.1 [modules], page 125). The files described in this section provide other, more flexible, ways to run programs whenever something is committed.

There are three kind of programs that can be run on commit. They are specified in files in the repository, as described below. The following table summarizes the file names and the purpose of the corresponding programs.

`'commitinfo'`

The program is responsible for checking that the commit is allowed. If it exits with a non-zero exit status the commit will be aborted.

`'verifymsg'`

The specified program is used to evaluate the log message, and possibly verify that it contains all required fields. This is most useful in combination with the `'rcsinfo'` file, which can hold a log message template (see Section C.8 [`rcsinfo`], page 133).

`'editinfo'`

The specified program is used to edit the log message, and possibly verify that it contains all required fields. This is most useful in combination with the `'rcsinfo'` file, which can hold a log message template (see Section C.8 [`rcsinfo`], page 133). (obsolete)

`'loginfo'`

The specified program is called when the commit is complete. It receives the log message and some additional information and can store the log message in a file, or mail it to appropriate persons, or maybe post it to a local newsgroup, or... Your imagination is the limit!

C.3.1 The common syntax

The administrative files such as `'commitinfo'`, `'loginfo'`, `'rcsinfo'`, `'verifymsg'`, etc., all have a common format. The purpose of the files are described later on. The common syntax is described here.

Each line contains the following:

- A regular expression. This is a basic regular expression in the syntax used by GNU emacs.
- A whitespace separator—one or more spaces and/or tabs.
- A file name or command-line template.

Blank lines are ignored. Lines that start with the character `'#'` are treated as comments. Long lines unfortunately can *not* be broken in two parts in any way.

The first regular expression that matches the current directory name in the repository is used. The rest of the line is used as a file name or command-line as appropriate.

C.4 Commitinfo

The `'commitinfo'` file defines programs to execute whenever `'cvs commit'` is about to execute. These programs are used for pre-commit checking to verify that the modified, added and removed files are really ready to be committed. This could be used, for instance, to verify that the changed files conform to to your site's standards for coding practice.

As mentioned earlier, each line in the `'commitinfo'` file consists of a regular expression and a command-line template. The template can include a program name and any number of arguments you wish to supply to it. The full path to the current source repository is appended to the template, followed by the file names of any files involved in the commit (added, removed, and modified files).

The first line with a regular expression matching the relative path to the module will be used. If the command returns a non-zero exit status the commit will be aborted.

If the repository name does not match any of the regular expressions in this file, the `'DEFAULT'` line is used, if it is specified.

All occurrences of the name `'ALL'` appearing as a regular expression are used in addition to the first matching regular expression or the name `'DEFAULT'`.

Note: when CVS is accessing a remote repository, `'commitinfo'` will be run on the *remote* (i.e., server) side, not the client side (see Section 3.9 [Remote repositories], page 21).

C.5 Verifying log messages

Once you have entered a log message, you can evaluate that message to check for specific content, such as a bug ID. Use the `'verifymsg'` file to specify a program that is used to verify the log message. This program could be a simple script that checks that the entered message contains the required fields.

The `'verifymsg'` file is often most useful together with the `'rcsinfo'` file, which can be used to specify a log message template.

Each line in the `'verifymsg'` file consists of a regular expression and a command-line template. The template must include a program name, and can include any number of arguments. The full path to the current log message template file is appended to the template.

One thing that should be noted is that the `'ALL'` keyword is not supported. If more than one matching line is found, the first one is used. This can be useful for specifying a default verification script in a module, and then overriding it in a subdirectory.

If the repository name does not match any of the regular expressions in this file, the `'DEFAULT'` line is used, if it is specified.

If the verification script exits with a non-zero exit status, the commit is aborted.

Note that the verification script cannot change the log message; it can merely accept it or reject it.

The following is a little silly example of a `'verifymsg'` file, together with the corresponding `'rcsinfo'` file, the log message template and an verification script. We begin with the log message template. We want to always record a bug-id number on the first line of the log message. The rest of log message is free text. The following template is found in the file `'/usr/cvssupport/tc.template'`.

```
BugId:
```

The script `'/usr/cvssupport/bugid.verify'` is used to evaluate the log message.

```
#!/bin/sh
#
```

```

#      bugid.verify filename
#
# Verify that the log message contains a valid bugid
# on the first line.
#
if head -1 < $1 | grep '^BugId:[ ]*[0-9][0-9]*$' > /dev/null; then
    exit 0
else
    echo "No BugId found."
    exit 1
fi

```

The `'verifymsg'` file contains this line:

```
^tc      /usr/cvssupport/bugid.edit
```

The `'rcsinfo'` file contains this line:

```
^tc      /usr/cvssupport/tc.template
```

C.6 Editinfo

NOTE: The `'editinfo'` feature has been rendered obsolete. To set a default editor for log messages use the `EDITOR` environment variable (see Appendix D [Environment variables], page 137) or the `-e` global option (see Section A.4 [Global options], page 86). See Section C.5 [verifymsg], page 129, for information on the use of the `'verifymsg'` feature for evaluating log messages.

If you want to make sure that all log messages look the same way, you can use the `'editinfo'` file to specify a program that is used to edit the log message. This program could be a custom-made editor that always enforces a certain style of the log message, or maybe a simple shell script that calls an editor, and checks that the entered message contains the required fields.

If no matching line is found in the `'editinfo'` file, the editor specified in the environment variable `$CVSEEDITOR` is used instead. If that variable is not set, then the environment variable `$EDITOR` is used instead. If that variable is not set a default will be used. See Section 2.2 [Committing your changes], page 9.

The `'editinfo'` file is often most useful together with the `'rcsinfo'` file, which can be used to specify a log message template.

Each line in the `'editinfo'` file consists of a regular expression and a command-line template. The template must include a program name, and can include any number of arguments. The full path to the current log message template file is appended to the template.

One thing that should be noted is that the `'ALL'` keyword is not supported. If more than one matching line is found, the first one is used. This can be useful for specifying a default edit script in a module, and then overriding it in a subdirectory.

If the repository name does not match any of the regular expressions in this file, the `'DEFAULT'` line is used, if it is specified.

If the edit script exits with a non-zero exit status, the commit is aborted.

Note: when CVS is accessing a remote repository, or when the `-m` or `-F` options to `cvcs commit` are used, `editinfo` will not be consulted. There is no good workaround for this; use `verifymsg` instead.

C.6.1 Editinfo example

The following is a little silly example of a `editinfo` file, together with the corresponding `rcsinfo` file, the log message template and an editor script. We begin with the log message template. We want to always record a bug-id number on the first line of the log message. The rest of log message is free text. The following template is found in the file `/usr/cvssupport/tc.template`.

BugId:

The script `/usr/cvssupport/bugid.edit` is used to edit the log message.

```
#!/bin/sh
#
#      bugid.edit filename
#
# Call $EDITOR on FILENAME, and verify that the
# resulting file contains a valid bugid on the first
# line.
if [ "x$EDITOR" = "x" ]; then EDITOR=vi; fi
if [ "x$CVSEEDITOR" = "x" ]; then CVSEEDITOR=$EDITOR; fi
$CVSEEDITOR $1
until head -1|grep '^BugId:[ ]*[0-9][0-9]*$' < $1
do echo -n "No BugId found. Edit again? ([y]/n)"
  read ans
  case ${ans} in
    n*) exit 1;;
  esac
$CVSEEDITOR $1
done
```

The `editinfo` file contains this line:

```
^tc      /usr/cvssupport/bugid.edit
```

The `rcsinfo` file contains this line:

```
^tc      /usr/cvssupport/tc.template
```

C.7 Logininfo

The `logininfo` file is used to control where `cvcs commit` log information is sent. The first entry on a line is a regular expression which is tested against the directory that the change is being made to, relative to the `$CVSROOT`. If a match is found, then the remainder of the line is a filter program that should expect log information on its standard input.

If the repository name does not match any of the regular expressions in this file, the `DEFAULT` line is used, if it is specified.

All occurrences of the name `ALL` appearing as a regular expression are used in addition to the first matching regular expression or `DEFAULT`.

The first matching regular expression is used.

See Section C.3 [commit files], page 127, for a description of the syntax of the ‘loginfo’ file.

The user may specify a format string as part of the filter. The string is composed of a ‘%’ followed by a space, or followed by a single format character, or followed by a set of format characters surrounded by ‘{’ and ‘}’ as separators. The format characters are:

s	file name
V	old version number (pre-checkin)
v	new version number (post-checkin)

All other characters that appear in a format string expand to an empty field (commas separating fields are still provided).

For example, some valid format strings are ‘%’, ‘%s’, ‘%{s}’, and ‘%{sVv}’.

The output will be a string of tokens separated by spaces. For backwards compatibility, the the first token will be the repository name. The rest of the tokens will be comma-delimited lists of the information requested in the format string. For example, if ‘/u/src/master’ is the repository, ‘%{sVv}’ is the format string, and three files (ChangeLog, Makefile, foo.c) were modified, the output might be:

```
/u/src/master ChangeLog,1.1,1.2 Makefile,1.3,1.4 foo.c,1.12,1.13
```

As another example, ‘%{ }’ means that only the name of the repository will be generated.

Note: when CVS is accessing a remote repository, ‘loginfo’ will be run on the *remote* (i.e., server) side, not the client side (see Section 3.9 [Remote repositories], page 21).

C.7.1 Loginfo example

The following ‘loginfo’ file, together with the tiny shell-script below, appends all log messages to the file ‘\$CVSROOT/CVSROOT/commitlog’, and any commits to the administrative files (inside the ‘CVSROOT’ directory) are also logged in ‘/usr/adm/cvsroot-log’.

```
ALL          /usr/local/bin/cvs-log $CVSROOT/CVSROOT/commitlog
^CVSROOT    /usr/local/bin/cvs-log /usr/adm/cvsroot-log
```

The shell-script ‘/usr/local/bin/cvs-log’ looks like this:

```
#!/bin/sh
(echo "-----";
 echo -n $USER" ";
 date;
 echo;
 sed '1s+${CVSROOT}'+) >> $1
```

C.7.2 Keeping a checked out copy

It is often useful to maintain a directory tree which contains files which correspond to the latest version in the repository. For example, other developers might want to refer to the latest sources without having to check them out, or you might be maintaining a web site with CVS and want every checkin to cause the files used by the web server to be updated.

The way to do this is by having `loginfo` invoke `cvs update`. Doing so in the naive way will cause a problem with locks, so the `cvs update` must be run in the background. Here is an example (this should all be on one line):

```
^cyclic-pages (date; cat; (sleep 2; cd /u/www/local-docs;
  cvs -q update -d) &) >> $CVSROOT/CVSROOT/update.log 2>&1
```

This will cause checkins to repository directories starting with `cyclic-pages` to update the checked out tree in `/u/www/local-docs`.

C.8 Rcsinfo

The `rcsinfo` file can be used to specify a form to edit when filling out the commit log. The `rcsinfo` file has a syntax similar to the `verifymsg`, `commitinfo` and `loginfo` files. See Section C.3.1 [syntax], page 128. Unlike the other files the second part is *not* a command-line template. Instead, the part after the regular expression should be a full pathname to a file containing the log message template.

If the repository name does not match any of the regular expressions in this file, the `DEFAULT` line is used, if it is specified.

All occurrences of the name `ALL` appearing as a regular expression are used in addition to the first matching regular expression or `DEFAULT`.

The log message template will be used as a default log message. If you specify a log message with `cvs commit -m message` or `cvs commit -f file` that log message will override the template.

See Section C.5 [verifymsg], page 129, for an example `rcsinfo` file.

When CVS is accessing a remote repository, the contents of `rcsinfo` at the time a directory is first checked out will specify a template which does not then change. If you edit `rcsinfo` or its templates, you may need to check out a new working directory.

C.9 Ignoring files via cvsignore

There are certain file names that frequently occur inside your working copy, but that you don't want to put under CVS control. Examples are all the object files that you get while you compile your sources. Normally, when you run `cvs update`, it prints a line for each file it encounters that it doesn't know about (see Section A.18.2 [update output], page 113).

CVS has a list of files (or `sh(1)` file name patterns) that it should ignore while running `update`, `import` and `release`. This list is constructed in the following way.

- The list is initialized to include certain file name patterns: names associated with CVS administration, or with other common source control systems; common names for patch files, object files, archive files, and editor backup files; and other names that are usually artifacts of assorted utilities. Currently, the default list of ignored file name patterns is:

```
RCS      SCCS      CVS      CVS.adm
RCSLOG   cvslog.*
tags     TAGS
.make.state      .nse_depinfo
```

```

*~      **      .##      ,*      _$*      *$
*.old   *.bak   *.BAK   *.orig  *.rej   .del-*
*.a     *.olb   *.o     *.obj   *.so    *.exe
*.Z     *.elc   *.ln
core

```

- The per-repository list in ‘`$CVSROOT/CVSROOT/cvsignore`’ is appended to the list, if that file exists.
- The per-user list in ‘`.cvsignore`’ in your home directory is appended to the list, if it exists.
- Any entries in the environment variable `$CVSIGNORE` is appended to the list.
- Any ‘`-I`’ options given to `cv`s is appended.
- As `cv`s traverses through your directories, the contents of any ‘`.cvsignore`’ will be appended to the list. The patterns found in ‘`.cvsignore`’ are only valid for the directory that contains them, not for any sub-directories.

In any of the 5 places listed above, a single exclamation mark (‘!’) clears the ignore list. This can be used if you want to store any file which normally is ignored by `cv`s.

Specifying ‘`-I !`’ to `cv`s `import` will import everything, which is generally what you want to do if you are importing files from a pristine distribution or any other source which is known to not contain any extraneous files. However, looking at the rules above you will see there is a fly in the ointment; if the distribution contains any ‘`.cvsignore`’ files, then the patterns from those files will be processed even if ‘`-I !`’ is specified. The only workaround is to remove the ‘`.cvsignore`’ files in order to do the import. Because this is awkward, in the future ‘`-I !`’ might be modified to override ‘`.cvsignore`’ files in each directory.

C.10 The history file

The file ‘`$CVSROOT/CVSROOT/history`’ is used to log information for the `history` command (see Section A.11 [history], page 101). This file must be created to turn on logging. This is done automatically if the `cv`s `init` command is used to set up the repository (see Section 3.6 [Creating a repository], page 20).

The file format of the ‘`history`’ file is documented only in comments in the `cv`s source code, but generally programs should use the `cv`s `history` command to access it anyway, in case the format changes with future releases of `cv`s.

C.11 Expansions in administrative files

Sometimes in writing an administrative file, you might want the file to be able to know various things based on environment `cv`s is running in. There are several mechanisms to do that.

To find the home directory of the user running `cv`s (from the `HOME` environment variable), use ‘`~`’ followed by ‘`/`’ or the end of the line. Likewise for the home directory of `user`, use ‘`~user`’. These variables are expanded on the server machine, and don’t get any reasonable expansion if `ps`erver (see Section 3.9.3 [Password authenticated], page 23) is in

used; therefore user variables (see below) may be a better choice to customize behavior based on the user running `CVS`.

One may want to know about various pieces of information internal to `CVS`. A `CVS` internal variable has the syntax `CVS_{variable}`, where *variable* starts with a letter and consists of alphanumeric characters and `_`. If the character following *variable* is a non-alphanumeric character other than `_`, the `{` and `}` can be omitted. The `CVS` internal variables are:

CVSROOT This is the value of the `CVS` root in use. See Chapter 3 [Repository], page 13, for a description of the various ways to specify this.

RCSBIN This is the value `CVS` is using for where to find RCS binaries. See Section A.4 [Global options], page 86, for a description of how to specify this.

CVSEEDITOR

VISUAL

EDITOR These all expand to the same value, which is the editor that `CVS` is using. See Section A.4 [Global options], page 86, for how to specify this.

USER Username of the user running `CVS` (on the `CVS` server machine).

If you want to pass a value to the administrative files which the user that is running `CVS` can specify, use a user variable. To expand a user variable, the administrative file contains `CVS_{=variable}`. To set a user variable, specify the global option `-s` to `CVS`, with argument *variable=value*. It may be particularly useful to specify this option via `.cvsrc` (see Section A.3 [`~/cvsrc`], page 86).

For example, if you want the administrative file to refer to a test directory you might create a user variable `TESTDIR`. Then if `CVS` is invoked as

```
CVS -s TESTDIR=/work/local/tests
```

and the administrative file contains `sh CVS_{=TESTDIR}/runtests`, then that string is expanded to `sh /work/local/tests/runtests`.

All other strings containing `$` are reserved; there is no way to quote a `$` character so that `$` represents itself.

Appendix D All environment variables which affect CVS

This is a complete list of all environment variables that affect CVS.

\$CVSIGNORE

A whitespace-separated list of file name patterns that CVS should ignore. See Section C.9 [cvsignore], page 133.

\$CVSWRAPPERS

A whitespace-separated list of file name patterns that CVS should treat as wrappers. See Section C.2 [Wrappers], page 126.

\$CVSREAD If this is set, `checkout` and `update` will try hard to make the files in your working directory read-only. When this is not set, the default behavior is to permit modification of your working files.

\$CVSUMASK

Controls permissions of files in the repository. See Section 3.2.2 [File permissions], page 15.

\$CVSROOT Should contain the full pathname to the root of the CVS source repository (where the RCS history files are kept). This information must be available to CVS for most commands to execute; if **\$CVSROOT** is not set, or if you wish to override it for one invocation, you can supply it on the command line: `'cvs -d cvsroot cvs_command...'` Once you have checked out a working directory, CVS stores the appropriate root (in the file `'CVS/Root'`), so normally you only need to worry about this when initially checking out a working directory.

\$EDITOR

\$CVSEEDITOR

Specifies the program to use for recording log messages during commit. **\$CVSEEDITOR** overrides **\$EDITOR**. See Section 2.2 [Committing your changes], page 9.

\$PATH If **\$RCSBIN** is not set, and no path is compiled into CVS, it will use **\$PATH** to try to find all programs it uses.

\$RCSBIN This is the value CVS is using for where to find RCS binaries. See Section A.4 [Global options], page 86, for a description of how to specify this. If not set, a compiled-in value is used, or your **\$PATH** is searched.

\$HOME

\$HOMEPATH

Used to locate the directory where the `'cvsrsrc'` file is searched (**\$HOMEPATH** is used for Windows-NT). see Section A.3 [`~/cvsrsrc`], page 86

\$CVS_RSH Specifies the external program which CVS connects with, when `:ext:` access method is specified. see Section 3.9.2 [Connecting via rsh], page 22.

\$CVS_SERVER

Used in client-server mode when accessing a remote repository using RSH. It specifies the name of the program to start on the server side when accessing

a remote repository using RSH. The default value is `cvs`. see Section 3.9.2 [Connecting via rsh], page 22

`$CVS_PASSFILE`

Used in client-server mode when accessing the `cvs login server`. Default value is `‘$HOME/.cvspass’`. see Section 3.9.3.2 [Password authentication client], page 25

`$CVS_CLIENT_PORT`

Used in client-server mode when accessing the server via Kerberos. see Section 3.9.4 [Kerberos authenticated], page 26

`$CVS_RCMD_PORT`

Used in client-server mode. If set, specifies the port number to be used when accessing the RCMD demon on the server side. (Currently not used for Unix clients).

`$CVS_CLIENT_LOG`

Used for debugging only in client-server mode. If set, everything send to the server is logged into `‘$CVS_CLIENT_LOG.in’` and everything send from the server is logged into `‘$CVS_CLIENT_LOG.out’`.

`$CVS_SERVER_SLEEP`

Used only for debugging the server side in client-server mode. If set, delays the start of the server child process the the specified amount of seconds so that you can attach to it with a debugger.

`$CVS_IGNORE_REMOTE_ROOT`

(What is the purpose of this variable?)

`$COMSPEC` Used under OS/2 only. It specifies the name of the command interpreter and defaults to `CMD.EXE`.

`$TMPDIR`

`$TMP`

`$TEMP`

Directory in which temporary files are located. Those parts of CVS which are implemented using RCS inspect the above variables in the order they appear above and the first value found is taken; if none of them are set, a host-dependent default is used, typically `‘/tmp’`. The CVS server uses `TMPDIR`. See Section A.4 [Global options], page 86, for a description of how to specify this. Some parts of CVS will always use `‘/tmp’` (via the `tmpnam` function provided by the system).

On Windows NT, `TMP` is used (via the `_tempnam` function provided by the system).

The `patch` program which is used by the CVS client uses `TMPDIR`, and if it is not set, uses `‘/tmp’` (at least with GNU patch 2.1).

CVS invokes RCS to perform certain operations. The following environment variables affect RCS. Note that if you are using the client/server CVS, these variables need to be set on the server side (which may or not may be possible depending on how you are connecting). There is probably not any need to set any of them, however.

\$LOGNAME

\$USER If set, they affect who RCS thinks you are. If you have trouble checking in files it might be because your login name differs from the setting of e.g. **\$LOGNAME**.

\$RCSINIT Options prepended to the argument list, separated by spaces. A backslash escapes spaces within an option. The **\$RCSINIT** options are prepended to the argument lists of most RCS commands.

Appendix E Troubleshooting

If you are having trouble with CVS, this appendix may help. If there is a particular error message which you are seeing, then you can look up the message alphabetically. If not, you can look through the section on other problems to see if your problem is mentioned there.

E.1 Partial list of error messages

Here is a partial list of error messages that you may see from CVS. It is not a complete list—CVS is capable of printing many, many error messages, often with parts of them supplied by the operating system, but the intention is to list the common and/or potentially confusing error messages.

The messages are alphabetical, but introductory text such as ‘**cv**s **update**: ’ is not considered in ordering them.

In some cases the list includes messages printed by old versions of CVS (partly because users may not be sure which version of CVS they are using at any particular moment).

cannot change permissions on temporary directory

Operation not permitted

This message has been happening in a non-reproducible, occasional way when we run the client/server test suite, both on Red Hat Linux 3.0.3 and 4.1. We haven’t been able to figure out what causes it, nor is it known whether it is specific to linux (or even to this particular machine!). If the problem does occur on other unices, ‘**Operation not permitted**’ would be likely to read ‘**Not owner**’ or whatever the system in question uses for the unix **EPERM** error. If you have any information to add, please let us know as described in [BUGS], page 2. If you experience this error while using CVS, retrying the operation which produced it should work fine.

cannot open CVS/Entries for reading: No such file or directory

This generally indicates a CVS internal error, and can be handled as with other CVS bugs (see [BUGS], page 2). Usually there is a workaround—the exact nature of which would depend on the situation but which hopefully could be figured out.

cvs [checkout aborted]: cannot rename file *file* to **CVS/**,, *file*: Invalid argument

This message has been reported as intermittently happening with CVS 1.9 on Solaris 2.5. The cause is unknown; if you know more about what causes it, let us know as described in [BUGS], page 2.

cvs [update aborted]: could not patch *file*: No such file or directory

This means that there was a problem finding the **patch** program. Make sure that it is in your **PATH**. Note that despite appearances the message is *not* referring to whether it can find *file*.

cvs **update**: could not patch *file*; will refetch

This means that for whatever reason the client was unable to apply a patch that the server sent. The message is nothing to be concerned about, because

inability to apply the patch only slows things down and has no effect on what CVS does.

dying gasps from server unexpected

This message seems to be caused by a hard-to-track-down bug in CVS or the systems it runs on (we don't know—we haven't tracked it down yet!). If you see it, you probably can just retry the operation which failed, or if you have discovered information concerning its cause, please let us know as described in [BUGS], page 2.

end of file from server (consult above messages if any)

The most common cause for this message is if you are using an external `rsh` program and it exited with an error. In this case the `rsh` program should have printed a message, which will appear before the above message. For more information on setting up a CVS client and server, see Section 3.9 [Remote repositories], page 21.

cvs commit: Executing 'mkmodules'

This means that your repository is set up for a version of CVS prior to CVS 1.8. When using CVS 1.8 or later, the above message will be preceded by

cvs commit: Rebuilding administrative file database

If you see both messages, the database is being rebuilt twice, which is unnecessary but harmless. If you wish to avoid the duplication, and you have no versions of CVS 1.7 or earlier in use, remove `-i mkmodules` every place it appears in your `modules` file. For more information on the `modules` file, see Section C.1 [modules], page 125.

rcs error: Unknown option: -x,v/

This message will be followed by a usage message for RCS. It means that you have an old version of RCS (probably supplied with your operating system). CVS only works with RCS version 5 and later.

cvs [server aborted]: received broken pipe signal

This message seems to be caused by a hard-to-track-down bug in CVS or the systems it runs on (we don't know—we haven't tracked it down yet!). It seems to happen only after a CVS command has completed, and you should be able to just ignore the message. However, if you have discovered information concerning its cause, please let us know as described in [BUGS], page 2.

cvs commit: Up-to-date check failed for 'file'

This means that someone else has committed a change to that file since the last time that you did a `cvs update`. So before proceeding with your `cvs commit` you need to `cvs update`. CVS will merge the changes that you made and the changes that the other person made. If it does not detect any conflicts it will report `'M cacErrCodes.h'` and you are ready to `cvs commit`. If it detects conflicts it will print a message saying so, will report `'C cacErrCodes.h'`, and you need to manually resolve the conflict. For more details on this process see Section 5.3 [Conflicts example], page 35.

Usage: `diff3 [-exEX3 [-i | -m] [-L label1 -L label3]] file1 file2 file3`

Only one of [exEX3] allowed

This indicates a problem with the installation of `diff3` and `rcsmerge`. Specifically `rcsmerge` was compiled to look for GNU `diff3`, but it is finding unix `diff3` instead. The exact text of the message will vary depending on the system. The solution is to make sure `rcsmerge` finds GNU `diff3`. Depending on how `rcsmerge` was compiled, it might be sufficient to place GNU `diff3` in your `PATH`, or it might be necessary to recompile `rcsmerge` or find a binary distribution of `rcsmerge` which looks in the `PATH`.

`cvs commit: warning: editor session failed`

This means that the editor which `cvs` is using exits with a nonzero exit status. Some versions of `vi` will do this even when there was not a problem editing the file. If so, point the `CVSEEDITOR` environment variable to a small script such as:

```
#!/bin/sh
vi $*
exit 0
```

E.2 Other common problems

Here is a list of problems which cannot be readily looked up based on an error message. They are in no particular order.

- If `cvs update` finds a conflict and tries to merge, as described in Section 5.3 [Conflicts example], page 35, but doesn't tell you there were conflicts, then you may have an old version of RCS. For more information on how to set this up, see the 'INSTALL' file in the `CVS` source distribution.

Appendix F GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright © 1989, 1991 Free Software Foundation, Inc.
59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING,

DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program’s source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions

for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore,

by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a

version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

one line to give the program's name and a brief idea of what it does.
 Copyright (C) 19yy *name of author*

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) 19yy *name of author*
 Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
 type 'show w'.
 This is free software, and you are welcome to redistribute it
 under certain conditions; type 'show c' for details.

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ‘show w’ and ‘show c’; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program
 ‘Gnomovision’ (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989
 Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Index

(Index is nonexistent)

Short Contents

About this manual	1
1 What is CVS?	5
2 A sample session	9
3 The Repository	13
4 Starting a project with CVS	29
5 Multiple developers	33
6 Revisions and branches	43
7 Merging	51
8 Recursive behavior	55
9 Adding files to a directory	57
10 Removing files	59
11 Removing directories	61
12 Tracking third-party sources	63
13 Moving and renaming files	67
14 Moving and renaming directories	69
15 History browsing	71
16 Keyword substitution	73
17 Handling binary files	77
18 How your build system interacts with CVS	79
19 Compatibility between CVS Versions	81
20 Revision management	83
Appendix A Guide to CVS commands	85
Appendix B Quick reference to CVS commands	115
Appendix C Reference manual for Administrative files	125
Appendix D All environment variables which affect CVS	137
Appendix E Troubleshooting	141
Appendix F GNU GENERAL PUBLIC LICENSE	145
Index	153

Table of Contents

About this manual	1
Checklist for the impatient reader	1
Credits	1
BUGS	2
1 What is CVS?	5
CVS is not	5
2 A sample session	9
2.1 Getting the source	9
2.2 Committing your changes	9
2.3 Cleaning up	10
2.4 Viewing differences	10
3 The Repository	13
3.1 Telling CVS where your repository is	13
3.2 How data is stored in the repository	14
3.2.1 Where files are stored within the repository	14
3.2.2 File permissions	15
3.2.3 The attic	16
3.3 How data is stored in the working directory	16
3.4 The administrative files	19
3.4.1 Editing administrative files	19
3.5 Multiple repositories	20
3.6 Creating a repository	20
3.7 Backing up a repository	21
3.8 Moving a repository	21
3.9 Remote repositories	21
3.9.1 Server requirements	22
3.9.2 Connecting with rsh	22
3.9.3 Direct connection with password authentication ..	23
3.9.3.1 Setting up the server for password authentication	23
3.9.3.2 Using the client with password authentication	25
3.9.3.3 Security considerations with password authentication	25
3.9.4 Direct connection with kerberos	26
3.10 Read-only repository access	26
3.11 Temporary directories for the server	27

4	Starting a project with CVS	29
4.1	Setting up the files	29
4.1.1	Creating a directory tree from a number of files . .	29
4.1.2	Creating Files From Other Version Control Systems	30
4.1.3	Creating a directory tree from scratch	30
4.2	Defining the module	30
5	Multiple developers	33
5.1	File status	33
5.2	Bringing a file up to date	34
5.3	Conflicts example	35
5.4	Informing others about commits	37
5.5	Several developers simultaneously attempting to run CVS . .	37
5.6	Mechanisms to track who is editing files	38
5.6.1	Telling CVS to watch certain files	38
5.6.2	Telling CVS to notify you	39
5.6.3	How to edit a file which is being watched	40
5.6.4	Information about who is watching and editing . .	41
5.6.5	Using watches with old versions of CVS	41
5.7	Choosing between reserved or unreserved checkouts	41
6	Revisions and branches	43
6.1	Revision numbers	43
6.2	Versions, revisions and releases	44
6.3	Assigning revisions	44
6.4	Tags—Symbolic revisions	45
6.5	What branches are good for	47
6.6	Creating a branch	47
6.7	Sticky tags	48
6.8	Magic branch numbers	50
7	Merging	51
7.1	Merging an entire branch	51
7.2	Merging from a branch several times	51
7.3	Merging differences between any two revisions	52
7.4	Merging can add or remove files	53
8	Recursive behavior	55
9	Adding files to a directory	57
10	Removing files	59
11	Removing directories	61

12	Tracking third-party sources	63
12.1	Importing a module for the first time	63
12.2	Updating a module with the import command	63
12.3	Reverting to the latest vendor release	64
12.4	How to handle binary files with cvs import	64
12.5	How to handle keyword substitution with cvs import	64
12.6	Multiple vendor branches	65
13	Moving and renaming files	67
13.1	The Normal way to Rename	67
13.2	Moving the history file	67
13.3	Copying the history file	68
14	Moving and renaming directories	69
15	History browsing	71
15.1	Log messages	71
15.2	The history database	71
15.3	User-defined logging	71
15.4	Annotate command	71
16	Keyword substitution	73
16.1	RCS Keywords	73
16.2	Using keywords	74
16.3	Avoiding substitution	74
16.4	Substitution modes	75
16.5	Problems with the \$Log\$ keyword.	75
17	Handling binary files	77
18	How your build system interacts with CVS	
	79
19	Compatibility between CVS Versions	81
20	Revision management	83
20.1	When to commit?	83

Appendix A Guide to CVS commands 85

A.1	Overall structure of CVS commands	85
A.2	CVS's exit status	85
A.3	Default options and the <code>~/cvsrc</code> file	86
A.4	Global options	86
A.5	Common command options	88
A.6	admin—Administration front end for <code>rcs</code>	90
	A.6.1 admin options	91
	A.6.2 admin examples	93
	A.6.2.1 Outdating is dangerous	93
	A.6.2.2 Comment leaders	94
A.7	checkout—Check out sources for editing	94
	A.7.1 checkout options	95
	A.7.2 checkout examples	96
A.8	commit—Check files into the repository	96
	A.8.1 commit options	97
	A.8.2 commit examples	97
	A.8.2.1 Committing to a branch	97
	A.8.2.2 Creating the branch after editing	98
A.9	diff—Show differences between revisions	98
	A.9.1 diff options	99
	A.9.2 diff examples	99
A.10	export—Export sources from CVS, similar to checkout	100
	A.10.1 export options	100
A.11	history—Show status of files and users	101
	A.11.1 history options	101
A.12	import—Import sources into CVS, using vendor branches	102
	A.12.1 import options	103
	A.12.2 import output	104
	A.12.3 import examples	104
A.13	log—Print out log information for files	104
	A.13.1 log options	104
	A.13.2 log examples	106
A.14	rdiff—'patch' format diffs between releases	106
	A.14.1 rdiff options	106
	A.14.2 rdiff examples	107
A.15	release—Indicate that a Module is no longer in use	107
	A.15.1 release options	108
	A.15.2 release output	108
	A.15.3 release examples	108
A.16	rtag—Add a symbolic tag to a module	109
	A.16.1 rtag options	109
A.17	tag—Add a symbolic tag to checked out versions of files	110
	A.17.1 tag options	110
A.18	update—Bring work tree in sync with repository	111
	A.18.1 update options	111

A.18.2	update output	112
Appendix B	Quick reference to CVS commands	
.....	115
Appendix C	Reference manual for	
Administrative files	125
C.1	The modules file	125
C.2	The cvswrappers file	126
C.3	The commit support files	127
C.3.1	The common syntax	128
C.4	Commitinfo	128
C.5	Verifying log messages	129
C.6	Editinfo	130
C.6.1	Editinfo example	131
C.7	Logininfo	131
C.7.1	Logininfo example	132
C.7.2	Keeping a checked out copy	132
C.8	Rcsinfo	133
C.9	Ignoring files via cvsignore	133
C.10	The history file	134
C.11	Expansions in administrative files	134
Appendix D	All environment variables which	
affect CVS	137
Appendix E	Troubleshooting	141
E.1	Partial list of error messages	141
E.2	Other common problems	143
Appendix F	GNU GENERAL PUBLIC	
LICENSE	145
Preamble	145
How to Apply These Terms to Your New Programs	150
Index	153

