

---

# *Perforce 2002.1 User's Guide*

**April 2002**

---

---

This manual copyright 1997-2002 Perforce Software.

All rights reserved.

Perforce software and documentation is available from <http://www.perforce.com>. You may download and use Perforce programs, but you may not sell or redistribute them. You may download, print, copy, edit, and redistribute the documentation, but you may not sell it, or sell any documentation derived from it. You may not modify or attempt to reverse engineer the programs.

Perforce programs and documents are available from our Web site as is. No warranty or support is provided. Warranties and support, along with higher capacity servers, are sold by Perforce Software.

Perforce Software assumes no responsibility or liability for any errors or inaccuracies that may appear in this book.

By downloading and using our programs and documents you agree to these terms.

Perforce and Inter-File Branching are trademarks of Perforce Software. Perforce software includes software developed by the University of California, Berkeley and its contributors.

All other brands or product names are trademarks or registered trademarks of their respective companies or organizations.

---

---

# Table of Contents

---

<b>Preface</b>	<b>About This Manual .....</b>	<b>9</b>
	Administering Perforce?.....	9
	Please Give Us Feedback .....	9
<b>Chapter 1</b>	<b>Perforce Concepts .....</b>	<b>11</b>
	Perforce Architecture .....	12
	Moving Files Between the Clients and the Server .....	12
	File Conflicts .....	13
	Labeling Groups of Files.....	13
	Branching Files.....	14
	Job Tracking.....	14
	Change Review, Daemons, and Triggers.....	14
	Protections .....	15
<b>Chapter 2</b>	<b>Connecting to the Perforce Server.....</b>	<b>17</b>
	Before you begin .....	17
	Setting up your environment to use Perforce.....	17
	Telling Perforce clients where the server is .....	18
	Verifying the connection to the Perforce server .....	19
<b>Chapter 3</b>	<b>Perforce Basics: Quick Start .....</b>	<b>21</b>
	Underlying concepts .....	21
	File configurations used in the examples.....	21
	Setting up a client workspace .....	22
	Naming the client workspace .....	22
	Describing the client workspace to the Perforce server .....	22
	Copying depot files into your workspace.....	24
	Updating the depot with files from your workspace .....	25
	Adding files to the depot.....	25

Editing files in the depot ..... 27  
 Deleting files from the depot ..... 28  
 Submitting with multiple operations ..... 29  
 Backing out: reverting files to their unopened states ..... 30  
 Basic reporting commands ..... 30

**Chapter 4**    **Perforce Basics:**  
**The Details..... 33**

Description of the Client Workspace..... 33  
 Wildcards ..... 34  
     Wildcards and “p4 add” ..... 34  
 Mapping the Depot to the Client Workspace ..... 34  
     Multiple depots..... 35  
     Using views ..... 35  
     Wildcards in views ..... 36  
     Types of mappings ..... 37  
 Editing Existing Client Specifications ..... 40  
     Deleting an existing client specification ..... 40  
     Client specification options ..... 41  
     Line-ending conventions (CR/LF translation) ..... 42  
 Referring to Files on the Command Line ..... 43  
     Local syntax ..... 43  
     Perforce syntax ..... 43  
     Providing files as arguments to commands ..... 44  
     Wildcards and Perforce syntax ..... 45  
 Name and String Limitations ..... 45  
     Illegal characters in filenames and Perforce objects ..... 45  
     Name and description lengths ..... 46  
 Specifying Older File Revisions ..... 46  
     Using revision specifications without filenames ..... 48  
 Revision Ranges ..... 49  
 File Types ..... 50  
 Forms and Perforce Commands ..... 55  
     Reading forms from standard input;  
     Writing forms to standard output ..... 55  
 General Reporting Commands ..... 56

<b>Chapter 5</b>	<b>Perforce Basics: Resolving File Conflicts.....</b>	<b>57</b>
	RCS Format: How Perforce Stores File Revisions.....	57
	Only the differences between revisions are stored.....	57
	Use of “diff” to determine file revision differences.....	59
	Scheduling Resolves of Conflicting Files .....	59
	Why “p4 sync” to Schedule a Resolve? .....	59
	How Do I Know When a Resolve is Needed? .....	60
	Performing Resolves of Conflicting Files.....	60
	File revisions used and generated by “p4 resolve” .....	61
	Types of conflicts between file revisions.....	61
	How the merge file is generated .....	62
	The “p4 resolve” options.....	62
	Using Flags with Resolve to Automatically Accept Particular Revisions.....	65
	Binary files and “p4 resolve” .....	66
	Locking Files to Minimize File Conflicts.....	66
	Preventing multiple resolves with p4 lock .....	67
	Preventing multiple checkouts with +l files.....	67
	Resolves and Branching.....	68
	Resolve Reporting.....	68
<b>Chapter 6</b>	<b>Perforce Basics: Miscellaneous Topics.....</b>	<b>71</b>
	Reconfiguring the Perforce Environment with \$P4CONFIG.....	71
	Perforce Passwords .....	72
	Command-Line Flags Common to All Perforce Commands.....	73
	Working Detached.....	74
	Finding changed files with “p4 diff” .....	75
	Using “p4 diff” to update the depot.....	75
	Refreshing files.....	75
	Recommendations for Organizing the Depot .....	76
	Renaming Files.....	77
	Revision histories and renamed files.....	77

<b>Chapter 7</b>	<b>Changelists .....</b>	<b>79</b>
	Working with the Default Changelist .....	80
	Creating Numbered Changelists Manually .....	81
	Working With Numbered Changelists .....	81
	Automatic Creation and Renumbering of Changelists .....	82
	When submit of the default changelist fails, the changelist is assigned a number .....	82
	Perforce May Renumber a Changelist upon Submission .....	83
	Deleting Changelists.....	83
	Changelist Reporting.....	84
<b>Chapter 8</b>	<b>Labels .....</b>	<b>85</b>
	Why Not Just Use Changelist Numbers? .....	85
	Creating a Label .....	85
	Adding and Changing Files Listed in a Label .....	86
	Previewing labelsync's results.....	88
	Preventing Accidental Overwrites of a Label's Contents .....	88
	Retrieving a Label's Contents into a Client Workspace .....	89
	Deleting Labels.....	89
	Label Reporting.....	90
<b>Chapter 9</b>	<b>Branching .....</b>	<b>91</b>
	What is Branching? .....	91
	When to Create a Branch .....	91
	Perforce's Branching Mechanisms: Introduction .....	92
	Branching and Merging, Method 1:	
	Branching with File Specifications.....	93
	Creating branched files.....	93
	Propagating changes between branched files .....	94
	Propagating changes from branched files to the original files.....	95
	Branching and Merging, Method 2:	
	Branching with Branch Specifications.....	95
	Branch Specification Usage Notes .....	97
	Integration Usage Notes .....	97
	Deleting Branches .....	99
	Advanced Integration Functions.....	99

Integrating specific file revisions .....	99
Re-integrating and re-resolving files .....	100
How Integrate Works .....	100
The yours, theirs, and base files .....	100
The integration algorithm .....	100
Integrate's actions.....	101
Integration Reporting.....	102
For More Information .....	102
<b>Chapter 10</b>	<b>Job Tracking .....</b>
	<b>103</b>
Job Usage Overview.....	103
Creating and editing jobs using the default job specification.....	104
Creating and editing jobs with custom job specifications .....	105
Viewing jobs by content with jobviews.....	106
Finding jobs containing particular words.....	107
Finding jobs by field values .....	107
Using and escaping wildcards in jobviews .....	107
Negating the sense of a query .....	108
Using dates in jobviews.....	108
Comparison operators and field types.....	108
Linking Jobs to Changelists.....	110
Automatically linking jobs to changelists with the p4 user form .....	110
Automatic update of job status .....	110
Manually associating jobs with changelists .....	111
What if there's no status field? .....	112
Deleting Jobs.....	112
Integrating with External Defect Tracking Systems .....	112
Job Reporting Commands .....	113
<b>Chapter 11</b>	<b>Reporting and Data Mining.....</b>
	<b>115</b>
Files.....	115
File metadata.....	115
Relationships between client and depot files .....	117
File contents .....	118
Changelists .....	120
Viewing changelists that meet particular criteria .....	120
Files and jobs affected by changelists.....	122

Labels.....	123
Branch and Integration Reporting.....	123
Job Reporting.....	124
Basic job information .....	124
Jobs, fixes, and changelists .....	125
Reporting for Daemons.....	126
System Configuration.....	126
Special Reporting Flags.....	127
Reporting with Scripting .....	127
Comparing the change content of two file sets.....	127
<b>Appendix A Installing Perforce .....</b>	<b>129</b>
Getting Perforce .....	129
Installing Perforce on UNIX.....	129
Download the files and make them executable .....	130
Creating a Perforce server root directory.....	130
Telling the Perforce server which port to listen to.....	130
Starting the Perforce server.....	131
Stopping the Perforce server.....	131
Telling Perforce clients which port to talk to.....	131
Installing Perforce on Windows.....	131
Terminology note: Windows services and servers .....	132
Starting and stopping Perforce on Windows .....	132
<b>Appendix B Environment Variables .....</b>	<b>135</b>
Setting and viewing environment variables.....	136
<b>Appendix C Glossary .....</b>	<b>137</b>
<b>Index.....</b>	<b>147</b>



# About This Manual

---

This is the *Perforce 2002.1 User's Guide*. It teaches the use of Perforce's command-line interface; other Perforce clients such as P4Win (the Perforce Windows client) are not discussed here. If you'd like documentation on other Perforce clients, please see our documentation pages, available from our web site at <http://www.perforce.com>.

Although you can use this guide as a reference manual, we intend it primarily as a guide/tutorial on using Perforce. The full syntax of most of the Perforce commands is explicitly not provided here; in particular, only a subset of the available flags are mentioned. For a complete guide to Perforce, please see the *Perforce Command Reference*, or the on-line help system. If you will be using Perforce on any operating system other than UNIX, please consult the Perforce platform notes for that OS.

Chapters 2 through 4 of this manual comprise our *Getting Started* guide. Newcomers to Perforce should start there, and move to subsequent chapters as needed.

## Administering Perforce?

---

If you're administering a Perforce server, you'll need the *Perforce System Administrator's Guide*, which contains all the system administration material formerly found in this manual. If you're installing Perforce, the *Perforce System Administrator's Guide* is also the place to start.

## Please Give Us Feedback

---

We are interested in receiving opinions on it from our users. In particular, we'd like to hear from users who have never used Perforce before. Does this guide teach the topic well? Please let us know what you think; we can be reached at [manual@perforce.com](mailto:manual@perforce.com).



Perforce facilitates the sharing of files among multiple users. It is a software configuration management tool, but software configuration management (SCM) has been defined in many different ways, depending on who's giving the definition. SCM has been described as providing version control, file sharing, release management, defect tracking, build management, and a few other things. It's worth looking at exactly what Perforce does and doesn't do:

- Perforce offers *version control*: multiple revisions of the same file are stored, and older revisions are always accessible.
- Perforce provides facilities for *concurrent development*; multiple users can edit their own copies of the same file.
- Some *release management* facilities are offered; Perforce will track which revisions of which files are part of a particular release.
- Bugs and system improvement requests can be tracked from entry to fix; this is known as *defect tracking*.
- Perforce supplies some *lifecycle management* functionality; files can be kept in release branches, development branches, or in any sort of needed file set.
- *Change review* functionality is provided by Perforce; this allows users to be notified by email when particular files are changed.
- Although a build management tool is not built into Perforce, we do offer a companion freeware product called `Jam`. The `Jam` tool and Perforce meet at the file system; source files managed by Perforce are easily built by `Jam`.

Although Perforce was built to manage source files, it can manage any sort of on-line documents. It can be used to store revisions of a manual, to manage Web pages, or to store old versions of operating system administration files.

**Note** | You don't need to read the rest of this chapter if you don't want to. All the material discussed here is also covered in the "how-to" chapters, which comprise the rest of the manual. This chapter is provided as a guide to what Perforce does, without the details of how to do it.

## Perforce Architecture

---

Perforce has a client/server architecture, in which many computers, called *clients*, are connected to one central machine, the *server*. Each user works on a client; at their command, files they've been editing are transferred to and from the server. The clients communicate with the server via TCP/IP.

The Perforce clients may be distributed around a local area network, wide area network, dialup network, or any combination of these. There can also be Perforce clients on the same host as the server.

Two programs do the bulk of Perforce's work:

- The `p4d` program runs on the Perforce server. It manages the shared file repository, and keeps track of users, clients, protections, and other Perforce metadata.

The `p4d` program must be run on a UNIX or Windows host.

- The `p4` program runs on each Perforce client. It sends the users' requests to the `p4d` server program for processing, and communicates with `p4d` via TCP/IP.

Perforce client programs can be run on many platforms, including UNIX, Linux, Windows, VMS, Macintosh, BeOS, and NeXT hosts.

This manual assumes that either you or your system administrator has already installed both `p4` and `p4d`. You'll find instructions for this in the *Perforce System Administrator's Guide*, available at our Web site.

## Moving Files Between the Clients and the Server

---

Users create, edit, and delete files in their own directories on the clients; these directories are called *client workspaces*. Perforce commands are used to move files to and from a shared file repository on the server known as the *depot*. Perforce users can retrieve files from the depot into their own client workspaces, where they can be read, edited, and resubmitted to the depot for other users to access. When a new *revision* of a file is stored in the depot, the old revisions are kept, and are still accessible.

Files that have been edited within a client workspace are sent to the depot via a *changelist*, which is a list of files and instructions that tell the depot what to do with those files. For example, one file might have been changed in the client workspace, another added, and another deleted. These file changes might be sent to the depot in a single changelist, which is processed *atomically*: either all the changes are made to the depot at once, or none of them are. This allows problem fixes that span multiple files to be updated in the depot at exactly the same time.

Each client workspace has its own *client view*, which determines which files in the depot can be accessed by that client workspace. One client workspace might be able to access all

the files in the depot, while another client workspace might access only a single file. The Perforce server is responsible for tracking the state of the client workspace; Perforce knows which files a client workspace has, where they are, and which files have write permission turned on.

See Chapter 3, Perforce Basics: Quick Start and Chapter 4, Perforce Basics: The Details for basic information about using Perforce.

## File Conflicts

---

When two users edit the same file, it is possible for their changes to conflict. For example, suppose two users copy the same file from the depot into their workspaces, and each edits his copy of the file in different ways. The first user sends his version of the file back to the depot, and then the second user tries to do the same thing. If Perforce were to unquestioningly accept the second user's file into the depot, the first user's changes would not be included in the latest revision of the file (known as the *head revision*).

When a file conflict is detected, Perforce allows the user experiencing the conflict to perform a *resolve* of the conflicting files. The resolve process allows the user to decide what needs to be done: should his file overwrite the other user's? Should his own file be thrown away? Or should the two conflicting files be merged into one? At the user's request, Perforce will perform a *three-way merge* between the two conflicting files and the single file that both were based on. This process generates a *merge* file from the conflicting files, which contains all the changes from both conflicting versions. This file can be edited and then submitted to the depot.

To learn how to resolve file conflicts, see Chapter 5, Perforce Basics: Resolving File Conflicts.

## Labeling Groups of Files

---

It is often useful to mark a particular set of file revisions for later access. For examples, the release engineers might want to keep a list of all the file revisions that comprise a particular release of their program. This list of files can be assigned a single mnemonic name, like `release2.0.1`; this name is a *label* for the user-determined list of files. At any subsequent time, the label can be used to copy the old file revisions into a client workspace.

For details about labels, see Chapter 8, Labels.

## Branching Files

---

Thus far, it has been assumed that all changes of files happen linearly, but this is not always the case. Suppose that one source file needs to evolve in two separate directions; perhaps one set of upcoming changes will allow the program to run under VMS, and another set will make it a Mac program. Clearly, two separately evolving copies of the same files are necessary.

Perforce's *Inter-File Branching*<sup>™</sup> mechanism allows any set of files to be copied within the depot. By default, the new file set, or *codeline*, evolves separately from the original files, but changes in either codeline can be propagated to the other.

For details about branching, see Chapter 9, Branching.

## Job Tracking

---

A *Job* is a generic term for a plain-text description of some change that needs to be made to the source code. A job might be a bug description, like “the system crashes when I press return”, or it might be a system improvement request, like “please make the program run faster.”

Whereas a job represents work that is intended to be performed, a changelist represents work actually done. Perforce's job tracking mechanism allows jobs to be linked to the changelists that implement the work requested by the job. A job can later be looked up to determine if and when it was fixed, which file revisions implemented the fix, and who fixed it. The fields contained in your system's jobs can be defined by the Perforce system administrator.

Perforce's job tracking mechanism does not implement all functionality normally supplied by full-scale defect tracking systems. Its simple functionality can be used as is, or it can be integrated with third-party job tracking systems through P4DTI - Perforce Defect Tracking and Integration.

To read more about jobs, please see Chapter 10, Job Tracking.

## Change Review, Daemons, and Triggers

---

Perforce's *change review* mechanism allows users to receive email notifying them when particular files have been updated in the depot. The files that a particular user receives notification on is determined by that user. Change review is implemented by an external program, or *daemon*, which can be modified by a knowledgeable user, allowing change review functionality to be customized.

Perforce can be made to run external scripts whenever changelists are submitted. These scripts, called *triggers*, allow files to be validated before they're submitted to the depot.

The change review daemon is not covered in this guide. To learn how to set up the change review daemon, integrate Perforce with third-party defect tracking systems, or write your own daemons, you'll need to consult the *Perforce System Administrator's Guide*.

## Protections

---

Perforce provides a protection scheme to prevent unauthorized or inadvertent access to the depot. The protection mechanism determines exactly which Perforce commands are allowed to be run by any particular client.

Permissions can be granted or denied based on users' usernames and IP addresses, or can be granted or denied to entire groups of users. Since Perforce usernames are easily changed, protections at the user level provide safety, not security. Protections at the IP address level are as secure as the host itself.

We discuss protections in the *Perforce System Administrator's Guide*.





# Connecting to the Perforce Server

---

Perforce uses a client/server architecture. Files are created and edited by users on their own client hosts; these files are transferred to and from a shared file repository located on a Perforce server. Every running Perforce system uses a single server and can have many clients.

As mentioned earlier, two programs do the bulk of Perforce's work:

- The `p4d` program runs on the Perforce server. It manages the shared file repository, and keeps track of users, clients, protections, and other Perforce metadata.
- The `p4` program runs on each Perforce client. It sends the users' requests to the `p4d` server program for processing, and communicates with `p4d` via TCP/IP.

Each Perforce client program needs to know the address and port of the Perforce server with which it communicates. This address is stored in the `P4PORT` environment variable.

## Before you begin

---

This chapter assumes that your system administrator has already set up a Perforce server (`p4d`) for you, and that it is already up and running.

If this is not the case (for instance, if you're installing Perforce from scratch), you'll also have to install the Perforce server before continuing. See the appendix, "Installing Perforce" on page 129, for information on how to install the server.

The information in the appendix is intended to help you install a server for evaluation purposes. If you're installing a production server, or are planning on extensive testing of your evaluation server, we strongly encourage you to read the full installation instructions in the *Perforce System Administrator's Guide*.

## Setting up your environment to use Perforce

---

A Perforce client program needs to know two things in order to talk to a Perforce server:

- the name of the host on which `p4d` is running, and
- the port on which `p4d` is listening

These are set via a single environment variable, `P4PORT`.

**Note** | See “Setting and viewing environment variables” on page 136 for information about how to set environment variables for most operating systems and shells.

If your site is already using Perforce, it’s possible that your system administrator has already set `P4PORT` for you; if not, you’ll need to set it yourself.

Either way, after setting `P4PORT` to point to your server, you should verify your client’s connection to the Perforce server with the `p4 info` command.

### Telling Perforce clients where the server is

To use Perforce, you’ll need to know the name of the host where `p4d` is located, and the number of the TCP/IP port on which it’s listening.

- If you’ve just installed the Perforce server yourself, you already know this, having configured the server on a specific host to listen to a specific port.
- If you’re connecting to an existing Perforce installation, you’ll have to ask your system administrator for the host and port of the Perforce server.

Once you’ve obtained the host and port information, set your `P4PORT` environment variable to `host:portNum`, where `host` is the name of the host on which `p4d` is running, and `portNum` is the port to which it is listening. For example:

If the server is running on...	and is listening to port...	set <code>P4PORT</code> to:
dogs	3435	dogs:3435
x.com	1818	x.com:1818

The definition of `P4PORT` can be shortened if the Perforce client is running on the same host as the server. In this case, only the port number need be provided to `p4`. If `p4d` is running on a host named or aliased `perforce`, listening on port 1666, the definition of `P4PORT` for the `p4` client can be dispensed with altogether. For example:

If the server is running on...	and is listening to port...	set <code>P4PORT</code> to:
<same host as the p4 client>	1543	1543
perforce	1666	<no value needed>

When `P4PORT` has been set, you should re-verify the connection with `p4 info`, as described below. Once this has been done, Perforce is ready to use.

## Verifying the connection to the Perforce server

To verify the connection, type `p4 info` at the command line. If the `P4PORT` environment variable is correctly set, you'll see something like this:

```
User name: edk
Client name: wrkstn12
Client host: wrkstn12
Client unknown.
Current directory: /usr/edk
Client address: 192.168.0.123:1818
Server address: p4server:1818
Server root: /usr/depot/p4d
Server date: 2000/07/28 12:11:47 PDT
Server version: P4D/FREEBSD/2000.1/16375 (2000/07/25)
Server license: P4 Admin <p4adm> 20 users on unix (expires 2001/01/01)
```

The `Server address:` field shows which Perforce server to which the client has connected; it displays the host and port number on which `p4d` is listening.

In the above example, everything is fine. If, however, you receive a variant of this message:

```
Perforce client error:
    Connect to server failed; check $P4PORT.
    TCP connect to perforce:1666 failed.
    perforce: host unknown.
```

then `P4PORT` has not been correctly set. If the value you see in the third line of the error message is `perforce:1666` (as above), then `P4PORT` has not been set at all; if the value is anything else, `P4PORT` has been incorrectly set. In either case, you'll need to set the value of `P4PORT`.

**Windows** | On Windows platforms, registry variables are preferred over environment variables, and you can set these with command `p4 set`.



---

## Chapter 3 **Perforce Basics: Quick Start**

---

This chapter teaches basic Perforce usage. You'll learn how to set up a workspace, populate it with files from the common file repository (the *depot*), edit these files and submit the changes back to the repository, back out of any unwanted changes, and the basic Perforce reporting commands.

This chapter gives a broad overview of these concepts and commands; for details, see Chapter 4, Perforce Basics: The Details.

### Underlying concepts

---

The basic ideas behind Perforce are quite simple: files are created, edited, and deleted in the user's own directories, which are called *client workspaces*. Perforce commands are used to move files to and from a shared file repository known as the *depot*. Perforce users can retrieve files from the depot into their own client workspaces, where they can be read, edited, and resubmitted to the depot for other users to access. When a new revision of a file is stored in the depot, the old revisions are kept and remain accessible.

Perforce was written to be as unobtrusive as possible, so that very few changes to your normal work habits are required. Files are still created in your own directories with your tool of choice; Perforce commands supplement your normal work actions instead of replacing them.

Perforce commands are always entered in the form `p4 command [arguments]`.

### File configurations used in the examples

This manual makes extensive use of examples based on the source code set for a program called Elm. The Elm examples used in this manual are set up as follows:

- A single depot is used to store the Elm files, and perhaps other projects as well. The elm files will be shared by storing them under an elm subdirectory within the depot.
- Each user will store his or her client workspace Elm files in a different subdirectory. The two users we'll be following most closely, Ed and Lisa, will work with their Elm files in the following locations:

User	Username	Client Workspace Name	Top of own Elm File Tree
Ed	edk	eds_elm	/usr/edk/elm
Lisa	lisag	lisas_ws	/usr/lisag/docs

## Setting up a client workspace

---

To move files between a client workspace and the depot, the Perforce server requires two pieces of information:

- A name that uniquely identifies the client workspace, and
- The top-level directory of this workspace.

### Naming the client workspace

To name your client workspace, or to use a different workspace, set the environment variable `P4CLIENT` to the name of the client workspace

**Example:** *Naming the client workspace*

*Ed is working on the code for Elm. He wants to refer to the collection of files he's working on by the name `eds_elm`. In the Korn or Bourne shells, he'd type:*

```
$ P4CLIENT=eds_elm ; export P4CLIENT
```

Each operating system or shell has its own method of defining environment variables. See the “Environment Variables” section of the *Perforce Command Reference* for details.

### Describing the client workspace to the Perforce server

Once the client workspace has been named, it must be identified and described to the Perforce server with the `p4 client` command. Typing `p4 client` brings up the client definition form in a standard text editor; once the form is filled in and the editor exited, the Perforce server is able to move files between the depot and the client workspace.

**Note** | Many `p4` commands, including `p4 client`, display a form for editing in a standard text editor. You can define the editor your client uses through the `P4EDITOR` environment variable.

The `p4 client` form has a number of fields; the two most important are the `Root` and `View`. The meanings of these fields are as follows:

Field	Meaning
Root:	Identifies the top subdirectory of the client workspace. This should be the lowest-level directory that includes all the files and directories that you'll be working with.
View:	Describes which files and directories in the depot are available to the client workspace, and where the files in the depot will be located within the client workspace.

**Example:** *Setting the client root and the client view:*

*Ed is working with his Elm files in a setting as described above. He's set the environment variable `P4CLIENT` to `eds_elm`. Now he types `p4 client` from his home directory, and sees the following form:*

```
Client: eds_elm
Owner:  ed
Description:
    Created by ed.
Root:   /usr/edk
Options:      nomodtime noclobber
View:
    //depot/...    //eds_elm/...
```

*With these default settings, all files in Ed's home directory of `/usr/edk` (including files unrelated to Ed's work) would be mapped to the depot, and all files in the depot would be mapped to Ed's home directory, likely cluttering it with files Ed has no interest in working with.*

*Ed would like to keep all Elm-related material in a subdirectory in his home directory; `/usr/edk/elm`, and he would like this directory to contain only files in the `elm_proj` portion of the depot. He therefore changes the values in the `Root:` and `View:` fields as follows:*

```
Client: eds_elm
Owner:  ed
Description:
    Created by ed.
Root:   /usr/edk/elm
Options:      nomodtime noclobber
View:
    //depot/elm_proj/... //eds_elm/...
```

*This specifies that `/usr/edk/elm` is the top level directory of Ed's client workspace, and that the files under this workspace directory are to be mapped to the depot's `elm_proj` subtree.*

*When Ed is done, he exits from the editor, and the `p4 client` command saves his changes.*

**Warning!** To use Perforce properly, it is crucial to understand how views work. Views are explained in more detail at the start of the next chapter.

**Windows** | In Windows environments, a single client workspace may span multiple drives by setting the client root to null and including the drive letter in the client view. Please see the platform notes on the Perforce web site for further details.

The read-only `Client:` field contains the string stored in the `P4CLIENT` environment variable. The `Description:` can be filled with anything at all; it's a place where you can enter text that describes the contents of this client workspace. The `View:` describes the relationship between files in the depot and files in the client workspace.

Creating a client specification has no immediate visible effect; no files are created when a client specification is created or edited. The client specification simply indicates where files will be located when subsequent Perforce commands are used.

You'll also use `p4 client` to change existing client specifications. This is described in "Perforce Basics: The Details" on page 33.

## Copying depot files into your workspace

---

Use `p4 sync` to retrieve files from the depot into a client workspace.

**Note** | If you're setting up a brand new depot, `p4 sync` won't do anything, since there are no files in the depot to copy to the client workspace yet. If this is the case, start by copying files from your client workspace to the depot with `p4 add`, as described in "Adding files to the depot" on page 25.

**Example:** *Copying files from the depot to a client workspace.*

*Lisa has been assigned to fix bugs in Ed's code. She creates a directory called `elm_ws` within her own directory, and sets up a client workspace; now she wants to copy all the existing `elm` files from the depot into her workspace.*

```
$ cd ~/elm_ws
$ p4 sync
//depot/elm_proj/doc/elmdoc.0#2 - added as /usr/lisag/elm_ws/doc/elmdoc.0
//depot/elm_proj/doc/elmdoc.1#2 - added as /usr/lisag/elm_ws/doc/elmdoc.1
<etc.>
```

*Once the command completes, the most recent revisions of all the files in the depot that are mapped through her client workspace view will be available in her workspace.*

The `p4 sync` command maps depot files through the client view, compares the result against the current client contents, and then adds, updates, or deletes files in the client workspace as needed to bring the client contents in sync with the depot. `p4 sync` can take filenames as parameters, with or without wildcards, to limit the files it retrieves.

If a file exists within a particular subdirectory in the depot, but that directory does not yet exist in the client workspace, the directory is created within the client workspace at sync time.



The job of `p4 sync` is to match the state of the client workspace to that of the depot; thus, if a file has been deleted from the depot, `p4 sync` deletes it from the client workspace.

## Updating the depot with files from your workspace

---

Any file in a client workspace can be added to, updated in, or deleted from the depot. This is accomplished in two steps:

1. Perforce is told the new state of client workspace files with the commands `p4 add filenames`, `p4 edit filenames`, or `p4 delete filenames`. When these commands are given, the corresponding files are listed in a Perforce changelist, which is a list of files and operations on those files to be performed in the depot.
2. The operations are performed on the files in the changelist when the `p4 submit` command is given.

**Note** | This chapter discusses only the *default changelist*, which is automatically maintained by Perforce. Changelists can also be created by the user; for a full discussion, see Chapter 7, Changelists.

The commands `p4 add`, `p4 edit`, and `p4 delete` do not immediately add, edit, or delete files in the depot. Instead, the affected file and the corresponding operation are listed in the *default changelist*, and the files in the depot are affected only when this changelist is submitted to the depot with `p4 submit`. This allows a set of files to be updated in the depot all at once: when the changelist is submitted, either all of the files in the changelist are affected, or none of them are.

When a file has been opened with `p4 add`, `p4 edit`, or `p4 delete`, but the corresponding changelist has not yet been submitted in the depot, the file is said to be *open* in the client workspace.

### Adding files to the depot

To add a file or files to the depot, type `p4 add filename(s)`. The `p4 add` command opens the file(s) for add and lists them in the default changelist, but they won't be added to the depot until the `p4 submit` command is given.

**Example:** *Adding files to a changelist.*

*Ed is writing a help manual for Elm. The files are named `elmdoc.0` through `elmdoc.3`, and they're sitting in the `doc` subdirectory of his client workspace root. He wants to add these files to the depot.*

```
$ cd ~/elm/doc
$ p4 add elmdoc.*
//depot/elm_proj/doc/elmdoc.0#1 - opened for add
//depot/elm_proj/doc/elmdoc.1#1 - opened for add
//depot/elm_proj/doc/elmdoc.2#1 - opened for add
//depot/elm_proj/doc/elmdoc.3#1 - opened for add
```

At this point, the files he wants to add to the depot have been added to his default changelist. However, the files are not actually stored in the depot until the `p4 submit` command is given.

**Example:** *Submitting a changelist to the depot.*

*Ed is ready to submit his added files to the depot. He types `p4 submit` and sees the following form in a standard text editor:*

```
Change: new
Client: edk
User: edk
Status: new
Description:
    <enter description here>
Files:
    //depot/elm_proj/doc/elmdoc.0 # add
    //depot/elm_proj/doc/elmdoc.1 # add
    //depot/elm_proj/doc/elmdoc.2 # add
    //depot/elm_proj/doc/elmdoc.3 # add
```

*Ed changes the contents of the `Description:` field to describe what these file updates do. When he's done, he quits from the editor, and the new files are added to the depot.*

The `Description:` field contents must be changed, or the depot update won't be accepted. Lines can be deleted from the `Files:` field; any files deleted from this list will carry over to the next default changelist, and will appear again the next time `p4 submit` is performed.

## Adding more than one file at once

Multiple file arguments can be provided on the command line.

**Example:** *Using multiple file arguments on a single command line.*

*Ed wants to add all of his Elm library, documentation, and header files to the depot.*

```
$ cd ~
$ p4 add elm/lib/* elm/hdrs/* elm/doc/*
//depot/elm_proj/lib/Makefile.SH#1 - opened for add
//depot/elm_proj/lib/add_site.c#1 - opened for add
//depot/elm_proj/lib/addrmchusr.c#1 - opened for add
<etc.>
```

*After p4 add is finished, Ed then does a p4 submit.*

If the directory containing a new file does not exist in the depot, it is automatically created within the depot at *submit* time.

The operating system's write permission on submitted files is turned off in the client workspace when `p4 submit` is performed. This helps ensure that file editing is done with Perforce's knowledge. The write permissions are turned back on by `p4 edit`, which is described below.

You might have noticed in the example above that the filenames are displayed as *filename#1*. Perforce always displays filenames with a *#N* suffix; the *#N* indicates that this is the *N*th revision of this file. Revision numbers are always assigned sequentially.

**Warning!** If a submit fails, the default changelist will be assigned a number, and you'll need to submit that changelist in a slightly different way. Please see Chapter 5, Perforce Basics: Resolving File Conflicts for instructions on resolving file conflicts.

## Populating empty depots

In Perforce, there's no difference between adding files to an empty depot and adding files to a depot that already contains other files. For this reason, populate new, empty depots by adding files from a client workspace with `p4 add`, as described above.

## Editing files in the depot

To open a file for `edit`, use `p4 edit`. This has two effects:

- The file(s) write permissions are turned on in the client workspace, and
- The file(s) to be edited are added to the default changelist.

Since the files must have their write permission turned back on before they can be edited, you must give the `p4 edit` command before attempting to edit the file.

To save the new file revision in the depot, use `p4 submit`, as described above.

**Note** | Before a file can be opened for `edit`, it must already have been added to the depot with `p4 add`, or copied into the client workspace from the depot with `p4 sync`.

**Example:** *Opening a file for edit:*

*Ed wants to make changes to his `elmdoc.3` file. He opens the file for edit.*

```
$ cd ~/elm
$ p4 edit doc/elmdoc.3
//depot/elm_proj/doc/elmdoc.3#1 - opened for edit
```

*He then edits the file with any text editor. When he's finished, he submits the file to the depot with `p4 submit`, as described above.*

## Deleting files from the depot

Files are deleted from the depot in a way similar to the way in which they are added and edited.

The `p4 delete` command opens the file for `delete` in the default changelist, and then `p4 submit` is used to delete the file from the depot. The `p4 delete` command also deletes the file from the client workspace; this occurs when the `p4 delete` command is given. The deletion of the file in the depot occurs only after the changelist with the delete operation is submitted.

**Example:** *Deleting a file from the depot.*

*Ed's file `doc/elmdoc.3` is no longer needed. He deletes it from both his client workspace and from the depot as follows:*

```
$ cd ~/elm/doc
$ p4 delete elmdoc.3
//depot/elm_proj/doc/elmdoc.3#1 - opened for delete
```

*The file is deleted from the client workspace immediately, but it is not deleted from the depot until he gives the `p4 submit` command.*

Once the changelist is submitted, it appears as if the file has been deleted from the depot; however, old file revisions are never actually removed. This makes it possible to read older revisions of “deleted” files back into the client workspace.

## Submitting with multiple operations

Multiple files can be included in any changelist. Submitting the changelist to the depot works atomically: either all the files are updated in the depot, or none of them are. (In Perforce's terminology, this is called an *atomic change transaction*). Changelists can be used to keep files together that have a common purpose.

**Example:** *Adding, updating, and deleting files in a single submit:*

*Ed is writing the portion of Elm that is responsible for multiple folders (multiple mailboxes). He has a new source file `src/newmbox.c`, and he needs to edit the header file `hdrs/s_elm.h` and the `doc/elmdoc` files. He adds the new file and prepares to edit the existing files:*

```
$ cd ~
$ p4 add elm/src/newmbox.c
//depot/elm_proj/src/newmbox.c#1 - opened for add
<etc.>
$ p4 edit elm/hdrs/s_elm.h doc/elmdoc.*
//depot/elm_proj/hdrs/s_elm.h#1 - opened for edit
//depot/elm_proj/doc/elmdoc.0#1 - opened for edit
//depot/elm_proj/doc/elmdoc.1#1 - opened for edit
//depot/elm_proj/doc/elmdoc.2#2 - opened for edit
```

*He edits the existing files and then performs a `p4 submit` of the default changelist:*

```
Change: new
Client: eds_elm
User: edk
Status: new
Description:
    Changes to Elm's "mult. mailbox" functionality
Files:
    //depot/elm_proj/doc/elmdoc.0 # edit
    //depot/elm_proj/doc/elmdoc.1 # edit
    //depot/elm_proj/doc/elmdoc.2 # edit
    //depot/elm_proj/hdrs/s_elm.h # edit
    //depot/elm_proj/src/newmbox.c # add
```

*All of his changes supporting multiple mailboxes are grouped together in a single changelist; when Ed quits from the editor, either all of these files are updated in the depot, or, if the submission fails for any reason, none of them are.*

Files can be deleted from the `Files:` field; these files are moved into the next default changelist, and appear again the next time `p4 submit` is performed.

## Backing out: reverting files to their unopened states

---

Any file opened for `add`, `edit`, or `delete` can be removed from its changelist with `p4 revert`. This command reverts the file in the client workspace back to its unopened state, and any local modifications to the file are lost.

**Example:** *Reverting a file back to the last synced version.*

*Ed wants to edit a set of files in his src directory: `leavembox.c`, `limit.c`, and `signals.c`. He opens the files for `edit`:*

```
$ cd ~elm/src
$ p4 edit leavembox.c limit.c signals.c
//depot/elm_proj/src/leavembox.c#2 - opened for edit
//depot/elm_proj/src/limit.c#2 - opened for edit
//depot/elm_proj/src/signals.c#1 - opened for edit
```

*and then realizes that `signals.c` is not one of the files he will be working on, and that he didn't mean to open it. He can revert `signals.c` to its unopened state with `p4 revert`:*

```
$ p4 revert signals.c
//depot/elm_proj/src/signals.c#1 - was edit, reverted
```

If `p4 revert` is used on a file that had been opened with `p4 delete`, it will appear back in the client workspace immediately. If `p4 add` was used to open the file, `p4 revert` removes it from the changelist, but leaves the client workspace file intact. If the reverted file was originally opened with `p4 edit`, the last synced version will be written back to the client workspace, overwriting the newly-edited version of the file. In this case, you may want to save a copy of the file before running `p4 revert`.

## Basic reporting commands

---

Perforce provides some 20+ reporting commands. Each chapter in this manual ends with a description of the reporting commands relevant to the chapter topic. All the reporting commands are discussed in greater detail in Chapter 11, Reporting and Data Mining.

The most basic reporting commands are `p4 help` and `p4 info`.

Command	Meaning
<code>p4 help commands</code>	Lists all Perforce commands with a brief description of each.
<code>p4 help command</code>	For any <i>command</i> provided, gives detailed help about that command. For example, <code>p4 help sync</code> provides detailed information about the <code>p4 sync</code> command.
<code>p4 help usage</code>	Describes command-line flags common to all Perforce commands.

<b>Command</b>	<b>Meaning</b>
<code>p4 help views</code>	Gives a discussion of Perforce view syntax
<code>p4 help</code>	Describes all the arguments that can be given to <code>p4 help</code> .
<code>p4 info</code>	Reports information about the current Perforce system: the server address, client root directory, client name, user name, Perforce version, and a few other tidbits.

Two other reporting commands are used quite often:

<b>Command</b>	<b>Meaning</b>
<code>p4 have</code>	Lists all file revisions that the Perforce server knows you have in the client workspace.
<code>p4 sync -n</code>	Reports what files would be updated in the client workspace by <code>p4 sync</code> without actually performing the sync operation.





---

## Chapter 4    **Perforce Basics: The Details**

---

This chapter covers the Perforce rules in detail. The topics discussed include views, mapping depots to client workspaces, Perforce wildcards, rules for referring to older file revisions, file types, and form syntax. For a brief overview of Perforce, refer to Chapter 3, Perforce Basics: Quick Start.

### **Description of the Client Workspace**

---

A Perforce client workspace is a collection of source files managed by Perforce on a host. Each collection is given a name which identifies the client workspace to the Perforce server. By default, the name is simply the host's name, but this can be overridden by the environment variable `P4CLIENT`. There can be more than one Perforce client workspace on a client host.

All files within a Perforce client workspace share a common root directory, called the *client root*. In the degenerate case, the client root can be the host's root, but in practice the client root is the lowest level directory under which the managed source files sit.

Perforce manages the files in a client workspace in three direct ways:

- It creates, updates, or deletes files as required in the workspace when the user requests Perforce to synchronize the client workspace with the depot,
- It turns on write permission when the user requests to edit a file, and
- It turns off write permission and submits updated versions back to the depot when the user has finished editing the file and submits his or her changes.

The entire Perforce client workspace state is tracked by the Perforce server. The server knows what files a client workspace has, where they are, and which files have write permission turned on.

Perforce's management of a client workspace requires a certain amount of cooperation from the user. Because client files are ordinary files with `write` permission turned off, willful users can circumvent the system by turning on write permission, directly deleting or renaming files, or otherwise modifying the file tree supposedly under Perforce's control.

Perforce counters this with two measures: first, Perforce has explicit commands to verify that the client workspace state is in accord with the server's recording of that state; second, Perforce tries to make using Perforce at least as easy as circumventing it. For example: to make a temporary modification to a file, it is easier to use Perforce than it is to copy and restore the file manually.

Files not managed by Perforce may also be under a client's root, but are ignored by Perforce. For example, Perforce may manage the source files in a client workspace, while the workspace also holds compiled objects, libraries, and executables, as well as a developer's temporary files.

In addition to accessing the client files, the `p4 client` program sometimes creates temporary files on the client host. Otherwise, Perforce neither creates nor uses any files on the client host.

## Wildcards

---

Perforce uses three wildcards for pattern matching. Any number and combination of these can be used in a single string:

Wildcard	Meaning
*	Matches anything except slashes; matches only within a single directory.
...	Matches anything including slashes; matches across multiple directories.
%d	Used for parametric substitution in views. See “Changing the order of filename substrings” on page 39 for a full explanation.

The “...” wildcard is passed by the `p4 client` program to the Perforce server, where it is expanded to match the corresponding files known to `p4d`. The \* wildcard is expanded locally by the OS shell before the `p4` command is sent to the server, and the files that match the wildcard are passed as multiple arguments to the `p4` command. To have Perforce match the \* wildcard against the contents of the depot, it must be escaped, usually with quotes or a backslash. Most command shells don't interfere with the other two wildcards.

### Wildcards and “p4 add”

The “...” wildcard can't be used with the `p4 add` command. The “...” wildcard is expanded by the Perforce server, and since the server doesn't know what files are being added (after all, they're not in the depot yet), it can't expand that wildcard. The \* wildcard may be used with `p4 add`, as it is expanded by the local OS shell, not by the `p4d` server.

## Mapping the Depot to the Client Workspace

---

Just as a client name is simply an alias for a particular directory on the client machine, a depot name is an alias for a directory on the Perforce server. The relationship between files in the depot and files in the client workspace is described in the *client view*, and it is set with the `p4 client` command.

When you type `p4 client`, Perforce displays a variation of the following form:

```
Client: eds_elm
Owner:  edk
Description:
    Created by ed.
Root:   /usr/edk/elm
Options:      nomodtime noclobber
View:
    //depot/...    //eds_elm/...
```

The contents of the `View:` field determine where client files get stored in the depot, and where depot files are copied to in the client.

**Note** | The `p4 client` form has more fields than are described here. For a full discussion, please see the *Perforce Command Reference*.

## Multiple depots

By default, there is a single depot in each Perforce server, and the name of the depot is `depot`. The Perforce system administrator can create multiple depots on the same Perforce server.

If your system administrator has created multiple depots on your server, the default client view will look something like this:

```
View:
    //depot/...    //eds_elm/depot/...
    //user_depot/... //eds_elm/user_depot/...
    //projects/... //eds_elm/projects/...
```

The *Perforce System Administrator's Guide* explains how to create multiple depots.

## Using views

Views consist of multiple lines, or *mappings*, and each mapping has two parts. The left-hand side specifies one or more files within the depot, and has the form:

```
//depotname/file_specification
```

The right-hand side of each mapping describes one or more files within the client workspace, and has the form:

```
//clientname/file_specification
```

The left-hand side of a client view mapping is called the *depot side*, and the right-hand side is the *client side*.

The default view in the example above is quite simple: it maps the entire depot to the entire client workspace. Views can contain multiple mappings, and can be much more complex, but all client views, no matter how elaborate, perform the same two functions:

- *Determine which files in the depot can be seen by a client workspace.*

This is determined by the sum of the depot sides of the mappings within a view. A view might allow the client workspace to retrieve every file in the depot, or only those files within two directories, or only a single file.

- *Construct a one-to-one mapping between files in the depot and files in the client workspace.*

Each mapping within a view describes a subset of the complete mapping. The one-to-one mapping might be straightforward; for example, the client workspace file tree might be identical to a portion of the depot's file tree. Or it can be oblique; for example, a file might have one name in the depot and another in the client workspace, or be moved to an entirely different directory in the client workspace. No matter how the files are named, there is always a one-to-one mapping.

To determine the exact location of any client file on the host machine, substitute the value of the `p4 client` form's `Root:` field for the client name on the client side of the mapping. For example, if the `p4 client` form's `Root:` field for the client `eds_elm` is set to `/usr/edk/elm`, then the file `//eds_elm/doc/elmdoc.1` will be found on the client host in `/usr/edk/elm/doc/elmdoc.1`.

**Windows** | On Windows machines, the Perforce client must be specified in a slightly different way if it is to span multiple drives. For details, please see the platform notes at <http://www.perforce.com/perforce/technical.html>

## Wildcards in views

Any wildcard used on the depot side of a mapping must be matched with an identical wildcard in the mapping's client side. Any string matched by the wildcard is identical on both sides.

In the client view

```
//depot/elm_proj/...    //eds_elm/...
```

the single mapping contains Perforce's "... " wildcard, which matches everything including slashes. The result is that any file in the `eds_elm` client workspace will be mapped to the same location within the depot's `elm_proj` file tree. For example, the file `//depot/elm_proj/nls/gencat/README` is mapped to the client workspace file `//eds_elm/nls/gencat/README`.

To properly specify file trees, use the "... " wildcard after a trailing slash. (If you specify only `//depot/elm_proj...`, then the resulting view also includes files and directories such as `//depot/elm_project_core.dumps`, which is probably not what you intended.)

## Types of mappings

By changing the `View:` field, it is possible to map only part of a depot to a client workspace. It is even possible to map files within the same depot directory to different client workspace directories, or to have files named differently in the depot and the client workspace. This section discusses Perforce's mapping methods.

### Direct client-to-depot views

The default view in the form presented by `p4 client` maps the entire client workspace tree into an identical directory tree in the depot. For example, the default view

```
//depot/... //eds_elm/...
```

indicates that any file in the directory tree under the client `eds_elm` will be stored in the identical subdirectory in the depot. This view is usually considered to be overkill, as most users only need to see a subset of the files in the depot.

### Mapping the full client to only part of the depot

Usually, you'll only want to see part of the depot. Change the left-hand side of the `View:` field to point to only the relevant portion of the depot.

**Example:** *Mapping part of the depot to the client workspace.*

*Bettie is rewriting the documentation for Elm, which is found in the depot within its `doc` subdirectory. Her client is named `elm_docs`, and her client root is `/usr/bes/docs`; she types `p4 client` and sets the `View:` field as follows:*

```
//depot/elm_proj/doc/... //elm_docs/...
```

*The files in `//depot/elm_proj/doc` are mapped to `/usr/bes/docs`. Files not beneath the `//depot/elm_proj/doc` directory no longer appear in Bettie's workspace.*

### Mapping files in the depot to different parts of the client

Views can consist of multiple mappings, which are used to map portions of the depot file tree to different parts of the client file tree. If there is a conflict in the mappings, later mappings have precedence over the earlier ones.

**Example:** *Multiple mappings in a single client view.*

*The `elm_proj` subdirectory of the depot contains a directory called `doc`, which has all of the Elm documents. Included in this directory are four files named `elmdoc.0` through `elmdoc.3`. Mike wants to separate these four files from the other documentation files in his client workspace, which is called `mike_elm`.*

To do this, he creates a new directory in his client workspace called `help`, which is located at the same level as his `doc` directory. The four `elmdoc` files will go here, so he fills in the `View:` field of the `p4 client` form as follows:

```
View:
      //depot/...                //mike_elm/...
      //depot/elm_proj/doc/elmdoc.*  //mike_elm/help/elmdoc.*
```

Any file whose name starts with `elmdoc` within the depot's `doc` subdirectory is caught by the later mapping and appears in Mike's workspace's `help` directory; all other files are caught by the first mapping and appear in their normal location. Conversely, any files beginning with `elmdoc` within Mike's client workspace `help` subdirectory are mapped to the `doc` subdirectory of the depot.

**Note** | Whenever you map two sections of the depot to different parts of the client workspace, some depot and client files will remain unmapped. See "Two mappings can conflict and fail" on page 39 for details.

### Excluding files and directories from the view

*Exclusionary mappings* allow files and directories to be excluded from a client workspace. This is accomplished by prefacing the mapping with a minus sign ( - ). Whitespace is not allowed between the minus sign and the mapping.

**Example:** *Using views to exclude files from a client workspace.*

*Bill, whose client is named `billm`, wants to view only source code; he's not interested in the documentation files. His client view would look like this:*

```
View:
      //depot/elm_proj/...        //billm/...
      -//depot/elm_proj/doc/...  //billm/doc/...
```

*Since later mappings have precedence over earlier ones, no files from the depot's `doc` subdirectory will ever be copied to Bill's client. Conversely, if Bill does have a `doc` subdirectory in his client, no files from that subdirectory will ever be copied to the depot.*

### Allowing filenames in the client to differ from depot filenames

Mappings can be used to make the filenames in the client workspace differ from those in the depot.

**Example:** *Files with different names in the depot and client workspace.*

*Mike wants to store the files as above, but he wants to take the `elmdoc.X` files in the depot and call them `helpfile.X` in his client workspace. He uses the following mappings:*

```
View:
    //depot/elm_proj...           //mike_elm/...
    //depot/elm_proj/doc/elmdoc.* //mike_elm/help/helpfile.*
```

Each wildcard on the depot side of a mapping must have a corresponding wildcard on the client side of the same mapping. The wildcards are replaced in the copied-to direction by the substring that the wildcard represents in the copied-from direction.

There can be multiple wildcards; the *n*th wildcard in the depot specification corresponds to the *n*th wildcard in the client description.

### Changing the order of filename substrings

The `%d` wildcard matches strings similarly to the `*` wildcard, but `%d` can be used to rearrange the order of the matched substrings.

**Example:** *Changing string order in client workspace names.*

*Mike wants to change the names of any files with a dot in them within his `doc` subdirectory in such a way that the file's suffixes and prefixes are reversed in his client workspace. For example, he'd like to rename the `Elm.cover` file in the depot `cover.Elm` in his client workspace. (Mike can be a bit difficult to work with). He uses the following mappings:*

```
View:
    //depot/elm_proj/...           //mike_elm/...
    //depot/elm_proj/doc/%1.%2     //mike_elm/doc/%2.%1
```

### Two mappings can conflict and fail

When you use multiple mappings in a single view, some files may map to two separate places in the depot or on the client. When a file doesn't map to the same place in both directions, Perforce ignores that file.

**Example:** *Mappings that conflict and fail.*

*Joe has constructed a view as follows:*

```
View:
    //depot/proj1/...           //joe/proj1/...
    //depot/proj1/foo          //joe/proj1/bar
```

*The depot file `//depot/proj1/bar` maps to `//joe/proj1/bar`, but that same client file `//joe/proj1/bar` maps back to the depot via the higher-precedence second line to `//depot/proj1/foo`. Because the file would be written back to a different location in the*

*depot than where it was read from, Perforce doesn't map this name at all, and the file is ignored.*

In older versions of Perforce, this was often used as a trick to exclude particular files from the client workspace. Because Perforce now has exclusionary mappings, this type of mapping is no longer useful, and should be avoided.

## Editing Existing Client Specifications

---

You can use `p4 client` at any time to change the client workspace specification.

Note that changing a client specification has no immediate effect on the locations of any files; the locations of files in your workspace are affected only when the client specification is *used* in subsequent commands.

This is particularly important for two types of client spec changes, specifically changes to the client view and changes to the client root:

- If you change only the value of the client workspace `View:` field with `p4 client`, perform a `p4 sync` immediately afterward. The files in the client workspace will be deleted from their old locations and written to their new locations.
- If you use `p4 client` to change the client workspace `Root:`, use `p4 sync #none` to remove the files from their old location in the workspace *before* using `p4 client` to change the client root. After using `p4 client` to change the root directory, perform a `p4 sync` to copy the files to their new locations within the client view. (If you forget to perform the `p4 sync #none` before changing the client view, you can always remove the files from their client workspace locations manually).

**Warning!** It's not a good idea to change both the client `Root:` and the client `View:` at the same time. Change either the `Root:` or the `View:` according to the instructions above, make sure that the files are in place in the client workspace, and then change the other.

## Deleting an existing client specification

An existing client workspace specification can be deleted with `p4 client -d clientname`. Deleting a client specification has no effect on any files in the client workspace or depot; it simply removes the Perforce server's record of the mapping between the depot and the client workspace.

To delete existing files from a client workspace, use `p4 sync #none` (described in "Specifying Older File Revisions" on page 46) on the files *before* deleting the client specification, or use the standard local OS deletion commands *after* deleting the client specification.



## Client specification options

The `Options:` field contains six values, separated by spaces. Each of the six options have two possible settings.

The following table provides the option values and their meanings:

Option	Choice	Default
<code>[no]allwrite</code>	Should unopened files be left writable on the client?	<code>noallwrite</code>
<code>[no]clobber</code>	Should <code>p4 sync</code> overwrite (clobber) writable but unopened files in the client with the same name as the newly synced files?	<code>noclobber</code>
<code>[no]compress</code>	Should the data sent between the client and the server be compressed? Both client and server must be version 99.1 or higher, or this setting will be ignored.	<code>nocompress</code>
<code>[no]crlf</code>	<b>Note: 2000.2 or earlier only!</b> Should CR/LF translation be performed automatically when copying files between the depot and the client workspace? (On UNIX, this setting is ignored).	<code>crlf</code>
<code>[un]locked</code>	Do other users have permission to edit the client specification? (To make a locked client specification truly effective, be sure to set a password for the client's owner with <code>p4 passwd</code> .)  If <code>locked</code> , only the owner is able to use, edit, or delete the client spec. Note that a Perforce superuser is still able to override the lock with the <code>-f</code> (force) flag.	<code>unlocked</code>

Option	Choice	Default
[no]modtime	<p>For files <i>without</i> the +m (modtime) file type modifier:</p> <ul style="list-style-type: none"> <li>• For Perforce clients at the 99.2 level or earlier, if modtime is set, the modification date (on the local filesystem) of a newly synced file is the date and time <i>at the server</i> when the file was submitted to the depot.</li> <li>• For Perforce clients at the 2000.1 level or higher, if modtime is set, the modification date (on the local filesystem) of a newly synced file is the datestamp <i>on the file</i> when the file was submitted to the depot.</li> <li>• If nomodtime is set, the modification date is the date and time of sync, regardless of Perforce client version.</li> </ul> <p>For files <i>with</i> the +m (modtime) file type modifier:</p> <ul style="list-style-type: none"> <li>• For Perforce clients at the 99.2 level or earlier, the +m modifier is ignored, and the behavior of modtime and nomodtime is as documented above.</li> <li>• For Perforce clients at the 2000.1 level or higher, the modification date (on the local filesystem) of a newly synced file is the datestamp <i>on the file</i> when the file was submitted to the depot, <i>regardless</i> of the setting of modtime or nomodtime on the client.</li> </ul>	<p>nomodtime (i.e. date and time of sync) for most files.</p> <p>Ignored for files with the +m file type modifier.</p>
[no]rmdir	Should p4 sync delete empty directories in a client if all files in the directory have been removed?	normdir

### Line-ending conventions (CR/LF translation)

The LineEnd: field controls the line-ending character(s) used for text files in the client workspace.

**Note** | The LineEnd: option is new to Perforce 2001.1, and replaces the old convention of specifying crlf or nocrlf in the Options: field.

The `LineEnd:` field accepts one of five values:

Option	Meaning
<code>local</code>	Use mode native to the client (default)
<code>unix</code>	UNIX-style line endings: <code>LF</code> only
<code>mac</code>	Macintosh-style: <code>CR</code> only
<code>win</code>	Windows-style: <code>CR</code> , <code>LF</code>
<code>share</code>	<p>Shared mode: Line endings are <code>LF</code> with any <code>CR/LF</code> pairs translated to <code>LF</code>-only style before storage or syncing with the depot.</p> <p>When you sync your client workspace, line endings will be <code>LF</code>. If you edit the file on a Windows machine, and your editor inserts <code>CR</code>s before each <code>LF</code>, the extra <code>CR</code>s will not appear in the archive file.</p> <p>The most common use of the <code>share</code> option is for users of Windows workstations who mount their UNIX home directories mounted as network drives; if you sync files from UNIX, but edit the files on a Windows machine, the <code>share</code> option eliminates problems caused by Windows-based editors' insertion of carriage returns in text files.</p>

## Referring to Files on the Command Line

File names provided as arguments to Perforce commands can be referred to in one of two ways: by using the names of the files in the client workspace, or by providing the names of the files in the depot. When providing client workspace file names, you may give the name in either *local* syntax or Perforce syntax.

### Local syntax

*Local syntax* is simply a file's name as specified by the local shell or OS. This name may be an absolute path, or may be specified relative to the current directory, although it can only contain relative components at the beginning of the file name (that is, it doesn't allow `sub/dir/./here/foo.c`). On UNIX, Ed could refer to the `README` file at Elm's top level as `/usr/edk/elm/README`, or in a number of other ways.

### Perforce syntax

Perforce provides its own filename syntax which remains the same across operating systems. Filenames specified in this way begin with two slashes and the client or depot name, followed by the path name of the file relative to the client or depot root directory.

The components of the path are separated by slashes.

#### Examples of Perforce syntax

```
//depot/...  
//elm_client/docs/help.1
```

Perforce syntax is sometimes called *depot syntax* or *client syntax*, depending on whether the file specifier refers to a file in the depot or on the client, but both forms of file specification work the same way.

## Providing files as arguments to commands

Because the client view provides a one-to-one mapping between any file in the client workspace and any file in the depot, *any file can be specified within any PERFORCE command in client syntax, depot syntax, or local syntax*. A depot's file specifier can be used to refer to a file in the client, and vice versa. Perforce will do the necessary mapping to determine which file is actually used.

Any filenames provided to Perforce commands can be specified in any valid local syntax, or in Perforce syntax by depot or client. If a client filename is provided, Perforce uses the client view to locate the corresponding file in the depot. If a depot filename is given, the client view is used to locate the corresponding file in the client workspace.

**Example:** *Uses of different syntaxes to refer to a file.*

*Ed wants to delete the src/lock.c file. He can give the p4 delete command in a number of ways:*

*While in his client root directory, he could type*

```
p4 delete src/lock.c
```

*or, while in the src subdirectory, he could type*

```
p4 delete lock.c
```

*or, while in any directory on the client host, he could type any of the following commands:*

```
p4 delete //eds_elm/src/lock.c  
p4 delete //depot/elm_proj/src/lock.c  
p4 delete /usr/edk/elm/src/lock.c
```

Client names and depot names in a single Perforce server share the same namespace, so Perforce will never confuse a client name with a depot name. Client workspace names and depot names can never be the same.

**Note** | The point of this section is worth repeating: any file can be specified within any Perforce command in client syntax, depot syntax, or local syntax. The examples in this manual will use these syntaxes interchangeably.

## Wildcards and Perforce syntax

Perforce wildcards may be mixed with both local or Perforce syntax. For example:

Wildcard	Meaning
<code>J*</code>	Files in the current directory starting with <code>J</code>
<code>*/help</code>	All files called <code>help</code> in current subdirectories
<code>...</code>	All files under the current directory and its subdirectories
<code>.../*.c</code>	All such files ending in <code>.c</code>
<code>/usr/edk/...</code>	All files under <code>/usr/edk</code>
<code>//weasel/...</code>	All files on client (or depot) <code>weasel</code>
<code>//depot/...</code>	All files in the depot named <code>depot</code>
<code>//...</code>	All files in all depots (when used to specify files on the command line)

## Name and String Limitations

### Illegal characters in filenames and Perforce objects

In order to support internationalization, Perforce allows the use of “unprintable” (i.e. non-ASCII) characters in filenames, label names, client workspace names, and other identifiers.

Perforce’s wildcards and the revision-specifying characters `@` and `#`, are not allowed in file names, label names, or other identifiers:

Character	Reason
<code>@</code>	Perforce revision specifier for date, label name, or changelist number.
<code>#</code>	Perforce revision specifier for revision numbers.
<code>*</code>	Perforce wildcard: matches anything, works within a single directory
<code>...</code>	Perforce wildcard: matches anything, works at the current directory level and includes files in all directory levels below the current level.
<code>%</code>	Perforce wildcard: <code>%0</code> through <code>%9</code> are used for positional substitutions.
<code>/</code>	Perforce separator for pathname components.

Observe that most of these characters tend to be difficult to use in filenames in cross-platform environments: UNIX separates path components with `/`, while many DOS commands interpret `/` as a command line switch. Most UNIX shells interpret `#` as the beginning of a comment. Both DOS and UNIX shells automatically expand `*` to match multiple files, and the DOS command line uses `%` to refer to variables.

### Caveats on non-ASCII filenames

Although non-ASCII characters are allowed in filenames and Perforce identifiers, entering them from the command line may require platform-specific solutions. Users of GUI-based file managers can manipulate such files with drag-and-drop operations.

For internationalization purposes, there are some limitations on how filenames with non-ASCII character sets are displayed. For Release 2001.1, all users should use a common code page setting (under Windows, use the **Regional Settings** applet in the **Control Panel**; under UNIX, set the `LOCALE` environment variable) in order to ensure that all filenames are displayed consistently across all machines in your organization.

If you are using Perforce in internationalized mode, all users must have `P4CHARSET` set properly. For details, see the *Command Reference*.

### Caveats on using spaces in filenames

Filenames can include spaces, but require some care as to their specification in client views and root mappings. If you use spaces in a client view or client root, quote the string in which the space appears.

Make sure not to quote spaces that divide two separate strings; for example, a client mapping containing a directory named “my projects” would be quoted like this:

```
//depot/projects/foo/... "//aclient/my projects/foo/..."
```

Other Perforce objects, such as branch names, client names, label names, and so on, may be *specified* with spaces, but these spaces will be automatically converted to underscores by the Perforce server.

## Name and description lengths

Descriptions in the forms used by `p4 client`, `p4 branch`, and so on, may be of any length. Similarly, all names given to Perforce objects such as branches, clients, and so on, are also limited to 1024 characters.

## Specifying Older File Revisions

---

All of the commands and examples we've seen thus far have been used to operate only on the most recent revisions of particular files, but many Perforce commands can act on older file versions. For example, if Ed types `p4 sync //eds_elm/src/lock.c`, the latest revision, or *head revision*, of `lock.c` is retrieved, but older revisions can be retrieved by tacking a revision specification onto the end of the file name.

**Warning!** Some OS shells treat the revision character # as a comment character if it starts a new word. If your shell is one of these, escape the # before use.

Revision Specifier	Meaning	Examples
<code>file#n</code>	Revision number	<code>p4 sync lock.c#3</code> Refers to revision 3 of file <code>lock.c</code>
<code>file@n</code>	A change number	<code>p4 sync lock.c@126</code> Refers to the version of <code>lock.c</code> when changelist 126 was submitted, even if it was not part of the change. <code>p4 sync //depot/...@126</code> Refers to the state of the entire depot at changelist 126 (numbered changelists are explained in Chapter 7, Changelists).
<code>file@labelname</code>	A label name	<code>p4 sync lock.c@beta</code> The revision of <code>lock.c</code> in the label called <code>beta</code> (labels are explained in Chapter 8, Labels).
<code>file@clientname</code>	A client name. The revision of <i>file</i> last taken into client workspace <i>clientname</i> .	<code>p4 sync lock.c@lisag_ws</code> The revision of <code>lock.c</code> last taken into workspace <code>lisag_ws</code>
<code>file#none</code>	The nonexistent revision.	<code>p4 sync lock.c#none</code> Says that there should be no version of <code>lock.c</code> in the client workspace, even if one exists in the depot.
<code>file#head</code>	The head revision, or latest version, of the file.	<code>p4 sync lock.c#head</code> Except for explicitly noted exceptions, this is identical to referring to the file with no revision specifier.
<code>file#have</code>	The revision on the current client. This is synonymous to <code>@client</code> where <code>client</code> is the current client name.	<code>p4 sync lock.c#have</code> The revision of <code>lock.c</code> found in the current client.

Revision Specifier	Meaning	Examples
<code>file@date</code>	The head revision of the file at 00:00:00 on the morning of that date. Dates are specified as YYYY/MM/DD.	<pre>p4 sync lock.c@1998/05/18</pre> <p>The head revision of <code>lock.c</code> as of 00:00:00, May 18, 1998.</p>
<code>file@"date time"</code>	The head revision of the file in the depot on the given date at the given time. The date is specified as above; the time is specified as HH:MM:SS.	<pre>p4 sync lock.c@"1998/05/18 15:21:34"</pre> <pre>p4 sync lock.c@1998/05/18:15:21:34</pre> <p>The head revision of <code>lock.c</code> as of May 18, 1998, at 3:21:34 pm. Both forms shown above are equivalent.</p> <p>The date and the time must be separated by a single space or a colon, and the entire string should be quoted. The time is specified on the 24-hour clock.</p> <p><b>Warning!</b> Perforce allows you to search on dates with two-digit years, but these years are assumed to fall in the twentieth century. For safety's sake, use four-digit years.</p>

In all cases, if a file doesn't exist at the given revision number, it appears as if the file doesn't exist at all. Thus, using a label to refer to a file that isn't in the label is indistinguishable from referring to a file that doesn't exist at all.

Date and time specifications are always interpreted with respect to the local time zone of the Perforce server. Note that because the server stores times internally in terms of number of seconds since the Epoch (00:00:00 GMT Jan. 1, 1970), if you move your server across time zones, the times recorded on the server will automatically be reported in the new timezone.

The date, time, and time zone in effect at your Perforce server are displayed in the "Server date:" line in the output of `p4 info`.

## Using revision specifications without filenames

Revision specifications can be provided without file names. This limits the command's action to the specified revision of all files in the depot or in the client's workspace. Thus, `#head` refers to the head revisions of all files in the depot, and `@labelname` refers to the revisions of all files in the named label.



**Example:** *Retrieving files using revision specifiers.*

*Ed wants to retrieve all the doc files into his Elm doc subdirectory, but he wants to see only those revisions that existed at change number 30. He types*

```
p4 sync //eds_elm/doc/*@30
```

*Later, he creates another client for a different user. The new client should have all of the file revisions that Ed last synced. Ed sets up the new client specification and types*

```
p4 sync //depot/elm_proj/...@eds_elm
```

*He could have typed*

```
p4 sync @eds_elm
```

*and the effect would have been the same.*

**Example:** *Removing files from the client workspace.*

*Another client needs all its files removed, but these files shouldn't be deleted from the depot. Ed sets P4CLIENT to the correct clientname, and types*

```
p4 sync #none
```

*The files are removed from his workspace ("synced to the nonexistent revision"), but not from the depot.*

## Revision Ranges

A few Perforce client commands can limit their actions to a range of revision numbers, rather than just a single revision. A revision range is two revision specifications, separated by a comma.

If only a single revision is given where a revision range is expected, the named revision specifies the end of the range, and the start of the range is assumed to be 1. If no revision number or range is given where a revision range is expected, the default is all revisions.

The commands that accept revision range specifications are:

p4 changes	p4 file	p4 integrate	p4 jobs
p4 print	p4 sync	p4 verify	

**Example:** *Listing changes with revision ranges.*

*A release manager needs to see a list of all changes to the Elm project in July. He types:*

```
p4 changes //depot/elm_proj/...@2000/7/1,2000/8/1
```

*The resulting list of changes looks like this:*

```
Change 632 on 2000/07/1 by edk@eds_elm 'Started work'  
Change 633 on 2000/07/1 by edk@eds_elm 'First build w/bug fix'  
...  
Change 673 on 2000/07/31 by edk@eds_elm 'Final build for QA'
```

*He can then use `p4 describe change#` against any desired change for a full description.*

## File Types

---

Perforce supports six base file types: `text` files, `binary` files, `unicode` files, `native apple` files on the Macintosh, `Mac resource` forks, and `UNIX symlinks`. File type modifiers are then applied to the base types allowing for support of RCS keyword expansion, file compression on the server, and more.

Perforce attempts to determine the type of the file automatically: when a file is opened with `p4 add`, Perforce first decides if the file is a regular file or a symbolic link, and then examines the first part of the file to determine whether it's `text` or `binary`. If any non-text characters are found, the file is assumed to be `binary`; otherwise, the file is assumed to be `text`. (Files of type `unicode` are detected only when the server is running in internationalized mode; for details, see your system administrator.)

Some file formats (for example, Adobe PDF files, and Rich Text Format files) are actually `binary` files, but can sometimes be mistakenly detected by Perforce as being of type `text`. For these files, your system administrator can use the `p4 typemap` command to set up a table matching Perforce file types to file name specifications.

Whenever you open a new file for `add`, Perforce checks the `typemap` table. If the file matches an entry in the table, Perforce uses the file type specified in the table, rather than the file type it would have otherwise used. You can override the file type specified in the `typemap` table by specifying it on the command line with the `-t filetype` modifier.

Once set, a file's type is normally inherited from one revision to the next, but can be overridden or changed with the `-t` flag:

- `p4 add -t filetype filespec` adds the files as the type you've specified.
- `p4 edit -t filetype filespec` opens the file for `edit`; when the files are submitted, their type will be changed.
- `p4 reopen -t filetype filespec` changes the type of a file that's already open for `add` or `edit`.

The `filetype` argument is specified as `basetype+modifiers`. File type modifiers may be combined; for instance, to change the file type of your Perl script `myscript.pl` to executable text with RCS keyword expansion, use `p4 edit -t text+kx myscript.pl`. Full listings of the base file types and their modifiers are provided below.

The type of an existing file can be determined with `p4 opened` or `p4 files`.

File revisions of binary files are normally stored in full within the depot, but only changes made to text files since the previous revision are normally stored. This is called *delta storage*, and Perforce uses RCS format to store its deltas. The file's type determines whether *full file* or *delta* storage is used. When delta storage is used, file merges and file compares can be performed. Files that are stored in their full form can't be merged or compared. (RCS format and delta storage are described in more detail at the start of the next chapter).

Some of the file types are compressed to `gzip` format for storage in the depot. The compression occurs during the submission process, and decompression happens while syncing. The client workspace always contains the file as it was submitted.

**Warning!** Do not try to fool Perforce into storing binary files in delta storage by changing the file type to `text`! If you add a file that contains a `^z` as text from a Windows client, only the part of the file up to the `^z` will be stored in the depot.

**Note** Versions of Perforce prior to 99.1 used a very different, somewhat limited, set of file types. These types have been maintained for backwards compatibility. Mappings between the old and new file types are described in the *Perforce Command Reference*.

The base Perforce file types are:

Keyword	Description	Comments	Server Storage Type
<code>text</code>	Text file	Treated as text on the client. Line-ending translations are performed automatically on Windows and Macintosh clients.	delta
<code>binary</code>	Non-text file	Accessed as binary files on the client. Stored compressed within the depot.	full file, compressed
<code>symlink</code>	Symbolic link	UNIX clients (and the BeOS client) access these as symbolic links. Non-UNIX clients treat them as (small) text files.	delta

Keyword	Description	Comments	Server Storage Type
apple	Multi-forked Macintosh file	AppleSingle storage of Mac data fork, resource fork, file type and file creator. New to Perforce 99.2.  For full details, please see the Mac platform notes at <a href="http://www.perforce.com/perforce/technical.html">http://www.perforce.com/perforce/technical.html</a>	full file, compressed, AppleSingle format.
resource	Macintosh resource fork	The only file type for Mac resource forks in Perforce 99.1 and before. This is still supported, but we recommend using the new apple file type instead.  For full details, please see the Mac platform notes at <a href="http://www.perforce.com/perforce/technical.html">http://www.perforce.com/perforce/technical.html</a>	full file, compressed
unicode	Unicode file	Perforce servers operating in internationalized mode support a Unicode file type. These files are translated into the local character set.  For details, see the <i>System Administrator's Guide</i> .	Stored as UTF-8

The file type modifiers are:

Modifier	Description	Comments
+x	Execute bit set on client	Used for executable files.
+w	File is always writable on client	
+ko	Old-style keyword expansion	Expands only the \$Id\$ and \$Header\$ keywords:  This pair of modifiers exists primarily for backwards compatibility with versions of Perforce prior to 2000.1, and corresponds to the +k (ktext) modifier in earlier versions of Perforce.

Modifier	Description	Comments
+k	RCS keyword expansion	<p>Expands RCS (Revision Control System) keywords.</p> <p>RCS keywords are case-sensitive.</p> <p>When using keywords in files, a colon after the keyword (e.g., \$Id:\$) is optional.</p> <p>Supported keywords are:</p> <ul style="list-style-type: none"> <li>• \$Id\$</li> <li>• \$Header\$</li> <li>• \$Date\$</li> <li>• \$DateTime\$</li> <li>• \$Change\$</li> <li>• \$File\$</li> <li>• \$Revision\$</li> <li>• \$Author\$</li> </ul>
+l	Exclusive open (locking)	<p>If set, only one user at a time will be able to open a file for editing.</p> <p>Useful for binary file types (e.g., graphics) where merging of changes from multiple authors is meaningless.</p>
+C	Server stores the full compressed version of each file revision	Default server storage mechanism for binary files.
+D	Server stores deltas in RCS format	Default server storage mechanism for text files.
+F	Server stores full file per revision	Useful for long ASCII files that aren't read by users as text, such as PostScript files.
+S	Only the head revision is stored on the server	Older revisions are overwritten within the depot. Useful for executable or .obj files.
+m	Preserve original modtime	The file's timestamp on the local filesystem is preserved upon submission and restored upon sync. Useful for third-party DLLs in Windows environments.

The modtime (+m) modifier is a special case: It is intended for use by developers who need to preserve a file's original timestamp. (Normally, Perforce updates the timestamp when a file is synced.) It allows a user to ensure that the timestamp of a file in a client workspace after a `p4 sync` will be the original timestamp existing *on the file* at the time of submission

(that is, *not* the time at the Perforce server at time of submission, and *not* the time on the client at the time of sync).

The most common case where this is useful is development involving the third-party DLLs often encountered in Windows environments. Because the timestamps on such files are often used as proxies for versioning information (both within the development environment and also by the operating system), it is sometimes necessary to preserve the files' original timestamps regardless of a Perforce user's client settings.

The `+m` modifier on a file allows this to happen; if set, Perforce ignores the `modtime` ("file's timestamp at time of submission") or `nomodtime` ("date and time on the client at time of sync") setting at the `p4 client` level when syncing the file, and always restore the file's original timestamp at the time of submit.

RCS keywords are expanded as follows:

Keyword	Expands To	Example
<code>\$Id\$</code>	File name and revision number in depot syntax	<code>\$Id: //depot/path/file.txt#3 \$</code>
<code>\$Header\$</code>	Synonymous with <code>\$Id\$</code>	<code>\$Header: //depot/path/file.txt#3 \$</code>
<code>\$Date\$</code>	Date of last submission in format <code>YYYY/MM/DD</code>	<code>\$Date: 2000/08/18 \$</code>
<code>\$DateTime\$</code>	Date and time of last submission in format <code>YYYY/MM/DD hh:mm:ss</code> Date and time are as of the local time on the Perforce server at time of submission.	<code>\$DateTime: 2000/08/18 23:17:02 \$</code>
<code>\$Change\$</code>	Perforce changelist number under which file was submitted	<code>\$Change: 439 \$</code>
<code>\$File\$</code>	File name only, in depot syntax (without revision number)	<code>\$File: //depot/path/file.txt \$</code>
<code>\$Revision\$</code>	Perforce revision number	<code>\$Revision: #3 \$</code>
<code>\$Author\$</code>	Perforce user submitting the file	<code>\$Author: edk \$</code>

## Forms and Perforce Commands

Certain Perforce commands, such as `p4 client` and `p4 submit`, present a form to the user to be filled in with values. This form is displayed in the editor defined in the environment variable `P4EDITOR`. When the user changes the form and exits the editor, the form is parsed by Perforce, checked for errors, and used to complete the command operation. If there are errors, Perforce gives an error message and you must try again.

The rules of form syntax are simple: keywords must be against the left margin and end with a colon, and values must either be on the same line as the keyword or indented on the lines beneath the keyword. Only the keywords already present on the form are recognized. Some keywords, such as the `Client:` field in the `p4 client` form, take a single value; other fields, such as `Description:`, take a block of text; and others, like `View:`, take a list of lines.

Certain fields, like `Client:` in `p4 client`, can't have their values changed; others, like `Description:` in `p4 submit`, must have their values changed. If you don't change a field that needs to be changed, or vice versa, the worst that will happen is that you'll get an error. We've done our best to make these cases as self-evident as possible. When in doubt, consult the *Perforce Command Reference* or use `p4 help command`.

### Reading forms from standard input; Writing forms to standard output

Any commands that require the user to fill in a form, such as `p4 client` and `p4 submit`, can read the form from standard input with the `-i` flag. Similarly, the `-o` flag can be used to write a form specification to standard output.

These two flags are primarily used in scripts that access Perforce: use the `-o` flag to read a form, process the strings representing the form within your script, and finally, use the `-i` flag to send the processed form back to Perforce.

For example, to create a new job within a script you could first use `p4 job -o > tempfile` to read a blank job specification, then add information to the proper lines in `tempfile`, and finally use a command like `p4 job -i < tempfile` to store the new job in Perforce.

The commands that display forms and can therefore use these flags are:

<code>p4 branch</code>	<code>p4 change</code>	<code>p4 client</code>
<code>p4 job</code>	<code>p4 label</code>	<code>p4 protect</code>
<code>p4 submit*</code>	<code>p4 typemap</code>	<code>p4 user</code>

\* `p4 submit` can take the `-i` flag, but not the `-o` flag.

## General Reporting Commands

Many reporting commands have specialized functions, and these are discussed in later chapters. The following reporting commands give the most generally useful information; all of these commands can take file name arguments, with or without wildcards, to limit reporting to specific files. Without the file arguments, the reports are generated for all files.

These reports always generate information on depot files, not files within the client workspace. As with any other Perforce command, when a client file is provided on the command line, Perforce maps it to the proper depot file.

Command	Meaning
<code>p4 filelog</code>	Generates a report on each revision of the file(s), in reverse chronological order.
<code>p4 files</code>	Lists file name, latest revision number, file type, and other information about the named file(s).
<code>p4 sync -n</code>	Tells you what <code>p4 sync</code> would do, without doing it.
<code>p4 have</code>	Lists all the revisions of the named files within the client that were last gotten from the depot. Without any files specifier, it lists all the files in the depot that the client has.
<code>p4 opened</code>	Reports on all files in the depot that are currently open for edit, add, delete, branch, or integrate within the client workspace.
<code>p4 print</code>	Lists the contents of the named file(s) to standard output.
<code>p4 where</code>	Given a file argument, displays the mapping of that file within the depot, the client workspace, and the local OS.

Revision specifiers can be used with all of these reporting commands, for example `p4 files @clientname` can be used to report on all the files in the depot that are currently found in client `clientname`. See Chapter 11, Reporting and Data Mining, for a more detailed discussion of each of these commands.



---

# Perforce Basics: Resolving File Conflicts

---

File conflicts can occur when two users edit and submit two versions of the same file. Conflicts can occur in a number of ways, but the situation is usually a variant of the following:

1. Ed opens file `f○○` for edit.
2. Lisa opens the same file in her client for edit.
3. Ed and Lisa both edit their client workspace versions of `f○○`.
4. Ed submits a changelist containing `f○○`, and the submit succeeds.
5. Lisa submits a changelist with her version of `f○○`; and her submit fails.

If Perforce were to accept Lisa's version into the depot, the head revision would contain none of Ed's changes. Instead, the changelist is rejected and a resolve must be performed. The resolve process allows a choice to be made: Lisa's version can be submitted in place of Ed's, Lisa's version can be dumped in favor of Ed's, a Perforce-generated merged version of both revisions can be submitted, or the Perforce-generated merged file can be edited and then submitted.

Resolving a file conflict is a two-step process: first the resolve is *scheduled*, then the resolve is *performed*. A resolve is automatically scheduled when a submit of a changelist fails because of a file conflict; the same resolve can be scheduled manually, without submitting, by syncing the head revision of a file over an opened revision within the client workspace. Resolves are always performed with `p4 resolve`.

Perforce also provides facilities for locking files when they are edited. This can eliminate file conflicts entirely.

---

## RCS Format: How Perforce Stores File Revisions

---

Perforce uses RCS format to store its text file revisions; binary file revisions are always saved in full. If you already understand what this means, you can skip to the next section of this chapter, as the remainder of this section explains how RCS format works.

### Only the differences between revisions are stored

A single file might have hundreds, even thousands, of revisions. Every revision of a particular file must be retrievable, and if each revision was stored in full, disk space problems could occur: one thousand 10KB files, each with a hundred revisions, would use a gigabyte of disk space. The scheme used by most SCM systems, including Perforce, is to

save only the latest revision of each file, and then store the differences between each file revision and the one previous.

As an example, suppose that a Perforce depot has three revisions of file `foo`. The head revision (`foo#3`) looks like this:

```
foo#3:
This is a test
of the
emergency
broadcast system
```

Revision two might be stored as a symbolic version of the following:

```
foo#2:
line 3 was "urgent"
```

And revision one might be a representation of this:

```
foo#1:
line 4 was "system"
```

From these partial file descriptions, any file revision can be reconstructed. The reconstructed `foo#1` would read

```
Reconstructed foo#1:
This is a test
of the
urgent
system
```

The RCS (*Revision Control System*) algorithm, developed by Walter Tichy, uses a notation for implementing this system that requires very little storage space and is quite fast. In RCS terminology, it is said that the full text of the head revisions are stored, along with the reverse deltas of each previous revision.

It is interesting to note that the full text of the first revision could be stored, with the deltas leading forward through the revision history of the file, but RCS has chosen the other path: the full text of the head revision of each file is stored, with the deltas leading backwards to the first revision. This is because the head revision is accessed much more frequently than previous file revisions; if the head revision of a file had to be calculated from the deltas each time it was accessed, any SCM utilizing RCS format would run much more slowly.

## Use of “diff” to determine file revision differences

RCS uses the “GNU diff” algorithm to determine the differences between two versions of the same file; p4d contains its own diff routine which is used by Perforce servers to determine file differences when storing deltas.

Because Perforce’s diff always determines file deltas by comparing chunks of text between newline characters, it is by default only used with `text` files. If a file is `binary`, each revision is stored in full.

## Scheduling Resolves of Conflicting Files

---

Whenever a file revision is to be submitted that is not an edit of the file’s current head revision, there will be a file conflict, and this conflict must be resolved.

In slightly more technical terms, we’ll call the file revision that was read into a client workspace the *base file revision*. If the base file revision for a particular file in a client workspace is not the same as the head revision of the same file in the depot, a *resolve* must be performed before the new file revision can be accepted into the depot.

Before resolves can be performed with `p4 resolve`, they must be scheduled; this can be done with `p4 sync`. An alternative is to submit a changelist that contains the newly conflicting files. If a resolve is necessary, the submit fails, and the resolve is scheduled automatically.

## Why “p4 sync” to Schedule a Resolve?

---

Remember that the job of `p4 sync` is to project the state of the depot onto the client. Thus, when `p4 sync` is performed on a particular file:

- If the file does not exist in the client, or it is found in the client but is unopened, it is copied from the depot to the client.
- If the file has been deleted from the depot, it is deleted from the client.
- If the file has been opened in the client with `p4 edit`, the Perforce server can’t simply copy the file onto the client: any changes that had been made to the current revision of the file in the client would be overwritten. Instead, a *resolve* is scheduled between the file revision in the depot, the file on the client, and the base file revision (the revision that was last read into the client).

**Example:** *Scheduling resolves with `p4 sync`*

*Ed is making a series of changes to the \*.guide files in the elm doc subdirectory. He has retrieved the //depot/elm\_proj/doc/\*.guide files into his client and has opened the files with `p4 edit`. He edits the files, but before he has a chance to submit them, Lisa submits new*

*versions of some of the same files to the depot. The versions Ed has been editing are no longer the head revisions, and resolves must be scheduled and performed for each of the conflicting files before Ed's edits can be accepted.*

*Ed schedules the resolves with `p4 sync //edk/doc/*.guide`. Since these files are already open in the client, Perforce doesn't replace the client files. Instead, Perforce schedules resolves between the client files and the head revisions in the depot.*

*Alternatively, Ed could have submitted the `//depot/elm_proj/doc/*.guide` files in a changelist; the file conflicts would cause the `p4 submit` to fail, and the resolves would be scheduled as part of the submission failure.*

## How Do I Know When a Resolve is Needed?

---

`p4 submit` fails when it determines that any of the files in the submitted changelist need to be resolved, and the error message includes the names of the files that need resolution. If the changelist provided to `p4 submit` was the default changelist, it is assigned a number, and this number must be used in all future references to the changelist. (Numbered changelists are discussed in Chapter 7, Changelists)

Another way of determining whether a resolve is needed is to run `p4 sync -n filenames` before performing the submit, using the files in the changelist as arguments to the command. If file conflict resolutions are necessary, `p4 sync -n` reports them. The advantage of this approach is that the files in the default changelist remain in the default changelist (that is, the default changelist will not be reassigned to a numbered changelist).

## Performing Resolves of Conflicting Files

---

File conflicts are fixed with `p4 resolve [filenames]`. Each file provided as an argument to `p4 resolve` is processed separately. `p4 resolve` starts with three revisions of the same file and generates a fourth version; the user can accept any of these revisions in place of the current client file, and can edit the generated version before accepting it. The new revisions must then be submitted with `p4 submit`.

`p4 resolve` is interactive; a series of options are displayed for the user to respond to. The dialog looks something like this:

```
/usr/edk/elm/doc/answer.1 - merging //depot/elm_proj/doc/answer.1#5
Diff chunks: 4 yours + 2 theirs + 1 both + 1 conflicting
Accept(a) Edit(e) Diff(d) Merge (m) Skip(s) Help(?) [e]:
```

The remainder of this section explains what this means, and how to use this dialog.

## File revisions used and generated by “p4 resolve”

`p4 resolve [filenames]` starts with three revisions of the same file, generates a new version that merges elements of all three revisions, allows the user to edit the new file, and writes the new file (or any of the original three revisions) to the client. The file revisions used by `p4 resolve` are these:

<i>yours</i>	The newly-edited revision of the file in the client workspace. This file is overwritten by <i>result</i> once the resolve process is complete.
<i>theirs</i>	The revision in the depot that the client revision conflicts with. Usually, this is the head revision, but <code>p4 sync</code> can be used to schedule a resolve with any revision between the head revision and <i>base</i> .
<i>base</i>	The file revision in the depot that <i>yours</i> was edited from. Note that <i>base</i> and <i>theirs</i> are different revisions; if they were the same, there would be no reason to perform a resolve.
<i>merge</i>	File variation generated by Perforce from <i>theirs</i> , <i>yours</i> , and <i>base</i> .
<i>result</i>	The file resulting from the resolve process. <i>result</i> is written to the client workspace, overwriting <i>yours</i> , and must subsequently be submitted by the user. The instructions given by the user during the resolve process determine exactly what is contained in this file. The user can simply accept <i>theirs</i> , <i>yours</i> , or <i>merge</i> as the result, or can edit <i>merge</i> to have more control over the result.

The remainder of this chapter will use the terms *theirs*, *yours*, *base*, *merge*, and *result* to refer to the corresponding file revisions. The definitions given above are somewhat different when `resolve` is used to integrate *branched* files.

## Types of conflicts between file revisions

The diff program that underlies the Perforce resolve mechanism determines differences between file revisions on a line-by-line basis. Once these differences are found, they are grouped into *chunks*: for example, three new lines that are adjacent to each other are grouped into a single chunk. *Yours* and *theirs* are both generated by a series of edits to *base*; for each set of lines in *yours*, *theirs*, and *base*, `p4 resolve` asks the following questions:

- Is this line set the same in *yours*, *theirs*, and *base*?
- Is this line set the same in *theirs* and *base*, but different in *yours*?

- Is this line set the same in *yours* and *base*, but different in *theirs*?
- Is this line set the same in *yours* and *theirs*, but different in *base*?
- Is this line set different in all three files?

Any line sets that are the same in all three files don't need to be resolved. The number of line sets that answer the other four questions are reported by `p4 resolve` in this form:

```
2 yours + 3 theirs + 1 both + 5 conflicting
```

In this case, two line sets are identical in *theirs* and *base* but are different in *yours*; three line sets are identical in *yours* and *base* but are different in *theirs*; one line set was changed identically in *yours* and *theirs*; and five line sets are different in *yours*, *theirs*, and *base*.

## How the merge file is generated

`p4 resolve` generates a preliminary version of the *merge* file, which can be accepted as is, edited and then accepted, or rejected. A simple algorithm is followed to generate this file: any changes found in *yours*, *theirs*, or both *yours* and *theirs* are applied to the base file and written to the merge file; and any conflicting changes will appear in the merge file in the following format:

```
>>>> ORIGINAL VERSION foo#n
(text from the original version)
==== THEIR VERSION foo#m
(text from their file)
==== YOUR VERSION foo
(text from your file)
<<<<
```

Thus, editing the Perforce-generated merge file is often as simple as opening the merge file, searching for the difference marker “>>>>”, and editing that portion of the text. This isn't always the case, as it is often necessary to examine the changes made to *theirs* to make sure they're compatible with other changes that you made. This can be facilitated by calling `p4 resolve` with the `-v` flag; `p4 resolve -v` tells Perforce to generate difference markers for all changes made in either file being resolved, instead of only for changes that are in conflict between the *yours* and *theirs* files.

## The “p4 resolve” options

The `p4 resolve` command offers the following options:

Option	Short Meaning	What it Does
<code>e</code>	edit merged	Edit the preliminary merge file generated by Perforce
<code>ey</code>	edit <i>yours</i>	Edit the revision of the file currently in the client

Option	Short Meaning	What it Does
et	edit <i>theirs</i>	Edit the revision in the depot that the client revision conflicts with (usually the head revision). This edit is read-only.
dy	diff <i>yours</i>	Diff line sets from <i>yours</i> that conflict with <i>base</i>
dt	diff <i>theirs</i>	Diff line sets from <i>theirs</i> that conflict with <i>base</i>
dm	diff <i>merge</i>	Diff line sets from <i>merge</i> that conflict with <i>base</i>
d	diff	Diff line sets from <i>merge</i> that conflict with <i>yours</i>
m	merge	Invoke the command <code>P4MERGE base theirs yours merge</code> . To use this option, you must set the environment variable <code>P4MERGE</code> to the name of a third-party program that merges the first three files and writes the fourth as a result.
?	help	Display help for <code>p4 resolve</code>
s	skip	Don't perform the resolve right now.
ay	accept <i>yours</i>	Accept <i>yours</i> into the client workspace as the resolved revision, ignoring changes that may have been made in <i>theirs</i> .
at	accept <i>theirs</i>	Accept <i>theirs</i> into the client workspace as the resolved revision. The revision that was in the client workspace is overwritten.
am	accept <i>merge</i>	Accept <i>merge</i> into the client workspace as the resolved revision. The version originally in the client workspace is overwritten.
ae	accept edit	If you edited the <i>merge</i> file (by selecting "e" from the <code>p4 resolve</code> dialog), accept the edited version into the client workspace. The version originally in the client workspace is overwritten.
a	accept	If <i>theirs</i> is identical to <i>base</i> , accept <i>yours</i> , if <i>yours</i> is identical to <i>base</i> , accept <i>theirs</i> , if <i>yours</i> and <i>theirs</i> are different from <i>base</i> , and there are no conflicts between <i>yours</i> and <i>theirs</i> ; accept merge, otherwise, there are conflicts between <i>yours</i> and <i>theirs</i> , so skip this file

Only a few of these options are visible on the command line, but all options are always accessible and can be viewed by choosing `help`. The `merge` file is generated by `p4d`'s internal diff routine. But the differences displayed by `dy`, `dt`, `dm`, and `d` are generated by a diff internal to the Perforce client program, and this diff can be overridden by specifying an external diff in the `P4DIFF` environment variable.

The command line has the following format:

```
Accept (a) Edit (e) Diff (d) Merge (m) Skip (s) Help (?) [am] :
```

Perforce's recommended choice is displayed in brackets at the end of the command line. Pressing `Return` or choosing `Accept` will perform the recommended choice. The recommended command is chosen by Perforce by the following algorithm: if there were no changes to *yours*, accept *theirs*. If there were no changes to *theirs*, accept *yours*. Otherwise, accept `merge`.

**Example: Resolving File Conflicts**

*In the last example, Ed scheduled the `doc/*.guide` files for resolve. This was necessary because both he and Lisa had been editing the same files; Lisa had already submitted versions, and Ed needs to reconcile his changes with Lisa's. To perform the resolves, he types `p4 resolve //depot/elm_proj/doc/*.guide`, and sees the following:*

```
/usr/edk/elm/doc/Alias.guide - merging
//depot/elm_proj/doc/Alias.guide#5
Diff chunks: 4 yours + 2 theirs + 1 both + 1 conflicting
Accept (a) Edit (e) Diff (d) Merge (m) Skip (s) Help (?) [e] :
```

*This is the resolve dialog for `doc/Alias.guide`, the first of the four `doc/*.guide` files. Ed sees that he's made four changes to the base file that don't conflict with any of Lisa's changes. He also notes that Lisa has made two changes that he's unaware of. He types `dt` (for "display theirs") to view Lisa's changes; he looks them over and sees that they're fine. Of most concern to him, of course, is the one conflicting change. He types `e` to edit the Perforce-generated merge file and searches for the difference marker "`>>>>`". The following text is displayed:*

```
Glyphic Technology
Mountain View, California
>>>> ORIGINAL VERSION
==== THEIR VERSION
94041
==== YOUR VERSION
98041
<<<<
```



*He and Lisa have both tried to add a zip code to an address in the file, but Ed had typed it wrong. He edits this portion of the merge file so it reads as follows:*

```
Glyphic Technology
Mountain View, California
94041
```

*The merge file is now acceptable to him: he's viewed Lisa's changes, seen that they're compatible with his own, and the only line conflict has been resolved. He quits from the editor and types `am`; the edited merge file is written to the client, and the resolve process continues on the next `doc/*.guide` file.*

When a version of the file is accepted onto the client, the previous client file is overwritten, and the new client file must still be submitted to the depot. Note that it is possible for another user to have submitted yet another revision of the same file to the depot between the time `p4 resolve` completes and the time `p4 submit` is performed; in this case, it would be necessary to perform another resolve. This can be prevented by performing a `p4 lock` on the file before performing the resolve.

## Using Flags with Resolve to Automatically Accept Particular Revisions

Five optional `p4 resolve` flags tell the command to work non-interactively. When these flags are used, particular revisions of the conflicting files are automatically accepted.

"p4 resolve" flag	Meaning
<code>-ay</code>	Automatically accept <i>yours</i> .
<code>-at</code>	Automatically accept <i>theirs</i> . Use this option with caution, as the file revision in the client workspace will be overwritten with no chance of recovery.
<code>-am</code>	Automatically accept the Perforce-recommended file revision: If <i>theirs</i> is identical to <i>base</i> , accept <i>yours</i> , if <i>yours</i> is identical to <i>base</i> , accept <i>theirs</i> , if <i>yours</i> and <i>theirs</i> are different from <i>base</i> , and there are no conflicts between <i>yours</i> and <i>theirs</i> , accept merge, otherwise, there are conflicts between <i>yours</i> and <i>theirs</i> , so skip this file.

“p4 resolve” flag	Meaning
-af	Accept the Perforce-recommended file revision, no matter what. If this option is used, the resulting file in the client should be edited to remove any difference markers.
-as	If <i>theirs</i> is identical to <i>base</i> , accept <i>yours</i> ; if <i>yours</i> is identical to <i>base</i> , accept <i>theirs</i> ; otherwise skip this file.

**Example:** *Automatically accepting particular revisions of conflicting files*

*Ed has been editing the files in `doc/*.guide`, and knows that some of them will require resolving. He types `p4 sync doc/*.guide`, and all of these files that conflict with files in the depot are scheduled for resolve.*

*He then types `p4 resolve -am`, and the merge files for all scheduled resolves are generated, and those merge files that contain no line set conflicts are written to his client workspace.*

*He'll still need to manually resolve all the other conflicting files, but the amount of work he needs to do is substantially reduced.*

## Binary files and “p4 resolve”

If any of the three file revisions participating in the merge are binary instead of text, a three-way merge is not possible. Instead, `p4 resolve` performs a two-way merge: the two conflicting file versions are presented, and you can edit and choose between them.

## Locking Files to Minimize File Conflicts

---

Once open, a file can be locked with `p4 lock` so that only the user who locked the file can submit the next revision of that file to the depot. Once the file is submitted, it is automatically unlocked. Locked files can also be unlocked manually by the locking user with `p4 unlock`.

The clear benefit of `p4 lock` is that once a file is locked, the user who locked it experiences no further conflicts on that file, and will not need to resolve the file. However, this comes at a price, as other users will not be able to submit the file until the file is unlocked, and will have to do their own resolves once they submit their revision. Under most circumstances, a user who locks a file is essentially saying to other users “I don’t want to deal with any resolves; you do them.”

There is an exception to this rule: Perforce also has a `+1` file type modifier to support exclusive-open. If you have a `+1` file open for edit, other users who attempt to edit the file will receive an error message.

The difference between `p4 lock` and `+l` is that `p4 lock` allows anyone to open a file for edit, but only the person who locked the file may submit it. By contrast, a file of type `+l` allows only one person to open the file for edit in the first place.

## Preventing multiple resolves with `p4 lock`

Without file locking, there is no guarantee that the resolve process will ever end. The following scenario demonstrates the problem:

1. Ed opens file `foo` for edit.
2. Lisa opens the same file in her client for edit.
3. Ed and Lisa both edit their client workspace versions of `foo`.
4. Ed submits a changelist containing that file, and his submit succeeds.
5. Lisa submits a changelist with her version of the file; her submit fails because of file conflicts with the new depot's `foo`.
6. Lisa starts a resolve.
7. Ed edits and submits a new version of the same file.
8. Lisa finishes the resolve and attempts to submit; the submit fails and must now be merged with Ed's latest file.

...and so on

File locking can be used in conjunction with resolves to avoid this sort of headache. The sequence would be implemented as follows:

1. Before scheduling a resolve, lock the file.
2. Then sync the file, resolve the file, and submit the file.

As long as the locked file is open, new versions can't be submitted by other users until the resolved file is either submitted or unlocked.

## Preventing multiple checkouts with `+l` files

If you know in advance that a certain file is only going to be worked on by a single user, you can use the `+l` (exclusive-open) filetype modifier to ensure that only one user at a time may work on the file.

You can change a file's type to exclusive-open by reopening it with the `+l` modifier. For instance:

```
p4 reopen -t binary+l file.gif
```

Although this prevents concurrent development, for some file types (usually binary files), merging and resolving may not be meaningful, and you may decide that you prefer to allow only one user to check these files out at a time.

If you find this style of work to be useful, you may want to ask your Perforce administrator to use the `p4 typemap` command (documented in the *Perforce Command Reference*) to ensure that *all* files matching a file specification (for instance, `//depot/.../*.gif` for all `.gif` files) are assigned type `+1` by default.

## Resolves and Branching

---

Files in separate codelines can be integrated with `p4 resolve`; for details about resolving branched files, see Chapter 9, Branching.

## Resolve Reporting

---

Four reporting commands are related to file conflict resolution: `p4 diff`, `p4 diff2`, `p4 sync -n`, and `p4 resolved`.

Command	Meaning
<code>p4 diff [filenames]</code>	<p>Runs a diff program between the file revision currently in the client and the revision that was last gotten from the depot. If the file is not open for edit in the client, the two file revisions should be identical, so <code>p4 diff</code> fails. Comparison of the revisions can be forced with <code>p4 diff -f</code>, even when the file in the client is not open for edit.</p> <p>Although <code>p4 diff</code> runs a diff routine internal to Perforce, this routine can be overridden by specifying an external diff in the <code>P4DIFF</code> environment variable.</p>
<code>p4 diff2 file1 file2</code>	<p>Runs <code>p4d</code>'s diff subroutine on any two Perforce depot files. The specified files can be any two file revisions, even revisions of entirely different files.</p> <p>The diff routine used by <code>p4d</code> cannot be overridden.</p>
<code>p4 sync -n [filenames]</code>	<p>Reports what the result of running <code>p4 sync</code> would be, without actually performing the sync. This is useful to see which files have conflicts and need to be resolved.</p>
<code>p4 resolved</code>	<p>Reports which files have been resolved but not yet submitted.</p>

The reporting chapter (Chapter 11, Reporting and Data Mining) has a longer description of each of these commands. `p4 help` provides a complete listing of the many flags for these reporting commands.



---

# Perforce Basics: Miscellaneous Topics

---

The manual thus far has provided an introduction to the basic functionality provided by Perforce, and subsequent chapters cover the more advanced features. In between are a host of other, smaller facilities; this chapter covers these topics. Included here is information on the following:

- Changing your Perforce environment with the `P4CONFIG` environment variable,
- Using passwords to prevent impersonation by other users,
- Command-line flags common to all Perforce commands,
- How to work on files while not connected to a Perforce server,
- Refreshing the client workspace,
- Renaming files, and
- Recommendations for organization of files within the depot.

---

## Reconfiguring the Perforce Environment with `$P4CONFIG`

---

Some Perforce users have multiple client workspaces, each of which may connect to a different port. There are three ways of changing your Perforce environment on the fly:

- Reset your environment or registry variables each time you want to move to a new workspace.
- Use command-line flags (discussed in the next section) to override the value of the environment variables `P4PORT`, `P4CLIENT`, and `P4USER`.
- Use the environment variable or registry variable `P4CONFIG` to point to a file containing a specification for the current Perforce environment.

`P4CONFIG` names a file (for example, `.p4env`) that is used to store variable settings. Whenever a Perforce command is executed, the current working directory and its parent directories are searched for a file with the name stored in `P4CONFIG`. If a file with that name is found, the values of `P4PORT`, `P4CLIENT`, and so on, are read from that file. If no file of the given name is found, the current values of the Perforce environment variables are used.

Each variable setting in the file stands alone on a line and must be in the form:

```
P4VAR=value
```

Values that can be stored in the `P4CONFIG` file are:

<code>P4CLIENT</code>	<code>P4DIFF</code>	<code>P4EDITOR</code>	<code>P4HOST</code>
<code>P4MERGE</code>	<code>P4PASSWD</code>	<code>P4PORT</code>	<code>P4USER</code>

**Example:** *Using `P4CONFIG` to automatically reconfigure the Perforce environment*

*Ed often switches between two workspaces on the same machine. The first workspace is `elmproj`. It has a client root of `/usr/edk/elm`, and connects to the Perforce server at `ida:1818`. The second workspace is called `graphicwork`. Its client root is `/usr/edk/other/graphics`, and it uses the Perforce server at `warhol:1666`.*

*Ed sets the `P4CONFIG` environment variable to `.p4settings`. He creates a file called `.p4settings` in `/usr/edk/elm` containing the following text:*

```
P4CLIENT=elmproj
P4PORT=ida:1818
```

*He creates a second `.p4settings` file in `/usr/edk/other/graphics`. It contains:*

```
P4PORT=warhol:1666
P4CLIENT=graphicwork
```

*He always works within the directories where his files are located. Whenever Ed is anywhere beneath `/usr/edk/other/graphics`, his Perforce client is `graphicwork`, and when he's in `/usr/edk/elmproj`, his client is `elmproj`.*

The values found in the file specified by `P4CONFIG` override any environment or registry variables that have been set. Command-line flags (discussed in the next section) override the values found in the `P4CONFIG` file.

`P4CONFIG` automates the process of changing the Perforce environment variables, as new settings take effect whenever you switch your current working directory from one client workspace directory to another. If you use multiple client workspaces, this is a very useful feature.

## Perforce Passwords

---

By default, any user can assume the identity of any other Perforce user by setting the value of `P4USER` to a different username, by using the `-u` flag with the `p4` command, or by setting the value of `P4USER` configuration within the file described by `P4CONFIG`. Although this makes it easy to perform tasks for other users when necessary, it can also lead to security problems.

To prevent another user from impersonating you within Perforce, use the `p4 passwd` command. No one, including the user who set the password, will be able to use any `p4` command without providing the password to Perforce. You can provide this password to Perforce in one of three ways:



1. Set the value of the environment or registry variable `P4PASSWD` to the correct password.
2. Use the `-P password` flag between `p4` and the actual command when giving a Perforce command (for instance, `p4 -P eds_password submit`).
3. Set the value of `P4PASSWD` within the file described by `P4CONFIG`.

For security's sake, the password should always be set while logged onto the server.

Be careful when setting passwords. Once you have set your Perforce password with `p4 user`, there is no way for you to use Perforce if you forget their password and the value of `P4PASSWD` has not been properly set. If this happens, the Perforce superuser will have to reset or remove your password.

**Note** | Passwords can also be created, changed, and deleted within the `p4 user` form.

If you need to have your password reset, contact your Perforce administrator. If you are a Perforce administrator, consult the *Perforce System Administrator's Guide* for information on resetting passwords and other common user-related tasks.

## Command-Line Flags Common to All Perforce Commands

Six flags are available for use with all Perforce commands. These flags are given between the system command `p4` and the command argument taken by `p4`. These flags are:

Flag	Meaning	Example
<code>-c clientname</code>	Runs the command on the specified client. Overrides the <code>P4CLIENT</code> environment variable.	<code>p4 -c joe edit //depot/foo</code> Opens file <code>foo</code> for editing under client workspace <code>joe</code> .
<code>-d directory</code>	Specifies the current directory, overriding the environment variable <code>PWD</code> .	<code>p4 -d ~elm/src edit foo bar</code> Opens files <code>foo</code> and <code>bar</code> for edit; these files are found relative to <code>~elm/src</code> .
<code>-p server_addr</code>	Gives the Perforce server's listening address, overriding <code>P4PORT</code> .	<code>p4 -p mama:1818 clients</code> Reports a list of clients on the server on host <code>mama</code> , port <code>1818</code> .

Flag	Meaning	Example
<code>-P password</code>	Supplies a Perforce password, overriding the value of <code>P4PASSWD</code> . Usually used in combination with the <code>-u user</code> flag.	<code>p4 -u ida -P idas_pw job</code> Create a new job as user <code>ida</code> , using <code>ida</code> 's Perforce password.
<code>-u username</code>	Specifies a Perforce user, overriding the <code>P4USER</code> environment variable.  The user may run only those commands to which he or she has access.	<code>p4 -u bill user</code> Presents the <code>p4 user</code> form to edit the specification for user <code>bill</code> . The command will work without the <code>-P</code> flag only if <code>bill</code> has not set a Perforce password.
<code>-x filename</code>	Instructs <code>p4</code> to read arguments, one per line, from the named file.	See the <i>Working Detached</i> section, directly below.

All Perforce commands can take these flags, even commands for which these flag usages are clearly useless (for instance, `p4 -u bill -d /usr/joe help`). Other flags are available as well; these additional flags are command dependent. Please consult the *Perforce Command Reference* or use `p4 help commandname` to see the flags for each command.

## Working Detached

---

Under normal circumstances, you work in your client workspace with a functioning network connection to a Perforce server. As you edit files, you announce your intentions to the server with `p4 edit`, and the server responds by noting the edit in the depot's metadata, and by unlocking the file in your client workspace. However, it is not always possible you to have a network connection to be present, and you may need a method for working entirely detached from the server.

The scheme is as follows:

- Work on files without giving Perforce commands. Instead, you use native OS commands to manually change the permissions on files, and then edit or delete the files.
- If you did not edit the files within the client workspace, copy them to the client workspace when the network connection is reestablished.
- Use `p4 diff` to find all files in your workspace that have changed or been added without Perforce's knowledge. Use the output from this command to bring the depot in sync with the client workspace.

## Finding changed files with “p4 diff”

Use the `p4 diff` reporting command to compare a file in the client workspace with the corresponding file in the depot. The behavior of `p4 diff` can be modified with two flags:

"p4 diff" Variation	Meaning
<code>p4 diff -se</code>	Tells the names of unopened files that are present on the client, but whose contents are different than the files last taken by the client with <code>p4 sync</code> . These files are candidates for <code>p4 edit</code> .
<code>p4 diff -sd</code>	Reports the names of unopened files missing from the client. These files are candidates for <code>p4 delete</code> .

**Note** You can use `p4 edit` on any file, even files you don't want to edit; this command gives the local file write permissions, but does not otherwise alter it.

## Using “p4 diff” to update the depot

The `p4 diff` variations described above can be used in combination with the `-x` flag to bring the state of the depot in sync with the changes made to the client workspace.

To open changed files for edit after working detached, use:

```
p4 diff -se > CHANGED_FILES
p4 -x CHANGED_FILES edit
```

To delete files from the depot that were removed from the client workspace, use:

```
p4 diff -sd > DEL_FILES
p4 -x DEL_FILES delete
```

As always, your edit and delete requests are stored in changelists, which Perforce does not process until you type `p4 submit`.

## Refreshing files

The process of syncing a depot with a formerly detached client workspace has a converse: Perforce can get confused about the contents of a client workspace if you accidentally use the local OS file deletion command.

For example, suppose that you accidentally delete a client workspace file via the UNIX `rm` command, and that the file is one that you wanted to keep. Even after a submit, `p4 have` will still list the file as being present in the workspace.

In such a situation, you can use `p4 sync -f files` to bring the client workspace in sync with the files the depot *thinks* you have. This command is mostly a recovery tool for bringing the client workspace back into sync with the depot after accidentally removing or damaging files managed by Perforce.

## Recommendations for Organizing the Depot

---

The default view brought up by `p4 client` maps the entire depot to the entire client workspace. If the client workspace is named `eds_elm`, the default view looks like this:

```
//depot/... //eds_elm/...
```

This is the easiest mapping, and can be used for the most simple Perforce depots, but mapping the entire depot to the workspace can lead to problems later on. Suppose your server currently stores files for only one project, but another project is added later: everyone who has a client workspace mapped as above will wind up receiving all the files from both projects into their workspaces. Additionally, the default view does not facilitate branch creation.

The safest way to organize the depot, even from the start, is to create one subdirectory per project within the depot. For example, if your company is working on three projects named `foo`, `bar`, and `zeus`, three subtrees might be created within the depot:

`//depot/foo`, `//depot/bar`, and `//depot/zeus`. If Joe is working on the `foo` project, his mapping might look like this:

```
//depot/foo/... //joe/...
```

And Sarah, who's working on the `bar` and `zeus` projects, might set up her client workspace as:

```
//depot/bar/... //sarah/bar/...  
//depot/zeus/... //sarah/zeus/...
```

This sort of organization can be easily extended to as many projects and branches as are needed.

Another way of solving the same problem would be to have the Perforce system administrator create one depot for each project or branch. Please see the *Perforce System Administrators Guide* for details about setting up multiple depots.

**Note** | The depot is divided into subdirectories simply by setting up the proper mappings within the client views.

## Renaming Files

Although Perforce doesn't have a rename command, you can rename files by using `p4 integrate` to copy the file from one location in the depot to another, deleting the file from the original location, and then submitting the changelist that includes the integrate and the delete. The process is as follows:

```
p4 integrate from_files to_files
p4 delete from_files
p4 submit
```

The *from\_file* is moved to the directory and renamed according to the *to\_file* specifier. For example, if *from\_file* is `d1/foo` and *to\_file* is `d2/bar`, then `foo` is moved to the `d2` directory, and is renamed `bar`. The *from\_file* and *to\_file* specifiers may include wildcards, as long as they are matched on both sides. Perforce `write` access (explained in the *Perforce System Administrator's Guide*) is needed on all the specified files.

### Revision histories and renamed files

When you rename a file (or move it from one directory to another) with `p4 integrate`, you create the new file by creating an integration record that links it to its deleted predecessor.

As a result, if you run `p4 changes newfile`, you'll see only changes to *newfile*. Only changes that have taken place after the renaming will be listed:

```
$ p4 changes newfile.c
Change 4 on 2000/10/24 by user@client 'Latest bugfix'
Change 3 on 2000/10/23 by user@client 'Renamed file'
```

In order to see the *full* history of changes to the file (including changes made before the file was renamed or moved), you must specify the `-i` (include changes from integrations) flag to `p4 changes`, as follows:

```
$ p4 changes -i newfile.c
Change 4 on 2000/10/24 by user@client 'Latest bugfix'
Change 3 on 2000/10/23 by user@client 'Renamed file'
Change 2 on 2000/10/21 by user@client 'second version'
change 1 on 2000/10/20 by user@client 'initial check-in'
```

Specifying the `-i` flag tells `p4 changes` to trace back through the integration records and retrieve all change information, regardless of how many times the file (or the directory in which it lives) has been renamed or moved.



A Perforce *changelist* is a list of files, their revision numbers, and operations to be performed on these files. Commands such as `p4 add filenames` and `p4 edit filenames` include the affected files in a changelist; the depot is not actually altered until the changelist is submitted with `p4 submit`.

When a changelist is submitted to the depot, the depot is updated *atomically*: either all of the files in the changelist are updated in the depot, or none of them are. This grouping of files as a single unit guarantees that code alterations spanning multiple files are updated in the depot simultaneously. To reflect the atomic nature of changelist submissions, submission of a changelist is sometimes called an *atomic change transaction*.

Perforce attempts to make changelist usage as transparent as possible: in the normal case, Perforce commands such as `p4 edit` add the affected files to the *default changelist*, and `p4 submit` sends the default changelist to the server for processing. However, there are two sets of circumstances that would require the user to understand and manipulate non-default changelists:

- Sometimes a user wants to split files into separate groups for submission.

For example, suppose a user is fixing two bugs, each of which spans a separate set of files. Rather than submit the fixes to both bugs in a single changelist, the user might elect to create one changelist for the files that fix the first bug, and another changelist for the files that fix the second bug. Each changelist would be submitted to the depot via separate `p4 submit` operations.

- Under certain circumstances, the `p4 submit` command can fail.

For example, if one user has a file locked and another user submits a changelist that contains that file, the submit fails. When a submit of the default changelist fails, the changelist is assigned a number, is no longer the default changelist, and must be referred to by its number.

In the above circumstances, the user must understand how to work with *numbered changelists*.

## Working with the Default Changelist

---

**Note** | The material in this subsection has already been presented in a slightly different form in earlier chapters. It is presented again here to provide a complete discussion of changelists.

A changelist is a list of files, revision numbers of those files, and operations to be performed on those files. For example, a single changelist might contain the following:

```
/doc/elm-help.1      revision 3      edit
/utils/elmalias.c   revision 2      delete
```

Each of the files in the changelist is said to be *open* within the client workspace: the first of the files above was opened for edit with `p4 edit`, and the second was opened for deletion with `p4 delete`.

The files in the changelist are updated within the depot with `p4 submit`, which sends the changelist to the server; the server processes the files contained in the changelist and alters the depot accordingly.

The commands that add or remove files from changelists are:

```
p4 add                p4 delete          p4 edit
p4 integrate         p4 reopen          p4 revert
```

By default, these commands, and `p4 submit`, act on the default changelist. For example, if you type `p4 add filename`, this file is added to the default changelist. When you type `p4 submit`, the default changelist is submitted.

When you type `p4 submit`, a change form is displayed, containing the files in the default changelist. Any file can be deleted from this list; when a file is deleted, it is moved to the next default changelist, and will appear again the next time you type `p4 submit`. A changelist must contain a user-entered description, which describes the nature of the changes being made.

`p4 submit` can take an optional, single file pattern as an argument. In this case, only those files in the default changelist that match the file pattern are included in the submitted changelist. Since the `p4d` server program must receive this file pattern as a single argument, make sure to escape the `*` wildcard if it is used.

When the user quits from the `p4 submit` editor, the changelist is submitted to the server and the server attempts to update the files in the depot. If there are no problems, the changelist is assigned a sequential number, and its status changes from *new* or *pending* to *submitted*. Once a changelist has been submitted, it becomes a permanent part of the depot's metadata, and is unchangeable except by Perforce superusers.



## Creating Numbered Changelists Manually

A user can create a changelist in advance of submission with `p4 change`. This command brings up the same form that you see during `p4 submit`.

All files in the default changelist are moved to this new changelist. When you quit from the form, the changelist is assigned the next changelist number in sequence, and this changelist must be subsequently referred to by this change number. Files can be deleted from the changelist by editing the form; files deleted from this changelist are moved to the next default changelist. The status for a changelist created by this method is `pending` until you submit the files in the changelist.

Any single client file can be included in only one pending changelist.

## Working With Numbered Changelists

You can use commands such as `p4 edit filename`, which by default adds the files to the default changelist, to append a file to a pending numbered changelist with the `-c changenum` flag. For example, to edit a file and submit it in changelist number 4, use `p4 edit -c 4 filename`.

You can move files from one changelist to another with `p4 reopen -c changenum filenames`, where `changenum` is the number of the moving-to changelist. If you are moving files to the default changelist, use `p4 reopen -c default filenames`.

**Example:** *Working with multiple changelists.*

*Ed is working on two bug fixes simultaneously. One of the bugs involves mail filtering and requires updates of files in the `filter` subdirectory; the other problem is in the aliasing system, and requires an update of `utils/elmalias.c`.*

*Ed wants to update each bug separately in the depot; this will allow him to refer to one bug fix by one change number and the other bug fix by another change number. He's already started fixing both bugs, and has opened some of the affected files for edit. He types `p4 change`, and sees*

```
Change: new
Client: eds_elm
User: edk
Status: new
Description:
    <enter description here>
Files:
    //depot/elm_proj/filter/filter.c # edit
    //depot/elm_proj/filter/lock.c # edit
    //depot/elm_proj/utils/elmalias.c # edit
```

*Ed wants to use this changelist to submit only the fix to the filter problems. He changes the form, deleting the last file revision from the file list. When he's done, the form looks like this:*

```
Change: new
Client: eds_elm
User: edk
Status: new
Description:
    Fixes filtering problems
Files:
    //depot/elm_proj/filter/filter.c # edit
    //depot/elm_proj/filter/lock.c # edit
```

*When he quits from the editor, he's told*

Change 29 created with 2 open file(s).

*The file that he removed from the list, utils/elmaliases.c, is now found in the default changelist. He could include that file in another numbered changelist, but decides to leave it where it is.*

*He fixes both bugs at his leisure. He realizes that the filter problem requires updates to another file: filter/lock.c. He opens this file for edit with `p4 edit -c 29 filter/lock.c`; opening the file with the `-c 29` flag puts the file in changelist 29, which he created above. (If the file had already been open for edit in the default changelist, he could have moved it to changelist 29 with `p4 reopen -c 29 filter/lock.c`).*

*Ed finishes fixing the aliasing bug, and, since the affected files are in the default changelist, he submits the changelist with a straightforward `p4 submit`. He'll finish fixing the filtering bug later.*

## Automatic Creation and Renumbering of Changelists

---

### When submit of the default changelist fails, the changelist is assigned a number

Submits of changelists occasionally fail. This can happen for a number of reasons:

- A file in the changelist has been locked by another user with `p4 lock`.
- The client workspace no longer contains a file included in the changelist.
- There is a server error, such as not enough disk space.
- The user was not editing the head revision of a particular file. The following sequence shows an example of how this can occur:
  - User A types `p4 edit //depot/foo`,
  - User B types `p4 edit //depot/foo`,

- User B submits her default changelist, and
- User A submits his default changelist.

User A's submit is rejected, since the file revision of `foo` that he edited is no longer the head revision of that file.

If any file in a changelist is rejected for any reason, the entire changelist is backed out, and none of the files in the changelist are updated in the depot. If the submitted changelist was the default changelist, Perforce assigns the changelist the next change number in sequence, and this change number must be used from this point on to refer to the changelist. Perforce also locks the files to prevent others from changing them while the user resolves the reason for the failed submit.

If the submit failed because the client-owned revision of the file is not the head revision, a *merge* must be performed before the changelist will be accepted. (File merging is described in Chapter 5, Perforce Basics: Resolving File Conflicts).

## Perforce May Renumber a Changelist upon Submission

---

The change numbers of submitted changelists always reflect the order in which the changelists were submitted. Thus, when a changelist is submitted, it may be renumbered.

**Example:** *Automatic renumbering of changelists*

*Ed has finished fixing the filtering bug that he's been using changelist 29 for. Since he created that changelist, he's submitted another changelist (change 30), and two other users have submitted changelists. Ed submits change 29 with `p4 submit -c 29`, and is told:*

*Change 29 renamed change 33 and submitted.*

## Deleting Changelists

---

To remove a pending changelist that has no files or jobs associated with it, use `p4 change -d changenum`. Pending changelists that contain open files or jobs must have the files and jobs removed from them before they can be deleted: use `p4 reopen` to move files to another changelist, `p4 revert` to remove files from the changelist (and revert them back to their old versions), or `p4 fix -d` to remove jobs from the changelist.

Changelists that have already been submitted can be deleted by the superuser only under very specific circumstances. Please see the *Perforce System Administrator's Guide* for more information.

## Changelist Reporting

---

The two reporting commands associated with changelists are `p4 changes` and `p4 describe`. The former is used to view lists of changelists with short descriptions, while the latter prints verbose information for a single changelist.

Command	Meaning
<code>p4 changes</code>	Displays a list of all pending and submitted changelists, one line per changelist, and an abbreviated description.
<code>p4 changes -m count</code>	Limits the number of changelists reported on to the last <i>count</i> changelists.
<code>p4 changes -s status</code>	Limits the list to those changelists with a particular status; for example, <code>p4 changes -s submitted</code> will list only already submitted changelists.
<code>p4 changes -u user</code>	Limits the list to those changelists submitted by a particular user.
<code>p4 changes -c workspace</code>	Limits the list to those changelists submitted from a particular client workspace.
<code>p4 describe changenum</code>	Displays full information about a single changelist. If the changelist has already been submitted, the report includes a list of affected files and the diffs of these files. (You can use the <code>-s</code> flag to exclude the file diffs.)

---

## Chapter 8    **Labels**

---

A Perforce *label* is a user-determined list of files and revisions. The label can later be used to reproduce the state of these files within a client workspace.

Labels provide a method of naming important combinations of file revisions for later reference. For example, the file revisions that comprise a particular release of your software might be given the label `release2.0.1`. At a later time, you can retrieve all the files in that label into a client workspace with a single command.

Create a label when:

- You want to keep track of all the file revisions contained in a particular release of the software,
- There exists a particular set of file revisions that you want to give to other users, or
- You have a set of file revisions that you want to branch from, but you don't want to perform the branch yet. In this case, you would create a label for the file revisions that will form the base of the branch.

### **Why Not Just Use Changelist Numbers?**

---

Labels share certain important characteristics with change numbers, as both refer to particular file sets, and both act as handles to refer to all the files in the set. But labels have some important advantages over change numbers:

- The file revisions referenced by a particular label can come from different changelists.
- A change number refers to the state of all the files in the depot at the time the changelist was submitted, while a label can refer to any arbitrary set of files and revisions.
- The files and revisions referenced by a label can be arbitrarily changed at any point in the label's existence.
- Changelists are always referred to by Perforce-assigned numbers, while labels are named by the user.

### **Creating a Label**

---

You can create a label with `p4 label labelname`. This command brings up a form similar to the `p4 client` form. Like clients, labels have associated views; the label view limits which files can be referenced by the label. Once you have created the label, use `p4 labelsync` to load the label with file references.

Label names share the same namespace as clients, branches, and depots. Thus, a label name cannot be the same as any existing client, branch, or depot name.

**Example: Creating a Label**

*Ed has finished the first version of filtering in elm. He wants to create a label that references only those files in the filter and hdrs subdirectories. He wants to name the label filters.1. He types p4 label filters.1 and fills in the resulting form as follows:*

```
Label: filters.1
Owner: edk
Description:
    Created by edk.
Options:      unlocked
View:
    //depot/elm_proj/filter/...
    //depot/elm_proj/hdrs/...
```

*When he quits from the editor, the label is created.*

Before following this example further, it's worth stopping for a moment to examine exactly what has and hasn't been accomplished. So far, a label called filters.1 has been created. It can contain files only from the depot's elm\_proj filter and hdrs subdirectories. But the label filters.1 is empty; it contains no file references. It will be loaded with its file references with p4 labelsync.

The View: field limits the files that are included in the label. These files must be specified by their location in the depot; this view differs from other views in that only the depot side of the view is specified.

The locked / unlocked option in the Options: field can prevent p4 labelsync from overwriting previously synced labels (this is described in "Preventing Accidental Overwrites of a Label's Contents" on page 88).

## Adding and Changing Files Listed in a Label

---

Once you've created a label, you can include references to files within it by using p4 labelsync. This command's syntax is

```
p4 labelsync [-a -d -n] -l labelname [filename...]
```

The rules followed by labelsync to include files in a label are as follows:

1. You must be the owner of the label in order to use p4 labelsync on it, and the label must be unlocked.

If you are not the owner of a label, you may (assuming you have sufficient permissions) make yourself the owner by running:

```
p4 label labelname
```

and changing the `Owner:` field to your `userid` in the `p4 label` form. Similarly, you may unlock a label by setting the `Options:` field (also in the `p4 label` form) to `unlocked`.

2. All files listed in a label must be contained in the label view specified in the `p4 label` form. Any files or directories that are not mapped through the label view are ignored by `p4 labelsync`. All the following rules assume this, without further mention.
3. When `p4 labelsync` is used to include a particular file in a label's file list, the file is added to the label if it is not already included in the label. If a different revision of the file is already included in the label's file list, it is replaced with the newly specified revision. Only one revision of any file is ever included in a label's file list.
4. If `labelsync` is called with no filename arguments, as in:

```
p4 labelsync -l labelname
```

then all the files mapped by the label view are listed in the label. The revisions added to the label are those last synced into the client; these revisions can be seen in the `p4 have` list. Calling `p4 labelsync` this way replaces *all* existing file references with the new ones.

**Example:** *Including file references in a label with `p4 labelsync`.*

*Ed has created a label called `filters.1` as specified above, and now he wants to load the `filters.1` label with the proper file revisions. He types:*

```
p4 labelsync -l filters.1
```

*and sees the following:*

```
//depot/elm_proj/filter/Makefile.SH#20 - added
//depot/elm_proj/filter/actions.c#25 - added
<etc.>
```

*Perforce adds those file revisions to the label that are those contained in the intersection of the label view and the current client `have` list.*

5. When you call `p4 labelsync` with file pattern arguments, and the arguments contain no revision specifications, the head revisions of these files are included in the label's file list.

**Example:** *Changing file references in a label with `p4 labelsync`.*

*After performing the above `labelsync` command, Ed finds that the file `filter/filter.c` is buggy. He fixes it, submits the new version, and wants to replace the old revision of this file in the label `filters.1` with the new revision. From the `filter` subdirectory, he types*

```
p4 labelsync -l filters.1 filter/filter.c
```

*and sees*

```
//depot/elm_proj/filter/filter.c#15 - updated
```

*The head revision of `filter.c` replaces the revision that had been previously included in the label.*

6. If you call `p4 labelsync` with file pattern arguments containing revision specifications, these file revisions are included in the label's file list.

**Example:** *Including a different file revision in a label*

*Ed realizes that the version of `filter/audit.c` contained in his label `filters.1` is not the version he wants to include in the label. He'd prefer to include revision 12 of that file. From the main `Elm` directory, he types*

```
p4 labelsync -l filters.1 filter/audit.c#12
```

*and sees:*

```
/depot/elm_proj/filter/audit.c#12 - updated
```

*This revision of `audit.c` replaces the revision that had been previously included in the label.*

## Previewing `labelsync`'s results

The results of `p4 labelsync` can be previewed with `p4 labelsync -n`. This lists the files that would be added, deleted, or replaced in the label without actually performing the operation.

## Preventing Accidental Overwrites of a Label's Contents

---

Since `p4 labelsync` with no file arguments overwrites all the files that are listed in the label, it is possible to accidentally lose the information that a label is meant to contain. To prevent this, call `p4 label labelname` and set the value of the `Options:` field to `locked`. It will be impossible to call `p4 labelsync` on that label until the label is subsequently unlocked.



## Retrieving a Label's Contents into a Client Workspace

To retrieve all the files listed in a label into a client workspace, use `p4 sync files@labelname`. Rather than simply adding the files to the client workspace, this command matches the state of the client workspace to the state of the label. Thus, files in the client workspace that aren't in the label may be deleted from the client workspace.

**Example:** *Retrieving files into a client workspace from a label.*

*Lisa wants to retrieve all the files listed in Ed's `filters.1` label into her client workspace. She can type:*

```
p4 sync //depot/...@filters.1
```

*or even:*

```
p4 sync @filters.1
```

*Instead, she's interested in seeing only the header files from that label, so she types:*

```
p4 sync //depot/elm_proj/hdrs/*@filters.1
```

*and sees:*

```
//depot/elm_proj/hdrs/curses.h#1 - added as /usr/lisag/elm/hdrs/curses.h
//depot/elm_proj/hdrs/defs.h#1 - added as /usr/lisag/elm/hdrs/defs.h
//depot/elm_proj/hdrs/test.h#3 - deleted as /usr/lisag/elm/hdrs/test.h
<etc>
```

*All the files in the subdirectory `hdrs` that are within the intersection of Ed's `filters.1` label and Lisa's client view are retrieved into her workspace. But there is another effect as well: files that are not in the intersection of the label's contents and `//depot/elm_proj/hdrs/*` are deleted from her workspace.*

If `p4 sync @labelname` is called with no file parameters, all files in the client that are not in the label will be deleted from the client. If this command is called with file arguments, as in `p4 sync files@labelname`, then the client workspace at the intersection of the depot, the client workspace view, and the file arguments will be updated to match the contents of the depot at that intersection.

## Deleting Labels

You can delete a label with:

```
p4 label -d labelname
```

You can delete files from labels with:

```
p4 labelsync -d -l labelname filepattern
```

To delete all the files from the label's file list, but leave the empty label in the Perforce database, use:

```
p4 labelsync -d -l labelname
```

## Label Reporting

---

The commands that provide reports on labels are:

Command	Description
<code>p4 labels</code>	Report the names, dates, and descriptions of all labels known to the server
<code>p4 labels <i>file#revrange</i></code>	Report the names, dates, and descriptions of all labels that include the specified revision range of <i>file</i> .
<code>p4 files @<i>labelname</i></code>	Lists all files and revisions contained in the given label.
<code>p4 sync -n @<i>labelname</i></code>	Shows what <code>p4 sync</code> would do when retrieving files from a particular label into your client workspace, without actually performing the sync.

---

## Chapter 9    **Branching**

---

Perforce's *Inter-File Branching*<sup>™</sup> mechanism allows any set of files to be copied within the depot, and allows changes made to one set of files to be copied, or *integrated*, into another. By default, the new file set (or *codeline*) evolves separately from the original files, but changes in either codeline can be propagated to the other with the `p4 integrate` command.

### **What is Branching?**

---

Branching is a method of keeping two or more sets of similar (but not identical) files synchronized. Most software configuration management systems have some form of branching; we believe that Perforce's mechanism is unique because it mimics the style in which users create their own file copies when no branching mechanism is available.

Suppose for a moment that you're writing a program and are not using an SCM system. You're ready to release your program: what would you do with your code? Chances are that you'd copy all your files to a new location. One of your file sets would become your release codeline, and bug fixes to the release would be made to that file set; your other file files would be your development file set, and new functionality to the code would be added to these files.

What would you do when you find a bug that's shared by both file sets? You'd fix it in one file set, and then copy the edits that you made into the other file set.

The only difference between this homegrown method of branching and Perforce's branching methodology is that Perforce *manages the file copying and edit propagation for you*. In Perforce's terminology, copying the files is called *making a branch*. Each file set is known as a *codeline*, and copying an edit from one file set to the other is called *integration*. The entire process is called *branching*.

### **When to Create a Branch**

---

Create a branch when two sets of code files have different rules governing when code can be submitted, or whenever a set of code files needs to evolve along different paths. For example:

- The members of the development group want to submit code to the depot whenever their code changes, whether or not it compiles, but the release engineers don't want code to be submitted until it's been debugged, verified, and signed off on.

They would branch the release codeline from the development codeline. When the development codeline is ready, it is integrated into the release codeline. Afterwards, patches and bug fixes are made in the release code, and at some point in the future, integrated back into the development code.

- A company is writing a driver for a new multi-platform printer. It has written a UNIX device driver, and is now going to begin work on a Macintosh driver using the UNIX code as their starting point.

The developers create a branch from the existing UNIX code, and now have two copies of the same code. These two codelines can then evolve separately. If bugs are found in either codeline, bug fixes can be propagated from one codeline to the other with the Perforce `p4 integrate` command.

One basic strategy is to develop code in `//depot/main/` and create branches for releases (for example, `//depot/re11.1/`). Make bug fixes that affect both codelines in `//depot/main/` and integrate them into the release codelines. Make release-specific bug fixes in the release branch and, if required, integrate them back into the `//depot/main/` codeline.

## Perforce's Branching Mechanisms: Introduction

---

Perforce provides two mechanisms for branching. One method requires no special setup, but requires the user to manually track the mappings between the two sets of files. The second method remembers the mappings between the two file sets, but requires some additional work to set up.

In the first method, the user specifies both the files that changes are being copied from and the files that the changes are being copied into. The command looks like this:

```
p4 integrate fromfiles tofiles
```

In the second method, Perforce stores a mapping that describes which set of files get branched to other files, and this mapping, or branch specification, is given a name. The command the user runs to copy changes from one set of files to the other looks like this:

```
p4 integrate -b branchname [tofiles]
```

These methods are described in the following two sections.

## Branching and Merging, Method 1: Branching with File Specifications

Use `p4 integrate fromfiles tofiles` to propagate changes from one set of files (the *source files*) to another set of files (the *target files*). The target files need to be contained within the current client workspace view. The source files do not need to be, so long as the source files are specified in depot syntax. If the target files do not yet exist, the entire contents of the source files are copied to the target files. If the target files have already been created, changes can be propagated from one set of files to the other with `p4 resolve`. In both cases, `p4 submit` must be used to store the new file changes in the depot. Examples and further details are provided below.

### Creating branched files

To create a copy of a file that will be tracked for branching, use the following procedure:

1. Determine where you want the copied (or *branched*) file(s) to reside within the depot and within the client workspace. Add the corresponding mapping specification to your client view.
2. Run `p4 integrate fromfiles tofiles`. The source files are copied from the server to target files in the client workspace.
3. Run `p4 submit`. The new files are created within the depot, and are now available for general use.

**Example:** *Creating a branched file.*

*Version 2.0 of Elm has just been released, and work on version 3.0 is about to commence. Work on the current development release always proceeds in //depot/elm\_proj/..., and it is determined that maintenance of version 2.0 will take place in //depot/elm\_r2.0/... The files in //depot/elm\_proj/... need to be branched into //depot/elm\_r2.0/..., so Ed does the following:*

*He decides that he'll want to work on the new //depot/elm\_r2.0/... files within his client workspace at /usr/edk/elm\_proj/r2.0. He uses `p4 client` to add the following mapping to his client view:*

```
//depot/elm_r2.0/... //eds_elm/r2.0/...
```

*He runs*

```
p4 integrate //depot/elm_proj/... //depot/elm_r2.0/...
```

*which copies all the files under //depot/elm\_proj/... to //eds\_elm/v2.0 in his client workspace. Finally, he runs `p4 submit`, which adds the new branched files to the depot.*

### Why not just copy the files?

Although it is possible to accomplish everything that has been done thus far by copying the files within the client workspace and using `p4 add` to add the files to the depot, when you use `p4 integrate`, Perforce is able to track the connections between related files in an *integration record*, allowing easy propagation of changes between one set of files and another.

### Propagating changes between branched files

After a file has been branched from another with `p4 integrate`, Perforce can track changes that have been made in either set of files and merge them using `p4 resolve` into the corresponding branch files. (You'll find a general discussion of the resolve process in Chapter 5, *Perforce Basics: Resolving File Conflicts*. File resolution with branching is discussed in "How Integrate Works" on page 100).

The procedure is as follows:

1. Run `p4 integrate fromfiles tofiles` to tell Perforce that changes in the source files need to be propagated to the target files.
2. Use `p4 resolve` to copy changes from the source files to the target files. The changes are made to the target files in the client workspace.
3. Run `p4 submit` to store the changed target files in the depot.

**Example:** *Propagating changes between branched files.*

*Ed has created a release 2.0 branch of the Elm source files as above, and has fixed a bug in the original codeline's `src/elm.c` file. He wants to merge the same bug fix to the release 2.0 codeline. From his home directory, Ed types*

```
p4 integrate elm_proj/src/elm.c //depot/elm_r2.0/src/elm.c
```

*and sees*

```
//depot/elm_r2.0/src/elm.c#1 - integrate from //depot/elm_proj/src/elm.c#9
```

*The file has been scheduled for resolve. He types `p4 resolve`, and the standard merge dialog appears on his screen.*

```
/usr/edk/elm_r2.0/src/elm.c - merging //depot/elm_proj/src/elm.c#2  
Diff chunks: 0 yours + 1 theirs + 0 both + 0 conflicting  
Accept(a) Edit(e) Diff(d) Merge (m) Skip(s) Help(?) [at]:
```

*He resolves the conflict with the standard use of `p4 resolve`. When he's done, the result file overwrites the file in his branched client, and it still must be submitted to the depot.*

There is one fundamental difference between resolving conflicts in two revisions of the same file, and resolving conflicts between the same file in two different codelines. The difference is that Perforce will detect conflicts between two revisions of the same file and

then schedule a resolve, but there are *always* differences between two versions of the same file in two different codelines, and these differences usually don't need to be resolved.

In their day-to-day use, there is *no* difference between branched files and non-branched files. The standard Perforce commands like `sync`, `edit`, `delete`, `submit`, and so on, are used with all files, and evolution of both codelines proceeds separately. When changes to one codeline need to be propagated to another, you must tell Perforce to do this with `p4 integrate`, but if the codelines evolve separately, and changes never need to be propagated, you'll never need to integrate or resolve the files in the two codelines.

## Propagating changes from branched files to the original files

A change can be propagated in the reverse direction, from branched files to the original files, by supplying the branched files as the source files, and the original files as the target files.

**Example:** *Propagating changes from branched files to the original files.*

*Ed wants to integrate some changes in //depot/elm\_r2.0/src/screen.c file to the original version of the same file. He types*

```
p4 integrate //depot/elm_r2.0/src/screen.c //depot/elm_proj/src/screen.c
```

*and then runs p4 resolve. The changes in the branched file can now be merged into his source file.*

## Branching and Merging, Method 2: Branching with Branch Specifications

To map a set of source files to target files, you can create a *branch specification* and use it as an argument to `p4 integrate`. To create and use a branch specification, do the following:

1. Use `p4 branch branchname` to create a view that indicates which target files map to which source files.
2. Make sure that the new files and directories are included in the `p4 client` view of the client workspace that will hold the new files.
3. Use `p4 integrate -b branchname` to create the new files.
4. To propagate changes from source files to target files, use `p4 integrate -b branchname [tofiles]`. Perforce uses the branch specification to determine which files the merged changes come from.
5. Use `p4 submit` to submit the changes to the target files to the depot.

The following example demonstrates the same branching that was performed in the example above, this time using a branch specification.

**Example: Creating a branch.**

*Version 2.0 of Elm has just been released, and work on version 3.0 is about to commence. Work on the current development release always proceeds in //depot/elm\_proj/..., and it is determined that maintenance of version 2.0 will take place in //depot/elm\_r2.0/... The files in //depot/elm\_proj/... need to be branched into //depot/elm\_r2.0/..., so Ed does the following:*

*Ed creates a branch specification called elm2.0 by typing `p4 branch elm2.0`. The following form is displayed:*

```
Branch: elm2.0
Date:   1997/05/25 17:43:28
Owner:  edk
Description:
        Created by edk.
View:
        //depot/... //depot/...
```

*The view maps the original codeline's files (on the left) to branched files (on the right). Ed changes the View: and Description: fields as follows:*

```
Branch: elm2.0
Date:   1997/05/25 17:43:28
Owner:  edk
Description:
        Elm release 2.0 maintenance codeline
View:
        //depot/elm_proj/... //depot/elm_r2.0/...
```

*Ed wants to work on the new //depot/elm\_r2.0/... files within his client workspace at /usr/edk/elm\_proj/r2.0. He uses `p4 client` to add the following mapping to his client view: `//depot/elm_r2.0/... //eds_elm/r2.0/...`*

*He runs `p4 integrate -b elm2.0`, which copies all the files under //depot/elm\_proj/... to //eds\_elm/r2.0/... in his client workspace; then he runs `p4 submit`, which adds the newly branched files to the depot.*

Once the branch has been created and the files have been copied into the branched codeline, changes can be propagated from the source files to the target files with `p4 integrate -b branchname`.

**Example: Propagating changes to files with `p4 integrate`.**

*A bug has been fixed in the original codeline's `src/elm.c` file. Ed wants to propagate the same bug fix to the branched codeline he's been working on. He types*

```
p4 integrate -b elm2.0 ~edk/elm_r2.0/src/elm.c
```

*and sees:*



```
//depot/elm_r2.0/src/elm.c#1 - integrate from //depot/elm_proj/src/elm.c#9
```

*The file has been scheduled for resolve. He types `p4 resolve`, and the standard merge dialog appears on his screen.*

```
/usr/edk/elm_r2.0/src/elm.c - merging //depot/elm_proj/src/elm.c#9
Diff chunks: 0 yours + 1 theirs + 0 both + 0 conflicting
Accept(a) Edit(e) Diff(d) Merge (m) Skip(s) Help(?) [at]:
```

*He resolves the conflict with the standard use of `p4 resolve`. When he's done, the result file overwrites the file in his branched client, and it still must be submitted to the depot.*

## Branch Specification Usage Notes

- *Creating or altering a branch specification has absolutely no immediate effect on any files within the depot or client workspace.* The branch specification merely specifies which files are affected by subsequent `p4 integrate` commands.
- Branch specifications may contain multiple mappings, just as client views can. For example, the following branch specification branches the Elm 2.0 source code and documents to two separate locations within the depot:

```
Branch: elm2.0
Date:   1997/05/25 17:43:28
Owner:  edk
Description:
        Elm release 2.0 maintenance codeline
View:
        //depot/elm_proj/src/... //depot/elm_r2.0/src/...
        //depot/elm_proj/docs/... //depot/docs/2.0/...
```

- Exclusionary mappings can be used within branch specifications.
- To reverse the direction of an integration that uses a branch specification, use the `-r` flag.

## Integration Usage Notes

- `p4 integrate` only acts on files that are the intersection of target files in the branch view and the client view. If file patterns are given on the command line, `integrate` further limits its actions to files matching the patterns. The source files supplied as arguments to `integrate` need not be in the client view.
- The basic syntax of the `integrate` command when using a branch specification is:

```
p4 integrate -b branchname [tofiles]
```

If you omit the `tofiles` argument, all the files in the branch are affected.

- The direction of integration through a branch specification may be reversed with the `-r` flag. For example, to integrate changes from a branched file to the original source file, use `p4 integrate -b branchname -r [tofiles]`
- The `p4 integrate` command, like `p4 add`, `p4 edit`, and `p4 delete`, does not actually affect the depot immediately; instead, it adds the affected files to a changelist, which must be submitted with `p4 submit`. This keeps the `integrate` operation atomic: either all the named files are affected at once, or none of them are.
- The actual action performed by `p4 integrate` is determined by particular properties of the source files and the target files:

*If the target file doesn't exist*, the source file is copied to *target*, *target* is opened for *branch*, and Perforce begins tracking the integration history between the two files. The next integration of the two files will treat this revision of *source* as *base*.

*If the target file exists, and was originally branched from the source file with `p4 integrate`*, then a three-way merge is scheduled between *target* and *source*. The base revision is the previously integrated revision of *source*.

*If the target file exists, but was not branched from the source*, then these two file revisions did not begin their lives at a common, older file revision, so there can be no *base* file, and `p4 integrate` rejects the integration. This is referred to as a *baseless merge*. To force the integration, use the `-i` flag; `p4 integrate` will use the first revision of *source* as *base*. (Actually, `p4 integrate` uses the most recent revision of *source* that was added to the depot as *base*. Since most files are only opened for `add` once, this will almost always be the first revision of *source*.)

**Note** | In previous versions of Perforce (99.1 and earlier), integration of a target that was not originally branched from the source would schedule a two-way merge, in which the only resolve choices were `accept yours` and `accept theirs`. As of Perforce 99.2, it is no longer possible to perform a two-way merge of a text file (even when possible, it was never desirable).

- By default, a file that has been newly created in a client workspace by `p4 integrate` cannot be edited before its first submission. To make a newly-branched file available for editing before submission, perform a `p4 edit` of the file after the resolve process is complete.
- To run the `p4 integrate` command, you need Perforce `write` access on the target files, and `read` access on the source files. (See the *Perforce System Administrator's Guide* for information on Perforce protections).

## Deleting Branches

To delete a branch, use

```
p4 branch -d branchname
```

Deleting a branch deletes only the branch specification, making the branch specification inaccessible from any subsequent `p4 integrate` commands. The files themselves can still be integrated with `p4 integrate fromfiles tofiles`, and the branch specification can always be redefined. If the files in the branched codeline are to be removed, they must be deleted with `p4 delete`.

## Advanced Integration Functions

Perforce's branching mechanism also allows integration of specific file revisions, the re-integration and re-resolving of already integrated code, and merging of two files that were previously not related.

### Integrating specific file revisions

By default, the `integrate` command integrates into the target all the revisions of the source since the last source revision that `integrate` was performed on. A revision range can be specified when integrating; this prevents unwanted revisions from having to be manually deleted from the merge while editing. In this case, the revision used as `base` is the first revision below the specified revision range.

The argument to `p4 integrate` is the target, the file revision specifier is applied to the source.

**Example:** *Integrating Specific File Revisions.*

*Ed has made two bug fixes to his file `src/init.c`, and Kurt wants to integrate the change into his branched version, which is called `newinit.c`. Unfortunately, `init.c` has gone through 20 revisions, and Kurt doesn't want to have to delete all the extra code from all 20 revisions while resolving.*

*Kurt knows that the bug fixes he wants were made to file revisions submitted in changelist 30. From the directory containing his `newinit.c` file in his branched workspace, he types*

```
p4 integrate -b elm_r1 newinit.c@30,@30
```

*The target file is given as an argument, but the file revisions are applied to the source. When Kurt runs `p4 resolve`, only the revision of Ed's file that was submitted in changelist 30 is scheduled for resolve. That is, Kurt only sees the changes that Ed made to `init.c` in changelist 30. The file revision that was present in the depot at changelist 29 is used as `base`.*

## Re-integrating and re-resolving files

After a revision of a source file has been integrated into a target, that revision is usually skipped in subsequent integrations with the same target. If all the revisions of a source have been integrated into a particular target, `p4 integrate` returns the error message `All revisions already integrated`. To force the integration of already-integrated files, specify the `-f` flag to `p4 integrate`.

Similarly, a target that has been resolved but not (yet) submitted can be re-resolved by specifying the `-f` flag to `p4 resolve`, which forces re-resolution of already resolved files. When this flag is used, the original client target file has been replaced with the result file by the original resolve process; when you re-resolve, *yours* is the new client file, the result of the original resolve.

## How Integrate Works

---

The following sections describe the mechanism behind the integration process.

### The *yours*, *theirs*, and *base* files

The following table explains the terminology *yours*, *theirs*, and *base* files.

Term	Meaning
<i>yours</i>	The file to which changes are being propagated (also called the <i>target</i> file). This file in the client workspace is overwritten by the result when you resolve.
<i>theirs</i>	The file from which changes are read (also known as the <i>source</i> file). This file resides in the depot, and is not changed by the resolve process.
<i>base</i>	The last integrated revision of the source file. When you use <code>integrate</code> to create the branched copy of the file in the depot, the newly-branched copy is <i>base</i> .

### The integration algorithm

`p4 integrate` performs the following steps:

1. Apply the branch view to any target files provided on the command line to produce a list of source/target file pairs. If no files are provided on the command line, a list of all source/target file pairs is generated, including each revision of each source file that is to be integrated.

2. Discard any source/target pairs for which the source file revisions have already been integrated. Each revision of each file that has been integrated is recorded, to avoid making you merge changes more than once.
3. Discard any source/target pairs whose source file revisions have integrations pending in files that are already opened in the client.
4. Integrate all remaining source/target pairs. The target file is opened on the client for the appropriate action and merging is scheduled.

## Integrate's actions

The integrate command will take one of three actions, depending on particular characteristics of the source and target files:

Action	Meaning
branch	If the target file does not exist, it is opened for <code>branch</code> . The branch action is a variant of <code>add</code> , but Perforce keeps a record of which source file the target file was branched from. This allows three-way merges to be performed between subsequent source and target revisions with the original source file revision as <i>base</i> .
integrate	If both the source and target files exist, the target is opened for integration, which is a variant of <code>edit</code> . Before a user can submit a file that has been opened for integration, the source and target must be merged with <code>p4 resolve</code> .
delete	When the target file exists but no corresponding source file is mapped through the branch view, the target is marked for deletion. This is consistent with <code>integrate</code> 's semantics: it attempts to make the target tree reflect the source tree.

By default, when you integrate using a branch specification, the original codeline contains the source files, and the branched codeline is the target. However, if you reverse the direction of integration by specifying the `-r` flag, the branched codeline contains the source, and the original files are the targets.

## Integration Reporting

---

The branching-related reporting commands are:

Command	Function
<code>p4 integrate -n [filepatterns]</code>	Previews the results of the specified integration, but does not perform the integration. (To perform the integration, omit the <code>-n</code> flag.)
<code>p4 resolve -n [filepatterns]</code>	Displays files that are scheduled for resolve by <code>p4 integrate</code> , but does not perform the resolve. (To perform the resolve, omit the <code>-n</code> flag.)
<code>p4 resolved</code>	Displays files that have been resolved but not yet submitted.
<code>p4 branches</code>	Displays all branches.
<code>p4 integrated filepatterns</code>	Displays the integration history of the specified files.
<code>p4 filelog -i [filepatterns]</code>	Displays the revision histories of the specified files, including the integration histories of files from which the specified files were branched.

## For More Information

---

Although Perforce's branching mechanism is relatively simple, the theory of branching can be very complex. When should a branch be created? At what point should code changes be propagated from one codeline to another? Who is responsible for performing merges? These questions will arise no matter what SCM system you're using, and the answers are not simple. Three on-line documents can provide some guidance in these matters.

A white paper on *InterFile Branching*, which describes Perforce's branching mechanism in technical detail, is available from:

<http://www.perforce.com/perforce/branch.html>

Christopher Seiwald and Laura Wingerd's Best SCM Practices paper provides a discussion of many source configuration management issues, including an overview of basic branching techniques. This paper is available at:

<http://www.perforce.com/perforce/bestpractices.html>

*Streamed Lines: Branching Patterns for Parallel Software Development* is an extremely detailed paper on branching techniques. You'll find it at:

<http://www.enteract.com/~bradapp/acme/plop98/streamed-lines.html>

---

## Chapter 10 Job Tracking

---

A *job* is a written description of some modification to be made to a source code set. A job might be a bug description, like “the system crashes when I press `return`”, or it might be a system improvement request, like “please make the program run faster.”

Whereas a job represents work that is intended, a changelist represents work actually done. Perforce’s job tracking mechanism allows jobs to be linked to the changelists that implement the work requested by the job. A job can later be looked up to determine if and when it was fixed, which file revisions implemented the fix, and who fixed it. A job linked to a particular changelist is marked as completed when the changelist is submitted.

Jobs perform no functions internally to Perforce; rather, they are provided as a method of keeping track of information such as what changes to the source are needed, which user is responsible for implementing the job, and which file revisions contain the implementation of the job. The type of information tracked by the jobs system can be customized; fields in the job form can be added, changed, and deleted by Perforce superusers.

---

### Job Usage Overview

---

There are five related but distinct aspects of using jobs.

- The Perforce superuser decides what fields are to be tracked in your system’s jobs, the possible values of a job’s fields, and their default values. This job template is edited with the `p4 jobspec` command. (See the *Perforce System Administrator’s Guide* for details on how to edit the job specification. The job specification need not be changed before users can create jobs).
- Jobs can be created and edited by any user with `p4 job`.
- The `p4 jobs` command can be used to look up all the jobs that meet specified criteria.
- Jobs can be linked to changelists automatically or manually; when a job is linked to a changelist, the job is marked as `closed` when the changelist is submitted.
- The jobs that have been fixed can be displayed with Perforce reporting commands. These commands can list all jobs that fixed particular files or file revisions, all the jobs that were fixed in a particular changelist, or all the changelists that were linked to a particular job fix.

The remainder of this chapter discusses how these tasks are accomplished.

## Creating and editing jobs using the default job specification

Jobs are created with the `p4 job` command.

### Example: *Creating a job*

*Sarah's Perforce server uses Perforce's default jobs specification. Sarah knows about a job in Elm's filtering subsystem, and she knows that Ed is responsible for Elm filters. Sarah creates a new job with `p4 job` and fills in the resulting form as follows:*

```

Job:      new
User:     edk
Status:   new
Date:     1998/05/18 17:15:40
Description:
          Filters on the "Reply-To:" field
          don't work.
    
```

*Sarah has filled in a description and has changed `User:` to `edk`.*

Since job fields differ from site to site, the fields in jobs at your site may be very different than what you see above. The default `p4 job` form's fields are:

Field Name	Description	Default
Job	The name of the job. Whitespace is not allowed in the name.	new
User	The user whom the job is assigned to, usually the username of the person assigned to fix this particular problem.	Perforce username of the person creating the job.
Status	<p>open, closed, suspended, or new.</p> <p>An <code>open</code> job is one that has been created but has not yet been fixed.</p> <p>A <code>closed</code> job is one that has been completed.</p> <p>A <code>suspended</code> job is an open job that is not currently being worked on.</p> <p>Jobs with status <code>new</code> exist only while a new job is being created; they change to status <code>open</code> as soon as the form has been completed and the job added to the database.</p>	new; changes to open after job creation form is closed.



Field Name	Description	Default
Date	The date the job was created, displayed as <i>YYYY/MM/DD HH/MM/SS</i>	The date and time at the moment this job was created. Changes to this field are ignored.
Description	Arbitrary text assigned by the user. Usually a written description of the problem that is meant to be fixed.	Text that <i>must</i> be changed

If `p4 job` is called with no parameters, a new job is created. The name that appears on the form is `new`, but this can be changed by the user to any desired string. If the `Job:` field is left as `new`, Perforce will assign the job the name `jobN`, where `N` is a sequentially-assigned six-digit number.

Existing jobs can be edited with `p4 job jobname`. The user and description can be changed arbitrarily; the status can be changed to any of the three valid status values `open`, `closed`, or `suspended`. When you call `p4 job jobname` with a nonexistent `jobname`, Perforce creates a new job. (A job, if submitted with a `Status:` of `new`, has this status automatically changed to `open` upon completion of the job form.)

## Creating and editing jobs with custom job specifications

A Perforce superuser can add and change fields within your server's jobs template with the `p4 jobspec` command. If this has been done, there may be additional fields in your `p4 jobs` form, and the names of the fields described above may have changed.

A sample customized job specification might look like this:

```
# Sample customized jobspec
#
# Job:          Job number
# Type:         The type of request: "bug" or "feature"
# Status:       Has it been fixed: "open", "closed", or "inprogress"
# Priority:     How soon should this job be fixed?
#              Values are "a", "b", "c", or "unknown"
# Owned_by:    Who's fixing the bug
# Reported_by: Who reported the bug
# Reported_date: When the bug was first entered
# Mod_date:    Last time the bug was updated
# Description:  Textual description of the bug

Job:    new
Type:   unknown
Status: open
Priority:    unknown
User:       unowned
Reported_by:  edk
Reported_date: 2001/06/02 10:09:46
Mod_date:    2001/06/03 14:22:38
Description:
          <enter description here>
```

Some of the fields have been set by the superuser to allow one of a set of values; for example, `Priority:` must be one of a, b, c, or unknown. The `p4 job` fields don't tell you what the valid values of the fields are; your Perforce superuser can tell you this in comments at the top of the job form. If you find the information in the comments for your jobs to be insufficient to enter jobs properly, please tell your Perforce superuser.

## Viewing jobs by content with jobviews

---

Jobs can be reported with `p4 jobs`. In its simplest form, with no arguments, `p4 jobs` will list every job stored in your Perforce server. However, `p4 job -e jobview` will list all jobs that match the criteria contained in `jobview`.

Throughout the following discussion, the following rules apply:

- Textual comparisons within jobviews are case-insensitive, as are the field names that appear in jobviews,
- only alphanumeric text and punctuation can appear in a jobview,
- there is currently no way to search for particular phrases. Jobviews can search jobs only by individual words.

## Finding jobs containing particular words

The jobview `'word1 word2 ... wordN'` can be used to find jobs that contain all of `word1` through `wordN` in any field (excluding date fields).

**Example:** *Finding jobs that contain all of a set of words in any field.*

*Ed wants to find all jobs that contain the words filter, file, and mailbox. He types:*

```
p4 jobs -e 'filter file mailbox'
```

Spaces between search terms in jobviews act as boolean *and*'s. You can use ampersands instead of spaces in jobviews, so the jobviews `'joe sue'` and `'joe&sue'` are identical.

To find jobs that contain any of the terms, separate the terms with the `'|'` character.

**Example:** *Finding jobs that contain any of a set of words in any field.*

*Ed wants to find jobs that contains any of the words filter, file or mailbox. He types:*

```
p4 jobs -e 'filter|file|mailbox'
```

## Finding jobs by field values

Search results can be narrowed by matching values within specific fields with the jobview syntax `'fieldname=value'`. Value must be a single alphanumeric word.

**Example:** *Finding jobs that contain words in specific fields*

*Ed wants to find all open jobs related to filtering of which he is the owner. He types:*

```
p4 jobs -e 'status=open user=edk filter.c'
```

*This will find all jobs with a Status: of open, a User: of edk, and the word filter.c in any non-date field.*

## Using and escaping wildcards in jobviews

The wildcard `"*"` allows for partial word matches. The jobview `"fieldname=string*"` matches `"string"`, `"stringy"`, `"stringlike"`, and so on.

To search for words that happen to contain wildcards, escape them at the command line. For instance, to search for `"*string"` (perhaps in reference to `char *string`), you'd use the following:

```
p4 jobs -e '\*string'
```

## Negating the sense of a query

The sense of a search term can be reversed by prefacing it with `^`, the *not* operator.

**Example:** *Finding jobs that don't contain particular words.*

*Ed wants to find all open jobs related to filtering of which he is not the owner. He types:*

```
p4 jobs -e 'status=open ^user=edk filter'
```

*This displays all jobs with a Status: of open, a User: of anyone but edk, and the word filter in any non-date field.*

The *not* operator `^` can be used only directly after an *and* (space or `&`). Thus, `p4 jobs -e '^user=edk'` is not allowed.

You can use the `*` wildcard to get around this: `p4 jobs -e 'job=* ^user=edk'` returns all jobs with a `user` field not matching `edk`.

## Using dates in jobviews

Jobs can be matched by date by expressing the date as `yyyy/mm/dd` or `yyyy/mm/dd:hh:mm:ss`. If you don't provide a specific time, the equality operator `=` matches the entire day.

**Example:** *Using dates within jobviews.*

*Ed wants to view all jobs modified on July 13, 1998. He enters*

```
p4 jobs -e 'mod_date=1998/07/13'
```

## Comparison operators and field types

The usual comparison operators are available. They are:

```
=      >      <      >=     <=
```

The behavior of these operators depends upon the type of the field in the jobview. The field types are:

Field Type	Explanation	Examples
word	A single word	A user name: <code>edk</code>
text	A block of text	A job's description
line	A single line of text. Differs from text fields only in that line values are entered on the same line as the field name, and text values are entered on the lines beneath the field name.	An email address A user's real name: <code>Linda Hopper</code>

Field Type	Explanation	Examples
<code>select</code>	One of a set of values	A job's status: open/suspended/closed
<code>date</code>	A date value	The date and time of job creation: 1998/07/15:13:21:4

Field types are often obvious from context; a field called `mod_date`, for example, is most likely a date field. If you're not sure of a field's type, run `p4 jobspec -o`, which outputs the job specification your local jobspec. The field called `Fields:` lists the job fields' names and datatypes.

The jobview comparison operators behave differently depending upon the type of field they're used with. The comparison operators match the different field types as follows:

Field Type	Use of Comparison Operators in Jobviews
<code>word</code>	The equality operator = must match the value in the word field exactly.
<code>text</code>	The inequality operators perform comparisons in ASCII order. The equality operator = matches the job if the word given as the value is found anywhere in the specified field. The inequality operators are of limited use here, since they'll match the job if <i>any</i> word in the specified field matches the provided value. For example, if a job has a text field <code>ShortDescription:</code> that contains only the phrase <code>gui bug</code> , and the jobview is <code>'ShortDesc&lt;filter'</code> , the job will match the jobview, because <code>bug&lt;filter</code> .
<code>line</code>	See <code>text</code> , above.
<code>select</code>	The equality operator = matches a job if the value of the named field is the specified word. Inequality operators perform comparisons in ASCII order.
<code>date</code>	Dates are matched chronologically. If a specific time is not provided, the operators =, <=, and >= will match the whole day.

## Linking Jobs to Changelists

---

Perforce automatically changes the value of a job's `status` field to `closed` when the job is linked to a particular changelist, and the changelist is submitted.

Jobs can be linked to changelists in one of two ways:

- Automatically, by setting the `JobView:` field in the `p4 user` form to a jobview that matches the job, and
- manually, with the `p4 fix` command.

### Automatically linking jobs to changelists with the `p4 user` form

The `p4 user` form can be used to automatically include particular jobs on any new changelists created by that user. To do this, call `p4 user` and change the `JobView:` field value to any valid jobview.

**Example:** *Automatically linking jobs to changelists with the `p4 user` form's `JobView` field.*

*Ed wants to see all open jobs that he owns in all changelists he creates. He types `p4 user` and adds a `JobView:` field:*

```
User:      edk
Update:    1998/06/02 13:11:57
Access:    1998/06/03 20:11:07
JobView:   user=edk&status=open
```

*All of Ed's jobs that meet these `JobView:` criteria will automatically appear on all changelists he creates. He can, and should, delete jobs that aren't fixed by the changelist from the changelist form before submission. When a changelist is submitted, the jobs linked to it will have their `status:` field's value changed to `closed`.*

### Automatic update of job status

The value of a job's `status` field is automatically changed to `closed` when one of its associated changelists is successfully submitted. Jobs can be disassociated from changelists by deleting the job from the changelist's change form. Similarly, any job can be added to a changelist by adding it to a changelist's change form.

**Example:** *Including and excluding jobs from changelists.*

*Ed has set his p4 user's JobView: field as in the example above. He is unaware of a job that Sarah has made Ed the owner of (when she entered the job, she set the User: field to edk). He is currently working on an unrelated problem; he types p4 submit and sees the following:*

```
Change: new
Client: edk
User: edk
Status: new
Description:
    Updating "File" files
Jobs:
    job000125          # Filters on "Reply-To" field don't work
Files:
    //depot/src/file.c          # edit
    //depot/src/file_util.c    # edit
    //depot/src/fileio.c       # edit
```

*Since this job is unrelated to the work he's been doing, and since it hasn't been fixed, he deletes it from the form and then quits from the editor. The changelist is submitted, and the job is not associated with it.*

**Example:** *Submit a changelist with an attached job.*

*Ed uses the reporting commands to read the details about the job. He fixes this problem, and a number of other filtering bugs; when he next types p4 submit, he sees:*

```
Change: new
Client: edk
User: rlo
Status: new
Description:
    Fixes a number of filter problems
Jobs:
    job000125          # Filters on "Reply-To" field don't work
Files:
    //depot/filter/actions.c    # edit
    //depot/filter/audit.c     # edit
    //depot/filter/filter.c     # edit
```

*Since the job is fixed in this changelist, Ed leaves the job on the form. When he quits from the editor, the job's status is changed to closed.*

## Manually associating jobs with changelists

p4 fix -c changenum jobname can be used to link any job to any changelist. If the changelist has already been submitted, the value of the job's Status: field is changed to closed. Otherwise, the job keeps its current status.

**Example:** *Manually attaching jobs to changelists.*

*Sarah has submitted a job called `options-bug` to Ed. Unbeknownst to Sarah, the bug reported by the job was fixed in Ed's previously submitted changelist 18. Ed links the job to the previously submitted changelist with `p4 fix -c 18 options-bug`. Since the changelist has already been submitted, the job's status is changed to `closed`.*

It is never necessary to use `p4 fix` to link an open job to a changelist newly created by the owner of the job, since this is done automatically. However, you can use `p4 fix` to link a changelist to a job owned by another user.

## What if there's no status field?

The discussion in this section has assumed that the server's job specification still contains the default `Status:` field. If the job specification has been altered so that this is no longer true, jobs can still be linked to changelists, but nothing in the job changes when the changelist is submitted. (In most cases, this is not a desired form of operation.) Please see the chapter on editing job specifications in the *Perforce System Administrator's Guide* for more details.

## Deleting Jobs

---

A job can be unlinked from any changelist with `p4 fix -d -c changenum jobname`.

Jobs can be deleted entirely with `p4 job -d jobname`.

## Integrating with External Defect Tracking Systems

---

If you want to integrate Perforce with your in-house defect tracking system, or develop an integration with a third-party defect tracking system, P4DTI is probably the best place to start.

To get started with P4DTI, see the P4DTI product information page at:

<http://www.perforce.com/perforce/products/p4dti.html>

Available from this page are the TeamShare and Bugzilla implementations, an overview of the P4DTI's capabilities, and a kit (including source code and developer documentation) for building integrations with other products or in-house systems.

Even if you don't use the P4DTI kit as a starting point, you can still use Perforce's job system as the interface between Perforce and your defect tracker. See the *Perforce System Administrator's Guide* for more information.



## Job Reporting Commands

The job reporting commands can be used to show the correspondence between files and the jobs they fix.

To See a Listing of...	Use This Command:
...all jobs that match particular criteria	<code>p4 jobs -e jobview</code>
...all the jobs that were fixed by changelists that affected particular file(s)	<code>p4 jobs filespec</code>
...all changelists and file revisions that fixed a particular job	<code>p4 fixes -j jobname</code>
...all jobs linked to a particular changelist	<code>p4 fixes -c changenum</code>
...all jobs fixed by changelists that contain particular files or file revisions	<code>p4 fixes filespec</code>

Other job reporting variations are available. For more examples, please see “Job Reporting” on page 124, or consult the *Perforce Command Reference*.



---

## Chapter 11 Reporting and Data Mining

---

Perforce's reporting commands supply information on all data stored within the depot. There are many reporting commands; in fact, there are almost as many reporting commands as there are action commands. These commands have been discussed throughout this manual; this chapter presents the same commands and provides additional information for each command. Tables in each section contain answers to questions of the form "How do I find information about...?"

Many of the reporting commands have numerous options, but discussion of all options for each command is beyond the scope of this manual. For a full description of any particular command, please consult the *Perforce Command Reference*, or type `p4 help command` at the command line.

One previously mentioned note on syntax is worth repeating here: any *filespec* argument in Perforce commands, as in:

```
p4 files filespec
```

will match any file pattern that is supplied in local syntax, depot syntax, or client syntax, with any Perforce wildcards. Brackets around [*filespec*] mean that the file specification is optional. Additionally, many of the reporting commands can take revision specifiers as part of the filespec. Revision specifiers are discussed on "Specifying Older File Revisions" on page 46.

### Files

---

The commands that report on files fall into two categories: those that give information about file *contents*, (for instance, `p4 print`, `p4 diff`), and those that supply information on file *metadata*, the data that describe a file with no reference to content (e.g. `p4 files`, `p4 filelog`). The first set of reporting commands discussed in this section describe file metadata, while the second set describe file contents.

### File metadata

#### Basic file information

To view information about single revisions of one or more files, use `p4 files`. This command provides the locations of the files within the depot, the actions (`add`, `edit`, `delete`, and so on) on those files at the specified revisions, the changelists the specified file revisions were submitted in, and the files' types. The output has this appearance:

```
//depot/README#5 - edit change 6 (text)
```

The `p4 files` command requires one or more *filespec* arguments. Filespec arguments can, as always, be provided in Perforce or local syntax, but the output always reports on the corresponding files within the depot. If you don't provide a revision number, Perforce uses the head revision.

Unlike most other commands, `p4 files` also describes deleted revisions, rather than suppressing information about deleted files.

To View File Metadata for...	Use This Command:
...all files in the depot, whether or not visible through your client view	<code>p4 files //depot/...</code>
...all the files currently in any client workspace	<code>p4 files @clientname</code>
...all the files in the depot that are mapped through your current client workspace view	<code>p4 files //clientname/...</code>
...a particular set of files in the current working directory	<code>p4 files filespec</code>
...a particular file at a particular revision number	<code>p4 files filespec#revisionNum</code>
...all files at change <i>n</i> , whether or not the file was actually included in change <i>n</i>	<code>p4 files @n</code>
...a particular file within a particular label	<code>p4 files filespec@labelname</code>

### File revision history

The revision history of a file is provided by `p4 filelog`. One or more file arguments must be provided, and since the point of `p4 filelog` is to list information about each revision of particular files, file arguments to this command may not contain a revision specification.

The output of `p4 filelog` has this form:

```
... #3 change 23 edit on 1997/09/26 by edk@doc <ktext> 'Fix help system'
... #2 change 9 edit on 1997/09/24 by lisag@src <text> 'Change file'
... #1 change 3 add on 1997/09/24 by edk@doc <text> 'Added filtering bug'
```

For each file that matches the *filespec* argument, the complete list of file revisions is presented, along with the number of the changelist that the revision was submitted in, the date of submission, the user who submitted the revision, the file's type at that revision, and the first few characters of the changelist description. With the `-l` flag, the entire description of each changelist is printed:

```
#3 change 23 edit on 1997/09/26 by edk@doc
Updated help files to reflect changes
in filtering system & other subsystems
...<etc.>
```

## Opened files

To see which files are currently opened within a client workspace, use `p4 opened`. For each opened file within the client workspace that matches a file pattern argument, `p4 opened` prints a line like the following:

```
//depot/elm_proj/README - edit default change (text)
```

Each opened file is described by its depot name and location, the operation that the file is opened for (add, edit, delete, branch, or integrate), which changelist the file is included in, and the file's type.

To See...	Use This Command:
...a listing of all opened files in the current workspace	<code>p4 opened</code>
...a list of all opened files in all client workspaces	<code>p4 opened -a</code>
...a list of all files in a numbered pending changelist	<code>p4 opened -c changelist#</code>
...a list of all files in the default changelist	<code>p4 opened -c default</code>
...whether or not a specific file is opened by you	<code>p4 opened filespec</code>
...whether or not a specific file is opened by anyone	<code>p4 opened -a filespec</code>

## Relationships between client and depot files

It is often useful to know how the client and depot are related at a particular moment in time. Perhaps you simply want to know where a particular client file is mapped to within the depot, or you may want to know whether or not the head revision of a particular depot file has been copied to the client. The commands that express the relationship between client and depot files are `p4 where`, `p4 have`, and `p4 sync -n`. The first of these commands, `p4 where`, shows the mappings between client workspace files, depot files, and local OS syntax. `p4 have` tells you which revisions of files you've last synced to your client workspace, and `p4 sync -n` describes which files would be read into your client workspace the next time you perform a `p4 sync`.

All these commands can be used with or without *filespec* arguments. `p4 sync -n` is the only command in this set that allows revision specifications on the *filespec* arguments.

The output of `p4 where filename` looks like this:

```
//depot/elm_proj/doc/Ref.guide //edk/doc/Ref.guide /usr/edk/doc/Ref.guide
```

The first part of the output is the location of the file in depot syntax; the second part is the location of the same file in client syntax, and the third is the location of the file in local OS syntax.

`p4 have`'s output has this form:

```
//depot/doc/Ref.txt#3 - /usr/edk/elm/doc/Ref.txt
```

and `p4 sync -n` provides output like:

```
//depot/doc/Ref.txt#3 - updating /usr/edk/elm/doc/Ref.txt
```

The following table lists other useful commands:

To See...	Use This Command:
...which revisions of which files you have in the client workspace	<code>p4 have</code>
...which revision of a particular file is in your client workspace	<code>p4 have filespec</code>
...where a particular file maps to within the depot, the client workspace, and the local OS	<code>p4 where filespec</code>
...where a particular file in the depot maps to in the workspace	<code>p4 where //depot/.../filespec</code>
...which files would be synced into your client workspace from the depot when you do the next sync	<code>p4 sync -n</code>

## File contents

### Contents of a single revision

You can view the contents of any file revision within the depot with `p4 print`. This command simply prints the contents of the file to standard output, or to the specified output file, along with a one-line banner that describes the file. The banner can be removed by passing the `-q` flag to `p4 print`. When printed, the banner has this format:

```
//depot/elm_proj/README#23 - edit change 50 (text)
```

`p4 print` takes a mandatory file argument, which can include a revision specification. If a revision is specified, the file is printed at the specified revision; if no revision is specified, the head revision is printed.

To See the Contents of Files...	Use This Command:
...at the current head revision	<code>p4 print filespec</code>
...without the one-line file header	<code>p4 print -q filespec</code>
...at a particular change number	<code>p4 print filespec@changenum</code>

### File content comparisons

A client workspace file can be compared to any revision of the same file within in the depot with `p4 diff`. This command takes a `filespec` argument; if no revision specification is supplied, the workspace file is compared against the revision last read into the workspace.

The `p4 diff` command has many options available; only a few are described in the table below. For more details, please consult the *Perforce Command Reference*.

Whereas `p4 diff` compares a client workspace file against depot file revisions, `p4 diff2` can be used to compare *any* two revisions of a file. It can even be used to compare revisions of different files. `p4 diff2` takes two file arguments -- wildcards are allowed, but any wildcards in the first file argument must be matched with a corresponding wildcard in the second. This makes it possible to compare entire trees of files.

There are many more flags to `p4 diff` than described below. For a full listing, please type `p4 help diff` at the command line, or consult the *Perforce Command Reference*.

To See the Differences between...	Use This Command:
...an open file within the client workspace and the revision last taken into the workspace	<code>p4 diff file</code>
...any file within the client workspace and the revision last taken into the workspace	<code>p4 diff -f file</code>
...a file within the client workspace and the same file's current head revision	<code>p4 diff file#head</code>
...a file within the client workspace and a specific revision of the same file within the depot	<code>p4 diff file#revnumber</code>
...the <i>n</i> -th and head revisions of a particular file	<code>p4 diff2 filespec filespec#n</code>
...all files at changelist <i>n</i> and the same files at changelist <i>m</i>	<code>p4 diff2 filespec@n filespec@m</code>
...all files within two branched codelines	<code>p4 diff2 //depot/path1/... //depot/path2/...</code>
...a file within the client workspace and the revision last taken into the workspace, passing the context diff flag to the underlying diff	<code>p4 diff -dc file</code>

The last example above bears further explanation; to understand how this works, it is necessary to discuss how Perforce implements and calls underlying diff routines.

Perforce uses two separate diffs: one is built into the `p4d` server, and the other is used by the `p4` client. Both diffs contain identical, proprietary code, but are used by separate sets of commands. The client diff is used by `p4 diff` and `p4 resolve`, and the server diff is used by `p4 describe`, `p4 diff2`, and `p4 submit`.

Perforce's built-in diff routine allows three `-d<flag>` flags: `-du`, `-dc`, and `-dn`. Both `p4 diff` and `p4 diff2` allow any of these flags to be specified. These flags behave identically to the corresponding flags in the standard UNIX diff.

Although the server must always use Perforce's internal diff routine, the client diff can be set to any external diff program by pointing the `P4DIFF` environment variable to the full path name of the desired executable. Any flags used by the external diff can be passed to it with `p4 diff`'s `-d` flag. Flags are passed to the underlying diff according to the following rules:

- If the character immediately following the `-d` is not a single quote, then all the characters between the `-d` and whitespace are prepended with a dash and sent to the underlying diff.
- If the character immediately following the `-d` is a single quote, then all the characters between the opening quote and the closing quote are prepended with a dash and sent to the underlying diff.

The following examples demonstrate the use of these rules in practice.

<b>If you want to pass the following flag to an external client diff program:</b>	<b>Then call <code>p4 diff</code> this way:</b>
<code>-u</code>	<code>p4 diff -du</code>
<code>--brief</code>	<code>p4 diff -d-brief</code>
<code>-C 25</code>	<code>p4 diff -d'C 25'</code>

---

## Changelists

Two separate commands are used to describe changelists. The first, `p4 changes`, lists changelists that meet particular criteria, without describing the files or jobs that make up the changelist. The second command, `p4 describe`, lists the files and jobs affected by a single changelist. These commands are described below.

### Viewing changelists that meet particular criteria

To view a list of changelists that meet certain criteria, such as changelists with a certain status, or changelists that affect a particular file, use `p4 changes`.



The output looks like this:

```
Change 36 on 1997/09/29 by edk@eds_elm 'Changed filtering me'
Change 35 on 1997/09/29 by edk@eds_elm 'Misc bug fixes: fixe'
Change 34 on 1997/09/29 by lisag@lisa 'Added new header inf'
```

By default, `p4 changes` displays an aggregate report containing one line for every changelist known to the system, but command line flags and arguments can be used to limit the changelists displayed to those of a particular status, those affecting a particular file, or the last  $n$  changelists.

Currently, the output can't be restricted to changelists submitted by particular users, although you can write simple shell or Perl scripts to implement this (you'll find an example of such a script in the *Perforce System Administrator's Guide*).

To See a List of Changelists...	Use This Command:
...with the first 31 characters of the changelist descriptions	<code>p4 changes</code>
...with the complete description of each changelist	<code>p4 changes -l</code>
...including only the last $n$ changelists	<code>p4 changes -m n</code>
...with a particular status (pending or submitted)	<code>p4 changes -s status</code>
...from a particular user	<code>p4 changes -u user</code>
...from a particular client workspace	<code>p4 changes -c workspace</code>
...limited to those that affect particular files	<code>p4 changes filespec</code>
...limited to those that affect particular files, but including changelists that affect files which were later integrated with the named files	<code>p4 changes -i filespec</code>
...limited to changelists that affect particular files, including only those changelists between revisions $m$ and $n$ of these files	<code>p4 changes filespec#m,#n</code>
...limited to those that affect particular files at each files revisions between labels $lab1$ and $lab2$	<code>p4 changes filespec@lab1,@lab2</code>
...limited to those between two dates	<code>p4 changes @date1,@date2</code>
...between an arbitrary date and the present day	<code>p4 changes @date1,@now</code>

**Note** For details about Perforce commands that allow you to use revision ranges with file specifications, see “Revision Ranges” on page 49.

## Files and jobs affected by changelists

To view a list of files and jobs affected by a particular changelist, along with the *diffs* of the new file revisions and the previous revisions, use `p4 describe`.

The output of `p4 describe` looks like this:

```
Change 43 by lisag@warhols on 1997/08/29 13:41:07

    Made grammatical changes to basic Elm documentation

Jobs fixed...

job000001 fixed on 1997/09/29 by edk@edk
    Fix grammar in main Elm help file

Affected files...

... //depot/doc/elm.1#2 edit

Differences...

==== //depot/doc/elm.1#2 (text) ====
53c53
> Way number 2, what is used common-like when, you know, like
---
> The second method is commonly used when transmitting
...<etc.>
```

This output is quite lengthy, but a shortened form that eliminates the *diffs* can be generated with `p4 describe -s changenum`.

To See:	Use This Command:
...a list of files contained in a pending changelist	<code>p4 opened -c changelist#</code>
...a list of all files submitted and jobs fixed by a particular changelist, displaying the <i>diffs</i> between the file revisions submitted in that changelist and the previous revisions	<code>p4 describe changenum</code>
...a list of all files submitted and jobs fixed by a particular changelist, without the file <i>diffs</i>	<code>p4 describe -s changenum</code>

To See:	Use This Command:
...a list of all files and jobs affected by a particular changelist, while passing the context diff flag to the underlying diff program	<code>p4 describe -dc changenum</code>
...the state of particular files at a particular changelist, whether or not these files were affected by the changelist	<code>p4 files filespec@changenum</code>

For more commands that report on jobs, see “Job Reporting” on page 124.

## Labels

Reporting on labels is accomplished with a very small set of commands. The only command that reports only on labels, `p4 labels`, prints its output in the following format:

```
Label release1.3 1997/5/18 'Created by edk'
Label lisas_temp 1997/10/03 'Created by lisag'
...<etc.>
```

The other label reporting commands are variations of commands we’ve seen earlier.

To See:	Use This Command:
...a list of all labels, the dates they were created, and the name of the user who created them	<code>p4 labels</code>
...a list of all labels containing a specific revision (or range)	<code>p4 labels file#revrange</code>
...a list of files that have been included in a particular label with <code>p4 labelsync</code>	<code>p4 files @labelname</code>
...what <code>p4 sync</code> would do when retrieving files from a particular label into your client workspace	<code>p4 sync -n @labelname</code>

## Branch and Integration Reporting

The plural form command of branch, `p4 branches`, lists the different branches in the system, along with their owners, dates created, and descriptions. Separate commands are used to list files within a branched codeline, to describe which files have been integrated, and to perform other branch-related reporting.

The table below describes the most commonly used commands for branch- and integration-related reporting.

To See:	Use This Command:
...a list of all branches known to the system	<code>p4 branches</code>
...a list of all files in a particular branched codeline	<code>p4 files filespec</code>
...what a particular <code>p4 integrate</code> variation would do, without actually doing the <code>integrate</code>	<code>p4 integrate [args] -n [filespec]</code>
...a list of all the revisions of a particular file	<code>p4 filelog -i filespec</code>
...what a particular <code>p4 resolve</code> variation would do, without actually doing the <code>resolve</code>	<code>p4 resolve [args] -n [filespec]</code>
...a list of files that have been resolved but have not yet been submitted	<code>p4 resolved [filespec]</code>
...a list of integrated, submitted files that match the <code>filespec</code> arguments	<code>p4 integrated filespec</code>
...a list of all the revisions of a particular file, including revision of the file(s) it was branched from	<code>p4 filelog -i filespec</code>

## Job Reporting

---

Two commands report on jobs. The first, `p4 jobs`, reports on all jobs known to the system, while the second command, `p4 fixes`, reports only on those jobs that have been attached to changelists. Both of these commands have numerous options.

### Basic job information

To see a list of all jobs known to the system, use `p4 jobs`. This command produces output similar to the following:

```
job000302 on 1997/08/13 by saram *open* 'FROM: headers no'  
filter_bug on 1997/08/23 by edk *closed* 'Can't read filters w'
```

Its output includes the job's name, date entered, job owner, status, and the first 31 characters of the job description.

All jobs known to the system are displayed unless command-line options are supplied. These options are described in the table below.

To See a List of Jobs:	Use This Command:
...including all jobs known to the server	<code>p4 jobs</code>
...including the full texts of the job descriptions	<code>p4 jobs -l</code>
...for which certain fields contain particular values (For more about jobviews, see “Viewing jobs by content with jobviews” on page 106)	<code>p4 jobs -e jobview</code>
...that have been fixed by changelists that contain specific files	<code>p4 jobs filespec</code>
...that have been fixed by changelists that contain specific files, including changelists that contain files that were later integrated into the specified files	<code>p4 jobs -i filespec</code>

## Jobs, fixes, and changelists

Any jobs that have been linked to a changelist with `p4 change`, `p4 submit`, or `p4 fix` is said to be *fixed*, and can be reported with `p4 fixes`.

The output of `p4 fixes` looks like this:

```
job000302 fixed by change 634 on 1997/09/01 by edk@eds_elm
filter_bug fixed by change 540 on 1997/10/22 by edk@eds_elm
```

A number of options allow the reporting of only those changes that fix a particular job, jobs fixed by a particular changelist, or jobs fixed by changelists that are linked to particular files.

A fixed job will not necessarily have a status of `closed job`, since `open jobs` can be linked to pending changelists, and `pending jobs` can be reopened even after the associated changelist has been submitted. To list jobs with a particular status, use `p4 jobs`.

To See a Listing of...	Use This Command:
...all fixes for all jobs	<code>p4 fixes</code>
...all changelists linked to a particular job	<code>p4 fixes -j jobname</code>
...all jobs linked to a particular changelist	<code>p4 fixes -c changenum</code>
...all jobs fixed by changelists that contain particular files	<code>p4 fixes filespec</code>
...all jobs fixed by changelists that contain particular files, including changelists that contain files that were later integrated with the specified files	<code>p4 fixes -i filespec</code>

## Reporting for Daemons

The Perforce change review mechanism uses the following reporting commands. Any of these commands might also be used with user-created daemons. For further information on daemons, please see the *Perforce System Administrator's Guide*.

To list...	Use this Command:
...the names of all counter variables currently used by your Perforce system	<code>p4 counters</code>
...the numbers of all changelists that have not yet been reported by a particular counter variable	<code>p4 review -t countername</code>
...all users who have subscribed to review particular files	<code>p4 reviews filespec</code>
...all users who have subscribed to read any files in a particular changelist	<code>p4 reviews -c changenum</code>
...a particular user's email address	<code>p4 users username</code>

## System Configuration

Three commands report on the Perforce system configuration. One command reports on all Perforce users, another prints data describing all client workspaces, and a third reports on Perforce depots.

`p4 users` generates its data as follows:

```
edk <edk@eds_ws> (Ed K.) accessed 1997/07/13
lisag <lisa@lisas_ws> (Lisa G.) accessed 1997/07/14
```

Each line includes a username, an email address, the user's "real" name, and the date that Perforce was last accessed by that user.

To report on client workspaces, use `p4 clients`:

```
Client eds_elm 1997/09/12 root /usr/edk 'Ed's Elm workspace'
Client lisa_doc 1997/09/13 root /usr/lisag 'Created by lisag.'
```

Each line includes the client name, the date the client was last updated, the client root, and the description of the client.

Depots can be reported with `p4 depots`. All depots known to the system are reported on; the described fields include the depot's name, its creation date, its type (`local` or `remote`), its IP address (if `remote`), the mapping to the local depot, and the system administrator's description of the depot.

The use of multiple depots on a single Perforce server is discussed in the *Perforce System Administrator's Guide*.

To view:	Use This Command:
...user information for all Perforce users	<code>p4 users</code>
...user information for only certain users	<code>p4 users username</code>
...brief descriptions of all client workspaces	<code>p4 clients</code>
...a list of all defined depots	<code>p4 depots</code>

## Special Reporting Flags

Two special flags, `-o` and `-n`, can be used with certain action commands to change their behavior from action to reporting.

The `-o` flag is available with most of the Perforce commands that normally bring up forms for editing. This flag tells these commands to write the form information to standard output, instead of bringing the definition into the user's editor. This flag is supported by the following commands:

<code>p4 branch</code>	<code>p4 client</code>	<code>p4 label</code>
<code>p4 change</code>	<code>p4 job</code>	<code>p4 user</code>

The `-n` flag prevents commands from doing their job. Instead, the commands simply tell you what they would ordinarily do. You can use the `-n` flag with the following commands

<code>p4 integrate</code>	<code>p4 resolve</code>	<code>p4 labelsync</code>	<code>p4 sync</code>
---------------------------	-------------------------	---------------------------	----------------------

## Reporting with Scripting

Although Perforce's reporting commands are sufficient for most needs, there may be times when you want to view data in a format that Perforce doesn't directly support. In these situations, the reporting commands can be used in combination with scripts to print only the data that you want to see. Here are two examples.

### Comparing the change content of two file sets

To compare the "change content" of two sets of files, you have to diff them externally. To do this, run `p4 changes` twice, once on each set of files, and then use any external diff routine to compare them.

In the following example, `main` represents the main codeline, and `r3.2` is a codeline that was originally branched from `main`:

```
p4 changes //depot/main/... > changes-in-main
p4 changes //depot/r3.2/... > changes-in-r98.4
diff changes-in-main changes-in-r98.4
```

You can use this to uncover which changes have been made to `r98.4` that haven't been integrated back into `main`.



---

## Appendix A **Installing Perforce**

---

This appendix outlines how to install a Perforce server for the first time.

This appendix is mainly intended for people installing an evaluation copy of Perforce for trial use; if you're installing Perforce for production use, or are planning on extensive testing of your evaluation server, we strongly encourage you to read the detailed information in the *Perforce System Administrator's Guide*.

### **Getting Perforce**

---

Perforce requires at least two executables: the server (`p4d`), and any of the Perforce client programs (for instance, `p4` on UNIX, `p4.exe` or `p4win.exe` on Windows).

The programs are available from the Downloads page on the Perforce web site:

<http://www.perforce.com/perforce/loadprog.html>

Go to the web page, select the files for your platform, and save the files to disk.

### **Installing Perforce on UNIX**

---

Although `p4` and `p4d` can be installed in any directory, on UNIX the Perforce client programs typically reside in `/usr/local/bin`, and the Perforce server is usually located either in `/usr/local/bin` or in its own server root directory. Perforce client programs can be installed on any machine that has TCP/IP access to the `p4d` host.

To limit access to the Perforce server files, it is recommended that `p4d` be owned and run by a Perforce user account that has been created for that purpose.

Only a few additional steps need to be performed before `p4` and `p4d` can be run. They are described in detail in the following sections. Briefly:

- Download the files and make them executable,
- create a root directory for the Perforce files,
- provide a TCP/IP port to `p4d`,
- provide the hostname of the Perforce server and the `p4d` port number to the Perforce client program(s), and
- start the Perforce server (`p4d`).

## Download the files and make them executable

On UNIX (or MacOS X), you'll also have to make the Perforce executables (`p4` and `p4d`) executable. After downloading the programs, use the `chmod` command to make them executable:

```
chmod +x p4
chmod +x p4d
```

## Creating a Perforce server root directory

Perforce stores all of its data in files and subdirectories of its own root directory, which can reside anywhere on the server system. This directory is called the *server root*. This directory should be owned by the account that runs `p4d`, and can be named anything at all. The only necessary permissions are `read`, `write`, and `execute` for the user who invokes `p4d`.

For security purposes, `read` and `write` access to the server root should be restricted to prevent anyone but the account owner from reading, modifying or even listing the actual depot files. To ensure that temporary files cannot be read by unauthorized users, set the `umask(1)` file creation-mode mask of the account owner to a value that will not permit other users to read the contents of the server root directory or its files.

For security purposes, you are strongly advised not to run `p4d` as `root` or any other privileged user.

The environment variable `P4ROOT` should be set to point to the server root. Alternatively, the `-r root_dir` flag can be provided when `p4d` is started to specify a server root directory. The Perforce client programs never use this directory directly, and do not need to know the value of `P4ROOT`; the `p4d` server is the only process which uses the `P4ROOT` environment variable.

Unlike `P4ROOT`, the environment variable `P4PORT` is used by both the Perforce server and Perforce clients, and should be set on both. Its use is discussed in the next two sections.

## Telling the Perforce server which port to listen to

The `p4d` server and Perforce client programs communicate with each other via TCP/IP. When `p4d` starts, it listens (by default) on port `1666`. The Perforce client (also by default) assumes that its `p4d` server is located on host `perforce`, and is listening on port `1666`.

If `p4d` is to listen on a different port, the port can be specified with the `-p port_num` flag when starting `p4d` (for instance, `p4d -p 1818`), or the port can be set with the `P4PORT` environment or registry variable.

## Starting the Perforce server

After `p4d`'s `P4PORT` and `P4ROOT` environment variables have been set, `p4d` can be run in the background with the command:

```
p4d &
```

Although this command is sufficient to run `p4d`, other flags (for instance, those that control such things as error logging, checkpointing, and journaling), can be provided. These flags (and others) are discussed in the *Perforce System Administrator's Guide*.

## Stopping the Perforce server

If you are running Perforce 99.2 or later, use the command

```
p4 admin stop
```

to shut down the Perforce server.

If you are running an earlier version of Perforce, you'll have to find the process ID of the `p4d` server and kill it manually from the UNIX shell. The use of `kill -15 (SIGTERM)` is preferable to `kill -9 (SIGKILL)`, as the database could be left in an inconsistent state if `p4d` happened to be in the middle of updating a file when a `SIGKILL` signal was received.

With the introduction of `p4 admin stop` in Release 99.2, the practice of manually killing the `p4d` server has become obsolete.

## Telling Perforce clients which port to talk to

By this time, your Perforce server should be up and running; see “Connecting to the Perforce Server” on page 17 for information on how to set up your environment to allow Perforce's client programs to talk to the server.

## Installing Perforce on Windows

---

Installation of Perforce on Windows is handled by the installer. You can get the installer by downloading it from the Downloads page of the Perforce web site.

- Install Perforce client software (“User install”).

This option allows for the installation of `p4.exe` (the Perforce command line client), `p4win.exe` (P4Win, the Perforce Windows client), and `p4scc.dll` (Perforce's implementation of the Microsoft common SCM interface).

- Install Perforce as either a Windows server or service as appropriate. (“Administrator typical” and “Administrator custom” install).

These options allow for the installation of both the Perforce client software as well as the Perforce Windows server (`p4d.exe`) and service (`p4s.exe`) executables.

You can also use either of these options to automatically upgrade an existing Perforce server or service running under Windows.

- **Uninstall Perforce:** removes the Perforce server, service, and client executables, registry keys, and service entries. The database and depot files in your server root, however, are preserved.

## Terminology note: Windows services and servers

In most cases, it makes no difference whether the Perforce server program was installed on UNIX, as an NT service, or as an NT server. Consequently, the terms “Perforce server” and “p4d” are used interchangeably to refer to “the process which handles requests from Perforce clients”. In cases where the distinction between an NT server and an NT service is important, the distinction is made.

On UNIX systems, there is only one Perforce “server” program (p4d) responsible for this back-end task. On Windows, however, this back-end program can be started either as an NT service (p4s.exe), which can be set to run at boot time, or as an NT server (p4d.exe), which is invoked from an MS-DOS prompt.

The Perforce service (p4s.exe) and the Perforce server (p4d.exe) executables are copies of each other; they are identical apart from their filenames. When run, they use the first three characters of the name with which they were invoked (either p4s or p4d) to determine their behavior. (For instance, invoking copies named p4smyservice.exe or p4dmyserver.exe will invoke a service and a server, respectively.)

In most cases, you will want to install Perforce as a service, not a server. For a more detailed discussion of the distinction between the two options, see the *Perforce System Administrator's Guide*.

## Starting and stopping Perforce on Windows

If you're running Perforce as a service under Windows, it will start when the machine boots. You can configure it within the **Services** applet in the **Control Panel**.

If you're running Perforce as a server under Windows, invoke p4d.exe from an MS-DOS command prompt. The flags under Windows are the same as those under UNIX.

If you are running Perforce 99.2 or above, whether as a service or a server, use the command

```
p4 admin stop
```

to shut down the service or server. Only a Perforce superuser may use this command.

For older revisions of Perforce, you'll have to shut down services by using the **Services** applet in the **Control Panel**, and servers running in MS-DOS windows by typing Ctrl-C in the window or clicking on the icon to **Close** the window.

While these options will work with both Release 99.2 and earlier versions of Perforce, they are not necessarily “clean”, in the sense that the server or service is shut down abruptly. With the availability of the `p4 admin stop` command in 99.2, their use is no longer recommended.



---

## Appendix B Environment Variables

---

This table lists all the Perforce environment variables and their definitions.

You'll find a full description of each variable in the *Perforce Command Reference*.

Variable	Definition
P4CHARSET	For internationalized installations only, the character set to use for Unicode translations
P4CLIENT	Name of current client workspace
P4CONFIG	File name from which values for current environment variables are to be read
P4DIFF	The name and location of the diff program used by <code>p4 resolve</code> and <code>p4 diff</code>
P4EDITOR	The editor invoked by those Perforce commands that use forms
P4HOST	Name of host computer to impersonate. Only used if the <code>Host :</code> field of the current client workspace has been set in the <code>p4 client</code> form.
P4JOURNAL	A file that holds the database journal data, or <code>off</code> to disable journaling.
P4LANGUAGE	This variable is reserved for system integrators.
P4LOG	Name and path of the file to which Perforce server error and diagnostic messages are to be logged.
P4MERGE	A third-party merge program to be used by <code>p4 resolve</code> 's merge option
P4PAGER	The program used to page output from <code>p4 resolve</code> 's diff option
P4PASSWD	Stores the user's password as set in the <code>p4 user</code> form
P4PORT	For the Perforce server, the port number to listen on; for the <code>p4 client</code> , the name and port number of the Perforce server with which to communicate
PWD	The directory used to resolve relative filename arguments to <code>p4</code> commands
P4ROOT	Directory in which <code>p4d</code> stores its files and subdirectories
P4USER	The user's Perforce username
TMP	The directory to which Perforce writes its temporary files

## Setting and viewing environment variables

---

Each operating system and shell has its own syntax for setting environment variables. The following table shows how to set the `P4CLIENT` environment variable in each OS and shell:

OS or Shell	Environment Variable Example
UNIX: ksh, sh, bash	<code>P4CLIENT=value ; export P4CLIENT</code>
UNIX: csh	<code>setenv P4CLIENT value</code>
VMS	<code>def/j P4CLIENT "value"</code>
Mac MPW	<code>set -e P4CLIENT value</code>
Windows	<code>p4 set P4CLIENT=value</code>  (See the <code>p4 set</code> section of the <i>Perforce Command Reference</i> or run the command <code>p4 help set</code> to obtain more information about setting Perforce's registry variables in Windows).  Windows administrators running Perforce as a service can set variables for use by a specific service with <code>p4 set -S svcname var=value</code> .

To view a list of the values of all Perforce variables, use `p4 set` without any arguments.

On UNIX, this displays the values of the associated environment variables. On NT, this displays either the MS-DOS environment variable (if set), or the value in the registry and whether it was defined with `p4 set` (for the current user) or `p4 set -s` (for the local machine).



---

## Appendix C Glossary

---

Term	Definition
access level	A permission assigned to a user to control which Perforce commands the user can execute. See <i>protections</i> .
apple file type	Perforce file type assigned to Macintosh files that are stored using AppleSingle format, permitting the data fork and resource fork to be stored as a single file.
atomic change transaction	Grouping operations affecting a number of files in a single transaction. If all operations in the transaction succeed, all the files are updated. If any operation in the transaction fails, none of the files are updated.
base	The file revision on which two newer, conflicting file revisions are based.
binary file type	Perforce file type assigned to a non-text file. By default, the contents of each revision are stored in full and the file is stored in compressed format.
branch	<i>(noun)</i> A codeline created by copying another codeline, as opposed to a codeline that was created by adding original files. <i>branch</i> is often used as a synonym for <i>branch view</i> . <i>(verb)</i> To create a codeline branch with <code>p4 integrate</code> .
branch form	The Perforce form you use to modify a branch.
branch specification	Specifies how a branch is to be created by defining the location of the original codeline and the branch. The branch specification is used by the integration process to create and update branches. Client workspaces, labels, and branch specifications cannot share the same name.
branch view	A specification of the branching relationship between two codelines in the depot. Each branch view has a unique name, and defines how files are mapped from the originating codeline to the target codeline. See <i>branch</i> .
changelist	An atomic change transaction in Perforce. The changes specified in the changelist are not stored in the depot until the changelist is submitted to the depot.
changelist form	The Perforce form you use to modify a changelist.

Term	Definition
changelist number	The unique numeric identifier of a changelist.
change review	The process of sending email to users who have registered their interest in changes made to specified files in the depot.
checkpoint	A copy of the underlying server metadata at a particular moment in time. See <i>metadata</i> .
client form	The Perforce form you use to define a client workspace.
client name	A name that uniquely identifies the current client workspace.
client root	The root directory of a client workspace. If two or more client workspaces are located on one machine, they cannot share a root directory.
client side	The right-hand side of a mapping within a client view, specifying where the corresponding depot files are located in the client workspace.
client view	A set of mappings that specifies the correspondence between file locations in the depot and the client workspace.
client workspace	Directories on the client computer where you work on file revisions that are managed by Perforce. By default this name is set to the name of the host machine on which the client workspace is located; to override the default name, set the <code>P4CLIENT</code> environment variable. Client workspaces, labels, and branch specifications cannot share the same name.
codeline	A set of files that evolve collectively. One codeline can be branched from another, allowing each set of files to evolve separately.
conflict	<p>One type of conflict occurs when two users open a file for edit. One user submits the file; after which the other user can't submit because of a conflict. The cause of this type of conflict is two users opening the same file.</p> <p>The other type of conflict is when users try to merge one file into another. This type of conflict occurs when the comparison of two files to a common base yields different results, indicating that the files have been changed in different ways. In this case, the merge can't be done automatically and must be done by hand. The type of conflict is caused by non-matching <i>diffs</i>.</p> <p>See <i>file conflict</i>.</p>

Term	Definition
counter	A numeric variable used by Perforce to track changelist numbers in conjunction with the review feature.
default changelist	The changelist used by Perforce commands, unless a numbered changelist is specified. A default pending changelist is created automatically when a file is opened for edit.
default depot	The depot name that is assumed when no name is specified. The default depot name is <code>depot</code> .
deleted file	In Perforce, a file with its head revision marked as deleted. Older revisions of the file are still available.
delta	The differences between two files.
depot	A file repository on the Perforce server. It contains all versions of all files ever submitted to the server. There can be multiple depots on a single server.
depot root	The root directory for a depot.
depot side	The left side of any client view mapping, specifying the location of files in a depot.
depot syntax	Perforce syntax for specifying the location of files in the depot.
detached	A client computer that cannot connect to a Perforce server.
diff	A set of lines that don't match when two files are compared. A <i>conflict</i> is a pair of unequal diffs between each of two files and a common third file.
donor file	The file from which changes are taken when propagating changes from one file to another.
exclusionary mapping	A view mapping that excludes specific files.
exclusionary access	A permission that denies access to the specified files.
file conflict	In a three-way file merge, a situation in which two revisions of a file differ from each other and from their base file.  Also: an attempt to submit a file that is not an edit of the head revision of the file in the depot; typically occurs when another user opens the file for edit after you have opened the file for edit.
file pattern	Perforce command line syntax that enables you to specify files using wildcards.
file repository	The master copy of all files; shared by all users. In Perforce, this is called the <i>depot</i> .

Term	Definition
file revision	A specific version of a file within the depot. Each revision is assigned a number, in sequence. Any revision can be accessed in the depot by its revision number, for example: <code>testfile#3</code> .
file tree	All the subdirectories and files under a given root directory.
file type	An attribute that determines how Perforce stores and diffs a particular file. Examples of file types are <code>text</code> and <code>binary</code> .
fix	A job that has been linked to a changelist.
form	Screens displayed by certain Perforce commands. For example, you use the Perforce change form to enter comments about a particular changelist and to verify the affected files.
full-file storage	The method by which Perforce stores revisions of binary files in the depot: every file revision is stored in full. Contrast this with <i>reverse delta storage</i> , which Perforce uses for <code>text</code> files.
get	An obsolete Perforce term: replaced by <i>sync</i> .
group	A list of Perforce users.
have list	The list of file revisions currently in the client workspace.
head revision	The most recent revision of a file within the depot. Because file revisions are numbered sequentially, this revision is the highest-numbered revision of that file.
integrate	To compare two sets of files (for example, two codeline branches) and <ul style="list-style-type: none"><li>• determine which changes in one set apply to the other,</li><li>• determine if the changes have already been propagated,</li><li>• propagate any outstanding changes.</li></ul>
Inter-File Branching	Perforce's proprietary branching mechanism.
job	A user-defined unit of work tracked by Perforce. The job template determines what information is tracked. The template can be modified by the Perforce system administrator
job specification	A specification containing the fields and valid values stored for a Perforce job.
job view	A syntax used for searching Perforce jobs.
journal	A file containing a record of every change made to the Perforce server's metadata since the time of the last checkpoint.

Term	Definition
journaling	The process of recording changes made to the Perforce server's metadata.
label	A named list of user-specified file revisions.
label view	The view that specifies which file names in the depot can be stored in a particular label.
lazy copy	A method used by Perforce to make internal copies of files without duplicating file content in the depot. Lazy copies minimize the consumption of disk space by storing references to the original file instead of copies of the file.
license file	Ensures that the number of Perforce users on your site does not exceed the number for which you have paid.
list access	A protection level that enables you to run reporting commands but prevents access to the contents of files.
local depot	Any depot located on the current Perforce server.
local syntax	The operating-system-specific syntax for specifying a file name.
lock	A Perforce file lock prevents other clients from submitting the locked file. Files are unlocked with the <code>p4 unlock</code> command or submitting the changelist that contains the locked file.
log	Error output from the Perforce server. By default, error output is written to standard error. To specify a log file, set the <code>P4LOG</code> environment variable, or use the <code>p4d -L</code> flag.
mapping	A single line in a view, consisting of a left side and a right side that specify the correspondences between files in the depot and files in a client, label, or branch. The left side specifies the depot file and the right side specifies the client files.  (See also <i>client view</i> , <i>branch view</i> , <i>label view</i> ).
MD5 checksum	The method used by Perforce to verify the integrity of archived files.
merge	The process of combining the contents of two conflicting file revisions into a single file.
merge file	A file generated by Perforce from two conflicting file revisions.
metadata	The data stored by the Perforce server that describes the files in the depot, the current state of client workspaces, protections, users, clients, labels, and branches. Metadata includes all the data stored in the server except for the actual contents of the files.

<b>Term</b>	<b>Definition</b>
modification time	The time a file was last changed.
nonexistent revision	A completely empty revision of any file. Syncing to a nonexistent revision of a file removes it from your workspace. An empty file revision created by deleting a file and the <code>#none</code> revision specifier are examples of nonexistent file revisions.
numbered changelist	A pending changelist to which Perforce has assigned a number.
open file	A file that you are changing in your client workspace.
owner	The Perforce user who created a particular client, branch, or label.
P4	The Perforce command line interface program, and the command you issue to execute Perforce commands from the operating system command line.
P4D	The program on the Perforce server that manages the depot and the metadata.
P4Diff	A Perforce application that displays the differences between two files. <code>P4Diff</code> is the default application used to compare files during the file resolution process.
P4Win	The Perforce Windows client, a Windows Explorer-style application that enables you to perform Perforce operations and view results graphically.
pending changelist	A changelist that has not been submitted.
Perforce server	The Perforce depot and metadata on a central host. Also the program that manages the depot and metadata.
protections	The permissions stored in the Perforce server's protections table.
RCS format	Revision Control System format. Used for storing revisions of text files. RCS format uses reverse delta encoding for file storage. Perforce uses RCS format to store text files. See also <i>reverse delta storage</i> .
read access	A protection level that enables you to read the contents of files managed by Perforce.
remote depot	A depot located on a server other than the current Perforce server.
resolve	The process of resolving a file after the file is resolved and before it is submitted

Term	Definition
resolve	The process you use to reconcile the differences between two revisions of a file.
resource fork	One fork of a Macintosh file. (Macintosh files are composed of a resource fork and a data fork.) You can store resource forks in Performe depots as part of an AppleSingle file by using Performe's <code>apple</code> file type.
reverse delta storage	The method that Performe uses to store revisions of text files. Performe stores the changes between each revision and its previous revision, plus the full text of the head revision.
revert	To discard the changes you have made to a file in the client workspace.
review access	A special protections level that includes <code>read</code> and <code>list</code> accesses, and grants permission to run the <code>review</code> command.
review daemon	Any daemon process that uses the <code>p4 review</code> command. See also <i>change review</i> .
revision number	A number indicating which revision of the file is being referred to.
revision range	A range of revision numbers for a specified file, specified as the low and high end of the range. For example, <code>foo#5,7</code> specifies revisions 5 through 7 of file <code>foo</code> .
revision specification	A suffix to a filename that specifies a particular revision of that file. Revision specifiers can be revision numbers, change numbers, label names, date/time specifications, or client names.
server	In Performe, the program that executes the commands sent by client programs. The Performe server ( <code>p4d</code> ) maintains depot files and metadata describing the files, and tracks the state of client workspaces.
server root	The directory in which the server program stores its metadata and all the shared files. To specify the server root, set the <code>P4ROOT</code> environment variable.
status	For a changelist, a value that indicates whether the changelist is new, pending, or submitted. For a job, a value that indicates whether the job is open, closed, or suspended. You can customize job statuses.
submit	To send a pending changelist and changed files to the Performe server for processing.

Term	Definition
subscribe	To register to receive email whenever changelists that affect particular files are submitted.
super access	An access level that gives the user permission to run <i>every</i> Perforce command, including commands that set protections, obliterate files, verify files with an MD5 signature, and set up connections to remote depots.
symlink file type	A Perforce file type assigned to UNIX symbolic links. On non-UNIX clients, symlink files are stored as text files.
sync	To copy a file revision (or set of file revisions) from the depot to a client workspace.
target file	The file that receives the changes from the donor file when you are integrating changes between a branched codeline and the original codeline.
text file type	Perforce file type assigned to a file that contains only ASCII text. See also <i>binary file type</i> .
theirs	The revision in the depot with which the client file is merged when you resolve a file conflict. When you are working with branched files, <i>theirs</i> is the donor file.
three-way merge	The process of combining three file revisions. During a three-way merge, you can identify where conflicting changes have occurred and specify how you want to resolve the conflicts.
tip revision	In Perforce, the <i>head revision</i> . <i>Tip revision</i> is a term used by some other SCM systems.
trigger	A script automatically invoked by the Perforce server when changelists are submitted.
two-way merge	The process of combining two file revisions. In a two-way merge, you can see differences between the files but cannot see conflicts.
typemap	A Perforce table in which you assign Perforce file types to files.
user	The identifier that Perforce uses to determine who is performing an operation.
view	A description of the relationship between two sets of files. See <i>client view</i> , <i>label view</i> , <i>branch view</i> .



---

Term	Definition
wildcard	<p>A special character used to match other characters in strings. Perforce wildcards are:</p> <ul style="list-style-type: none"><li>• * matches anything except a slash</li><li>• . . . matches anything including slashes</li><li>• %d used for parameter substitution in views</li></ul>
write access	<p>A protection level that enables you to run commands that alter the contents of files in the depot. <code>write</code> access includes <code>read</code> and <code>list</code> accesses.</p>
yours	<p>The edited version of a file in the client workspace, when you resolve a file. Also, the target file when you integrate a branched file.</p>



---

# Index

---

## Symbols

#  
    forbidden in filenames 45  
#have  
    specifying the revision synced-to 47  
#head  
    specifying the latest revision 47  
#none 47  
%  
    forbidden in filenames 45  
%0 .. %n 34, 39  
\*  
    wildcard 34, 107  
+k flag  
    keyword expansion 50  
...  
    wildcard 34  
<  
    as operator in job queries > 108  
=  
    as operator in job queries 108  
>  
    as operator in job queries 108  
>>>>  
    as diff marker 62  
@  
    forbidden in filenames 45  
    syncing to a label's contents 89  
^  
    as operator in job queries 108  
**A**  
accepting  
    files when resolving 64  
access level 137  
adding files  
    to depot 25  
administration  
    depot configuration 76  
allwrite client option 41

architecture of Perforce 12  
atomic change transactions 79  
    branching and integration 98  
    defined 12, 137  
    example 29  
automating Perforce 55  
**B**  
base 137  
    *base*  
        defined, when resolving 61  
base file types 50  
baseless merge 98  
basics of Perforce 21  
BeOS  
    and symbolic links 51  
binary files  
    how files are stored 51  
    how revisions are stored 51  
    resolving 66  
branch specs  
    branching with 95  
    creating 95  
    deleting 99  
    example 95  
    exclusionary mappings allowed 97  
    usage notes 97  
    using with `p4 integrate -b` 95  
branch views  
    creating 95  
    defined 137  
branches  
    comparing files between 119  
    defined 137  
    deleting 99  
    listing files in 124  
    propagating changes between files 94, 96  
    when to create 91  
branching  
    automatic 95

- best practices 102
  - branch command 95
  - branch spec example 95
  - codelines and 91
  - copying files vs. 94
  - defined 91
  - files without common ancestors 98
  - introduced 14
  - manually, with `p4 integrate` 93
  - reporting 123
  - reverse integration and 98
  - two techniques 92
  - when to branch 91
  - white paper 102
- bug tracking 11
- build management 11
- C**
- carriage return 41
- change management 11
- change review 11, 138
- changelists
  - adding and removing files 80
  - associated jobs 83, 110, 122, 125
  - atomic change transactions and 79
  - automatic renumbering of 82
  - default 26, 79, 139
  - defined 137
  - deleting 83
  - files open in 80
  - introduced 12
  - jobs vs. 14, 103
  - moving files between 81
  - numbered 79, 81, 85, 138
  - `p4 reopen` command 83
  - pending 142
  - processed atomically 12
  - reporting 84, 120
  - scheduling resolves 59
  - status of 80
  - submitting 25
  - subscribing 144
- changes
  - conflicting 57, 62
  - how to propagate between codelines 95
  - propagating between codelines 94, 95
  - undoing with `p4 revert` 30
- changelists
  - reporting by user 84
- chunks, diff 61
- client
  - connecting to server 17
  - side of mapping 35
- client root 138
  - changing 40
  - defined 33
  - specifying 23
- client specification
  - defining 22
  - deleting 40
  - editing 40
- client syntax 44
- client view 138
  - changing 40
  - exclusionary mappings 38
  - introduced 12
  - specifying 23, 34
- client workspace
  - changing client root 40
  - changing view 40
  - comparing files against depot 68, 119
  - defined 21, 33, 138
  - displaying files 56
  - listing 126
  - moving files between client and server 12
  - options 41
  - `p4 have` command 31
  - populating 24, 89
  - refreshing 75
  - specifying 22
  - state of 33
  - switching between 22
- client/server 12
- clients
  - mapping to depot files 117
  - Perforce, purpose of 12, 17
  - user directories 12

- clobber client option 41
- codelines
  - branching and 91
  - comparing files between 119
  - defined 138
  - listing files in 124
  - propagating changes between 95
  - resolving differences between 94
  - when to branch 91
- command-line
  - common flags and `p4 help usage` 30
  - flags, common to all commands 73
  - specifying files on 43
- commands
  - applying to multiple revisions at once 49
  - forms and 55
  - reporting 30, 56, 68, 113, 115
- comparing
  - files 118
- compress client option 41
- concurrent development 11
- configuration
  - changing 71
- configuration files 71
- conflicting changes 57
- conflicts
  - file 13, 66, 139
- counter 139
- CR/LF translation
  - and `LineEnd` setting 42
- CR/LF translation 41
- creating
  - jobs 104
- `crLf` client option 41
- cross-platform development
  - line endings 42
- Customizing
  - job specification 105
- D**
- daemons
  - reporting 126
- default changelist
  - defined 139
  - introduced 25
  - using 79
- default depot 139
- default job specification 104
- defect tracking
  - interfacing with third-party products 112
  - jobs and 14
  - using jobs 103
- deleting
  - branch specs 99
  - client specifications 40
  - files 25, 28
  - jobs 112
  - labels 89
- delta
  - defined 139
  - how stored 143
- depot
  - adding files from workspace 25
  - changelists and 12
  - comparing against files in workspace 119
  - compressing files 51
  - copying files to workspace 24
  - default 139
  - defined 12, 21, 139
  - listing 126
  - local 141
  - mapping to client files 117
  - multiple 35
  - organizing 76
  - remote 142
  - side of mapping 35
  - syntax 44, 139
  - updating after working offline 75
- detached 139
- development
  - concurrent 11
- diff
  - chunks 61
  - differences between revisions 59
  - excluding 84
  - markers 62
  - suppressing display 122

- two algorithms used by Perforce 120

**directories**

- clients and 12
- removing empty 42

- donor file 139

**E****editing**

- client specifications 40
- files 25, 27

- email notification 138

**environment variables**

- how to set 136
- P4CHARSET 135
- P4CLIENT 22, 135
- P4CONFIG 71, 135
- P4DIFF 135
- P4EDITOR 22, 135
- P4HOST 135
- P4JOURNAL 135
- P4LOG 135
- P4MERGE 135
- P4PAGER 135
- P4PASSWD 72, 135
- P4PORT 18, 135
- P4ROOT 135
- P4USER 135
- PWD 135
- TMP 135

- error messages 19

**example**

- adding files to depot 25

**example set**

- for this manual 21

**examples**

- advanced integration 99
- branch spec 95
- combining file type modifiers 50
- creating a label 86
- filling forms with `-i` and `-o` 55
- linking jobs and changelists 110, 111
- `p4 job` 104
- populating a label 87
- propagating changes to branches 96

- RCS keyword expansion 54

- reporting and scripting 127

- reporting on jobs 113

- resolving file conflicts 64

- use of `%0 .. %n` wildcard 39

**exclusionary mappings**

- branch specs and 97

- client views and 38

- defined 139

**exclusive-open**

- vs. locking 66

**F****fields**

- jobviews and 108

**file conflict**

- defined 139

- introduced 13

- resolving 64

**file format**

- RCS 142

**file repository 139****file revision 140****file specifications**

- branching with 93

**file type modifiers**

- combining 50

- listed 52

**file types 50**

- `+l` 66

- `apple` 52

- `binary` 51

- compressed in depot 51

- determined by Perforce 50

- determining 51

- explained 50

- listed 51

- resource 52

- specifying 50, 52

- `symlink` 51

- `text` 51

**filenames**

- spaces in 45

**files**

- adding to depot 25
- binary 51, 66
- changelist revision specifier 47
- changelists and 12
- changing type 50
- client workspace 34
- client workspace revision specifier 47
- command line syntax 43
- command syntax 44
- commands for reporting 115
- comparing 68, 118
- conflicting 57, 68
- copying vs. branching 94
- deleting from depot 25, 28
- deleting from labels 89
- displaying branch contents 124
- displaying contents 118
- displaying integrated and submitted 124
- displaying label contents 90, 123
- displaying mappings 56
- displaying opened 56, 117
- displaying resolved but not submitted 68, 124
- displaying revision history 56
- displaying workspace contents 56
- donor 139
- editing 25, 27
- have revision 47
- head revision 47
- integrating 99
- label revision specifier 47
- listing with `p4 files` 56
- locking 66
- managed by Perforce 33
- merging 61, 62, 141
- modifying a changelist 80
- moving between changelists 81
- moving between client and server 12, 24
- multi-forked 52
- nonexistent revision 47
- opening 25, 80
- permissions 41, 52
- propagating changes between branches
  - 94
  - removing from workspace 47
  - renaming 77
  - reopening in other changelist 83
  - re-resolving 100
  - resolving 60, 64
  - result 61
  - specifying revision 46, 47, 48
  - specifying revision ranges 49
  - specifying type 50, 52
  - stored in RCS format 57
  - submitting changes to depot 25
  - target 144
  - text 50, 144
  - theirs 144
  - types of 50
  - undoing changes 30
  - wildcards 34
  - working offline 75
  - yours 145
- fix
  - defined 140
- fixes
  - jobs and changelists 125
- flags
  - common to all commands 30, 73
  - i flag 55
  - n flag 127
  - o flag 127
- forms
  - and standard input/output 55
  - automating processing of 55
  - `P4EDITOR` 22
  - using 55
- G**
  - getting started with Perforce 21
- H**
  - have list 140
  - have revision
    - defined 47
  - head revision
    - defined 47, 140
    - resolving conflicts 59

help

- displaying command help 30
- displaying view syntax 31
- p4 help command 30

history

- displaying revision history 116

host

- Perforce server 17

I

-i flag

- filling forms with standard input 55

installation

- on WIndows 131
- UNIX 129

integrate

- defined 140
- lazy copy 141

integration

- advanced functions 99
- displaying integrated files 124
- displaying submitted integrated files 124
- files without common ancestors 98
- forcing 100
- previewing 124
- reporting commands 102
- reverse 98
- specific file revisions 99
- specifying direction of 98
- technical explanation 100

Inter-File Branching 140

- defined 14
- use of 91

J

job specification

- customizing 105
- default 104
- defined 140

job tracking 11, 14, 103

jobs 103

- \* wildcard 107
- and third-party defect trackers 112
- changelists associated 110, 112, 122, 125
- changelists vs. 14, 103

creating 104

- defined 103
- deleting 112
- editing 104
- jobviews 106, 108, 110
- reporting 113, 124
- searching 106, 108

K

keywords

- expansion 50, 53
- RCS examples 54

L

label specifier

- without filenames 48

label view 86, 141

labels

- adding files to 86
- changelist numbers vs. 85
- changing owner of 86
- client workspaces and 89
- creating 85
- defined 85, 141
- deleting 89
- deleting files from 89
- displaying contents 123
- displaying contents of 90
- introduced 13
- locking 88
- reporting 90, 123
- unlocking 86

labelsync

- ownership required 86

lazy copy

- integration and 141

life cycle management 11

limitations

- description lengths 45
- valid filenames 45

line endings 42

LineEnd 42

linefeed convention 41

listing

- file contents 118



- files in a branch 124
- files in a label 123
- files resolved but not submitted 124
- integrated and submitted files 124
- jobs in system 124
- opened files 117
- local syntax 141
  - defined 43
  - wildcards and 45
- locked client option 41
- locking files 66, 141
  - p4 lock vs. +1 66
- M**
- Macintosh
  - file types 52
  - line-ending convention 43
  - linefeed convention 41
  - resource fork 52, 143
- mappings
  - client-side 138
  - conflicting 39
  - defined 141
  - depot and client sides 35
  - depot-side 139
  - displaying 56, 117
  - examples 37
  - exclusionary 97, 139
  - multiple 97
  - renaming client files 38
  - views and 35
- markers, difference 62
- merge 141
  - baseless 98
  - three-way 13, 62, 144
  - two-way 144
- merge*
  - defined 61
- merging
  - conflicting changes 62
  - files 141
- metadata 116, 141
- mode
  - of files in workspace 41, 52
- modtime client option 42
- moving files
  - between changelists 81
- multi-forked file 52
- multiple depots 35
- multiple mappings 97
- N**
- n flag
  - previewing commands 31, 127
- namespace
  - shared for labels, clients, branches, and depots 86
- network
  - data compression 41
- new changelist 80
- noallwrite client option 41
- noclobber client option 41
- nocompress client option 41
- nocrlf client option 41
- nomodtime client option 42
- normdir client option 42
- numbered changelists
  - creating 81
- O**
- o flag
  - scripting 55, 127
- offline
  - working with Perforce 74
- older revisions 46
- opened files
  - listing 56
- operators
  - in job queries 108
- options
  - client workspace 41
    - p4 resolve command 62
- overriding 50
- owner
  - changing label 86
- P**
- p4 admin
  - and Windows 132
- p4 commands

- add command 25
- admin command 131
- branch command 95, 99
- branches command 124
- change command 83
- changes command 84, 120
- client command 22, 34, 40
- common flags 73
- counters command 126
- delete command 25, 28
- depots command 126
- describe command 84, 122
- diff command 68, 74, 75, 118
- diff2 command 68, 118
- edit command 25, 27
- filelog command 56, 116, 124
- files command 51, 56, 90, 116
- fix command 83, 110, 111
- fixes command 125
- have command 31, 56, 118
- help command 30
- info command 19, 30
- integrate command 77, 93, 95, 96, 97, 98, 100, 124
- job command 104
- jobs command 124
- label command 85, 89
- labels command 90, 123
- labelsync command 86, 89
- lock command 66
- opened command 51, 56, 117
- passwd command 72
- print command 56, 118
- rename command 77
- reopen command 50, 83
- resolve command 60, 62, 66, 94, 124
- resolved command 68, 124
- revert command 30
- review command 126
- reviews command 126
- submit command 25, 60, 79, 80, 82
- sync command 24, 56, 59, 68, 75, 89, 118
- typemap command 50
- user command 73, 110
- users command 126
- where command 56, 118
- P4CHARSET 135
- P4CLIENT 22, 135
- P4CONFIG 71, 135
- p4d 142
  - host 17
  - port 17
  - purpose of 12, 17
- p4d.exe 132
- P4DIFF 135
- P4DTI 112
- P4EDITOR 22, 135
- P4HOST 135
- P4JOURNAL 135
- P4LOG 135
- P4MERGE 135
- P4PAGER 135
- P4PASSWD 72, 135
- P4PORT 18, 130, 135
- P4ROOT 130, 135
- p4s.exe 132
- P4USER 135
- parametric substitution 34, 39
- passwords 72, 73
- pending changelist 80, 83, 142
- Perforce clients
  - connecting to server 17
  - purpose of 12, 17
- Perforce labels 85
- Perforce server 142
  - connecting to 17
  - host 17
  - port 17
  - purpose of 12, 17
  - tracks state of client workspace 33
  - vs. service 132
  - working when disconnected from 74
- Perforce service
  - vs. server 132
- Perforce syntax
  - and wildcards 45

- defined 43
- permissions 137
  - on files in workspace 41, 52
- port
  - Perforce server 17
  - server 130
- pre-submit trigger 144
- previewing
  - integration results 124
  - label contents 88
    - n flag 127
  - resolve results 124
  - sync results 31, 56, 68
- propagating changes
  - to branches 96
- PWD 135
- R**
- RCS format
  - defined 142
  - files 57
- RCS keyword expansion 53
  - +k modifier 50
  - examples 54
- recent changelists
  - p4 changes command 84
- release management 11
- remote depot 142
- removing files
  - from depot 28
  - from workspace 47
- renaming files 77
- renumbering
  - of changelists, automatic 82
- reporting 115
  - basic commands 30, 56
  - branches 123
  - changelists 84, 120, 122
  - daemons and 126
  - file metadata 116
  - files 115
  - files affected by integration 124
  - integration 102, 123
  - jobs 113, 124
  - labels 90, 123
  - resolves 68
  - scripting 127
- repository
  - file 139
- resolve 60, 64, 124, 143
- resolves
  - between codelines 94
  - branching and 68
  - conflicting changes 57
  - defined 13
  - diff chunks 61
  - displaying files before submission 124
  - multiple 100
  - performing 60
  - preventing multiple 67
  - reporting 68
  - scheduling 59
- resource fork 52, 143
- result*
  - resolving 61
- reverse delta storage 143
- reverse integration 98
- revert
  - defined 143
  - example 30
- revision
  - base 137
  - diffs and 59
  - file 140
  - have 47
  - head 47, 140
  - history 56, 116
  - number 143
  - range 49, 143
  - specification 46, 48, 143
  - tip 144
- rmdir* client option 42
- root
  - client 33
- S**
- SCM 11
- scripting

- examples 55
  - o flag 127
- reporting 127
- searching
  - jobs 106
- server
  - connecting to 17
  - Perforce 12, 17, 142
  - port 130
  - stopping 131
  - verifying connection 19
  - vs. service 132
  - Windows 132
- server root
  - creating 130
  - defined 130
  - P4ROOT 130
- setting environment variables 136
- setting up
  - client workspaces 22
  - environment 17
  - p4 info 19
- shell
  - parsing wildcards 34
- software configuration management 11
- spaces
  - filenames 45
- special characters
  - filenames 45
- specification
  - revision 143
- standard input
  - filling forms with 55
- standard output
  - generating forms with 55
  - p4 print command 56
- stopping server
  - p4 admin 131
- storage
  - full-file vs. delta 140
  - reverse delta 143
- submit
  - defined 143

- submitted changelist 80
- submitting
  - multiple changes at once 29
- subscribe
  - to changelist 144
- symbolic links 51
  - file types and 50
  - non-UNIX systems and 51
- sync
  - forcing 75
  - preview 56, 68
- syntax
  - client 44
  - depot 44, 139
  - local 43, 141
  - Perforce syntax 43
  - specifying files 44
- system administration
  - checkpoint 138
  - groups 140
  - journal 140
  - reporting 126

**T**

- target files 144
- TCP/IP 12, 130
- text files 50, 51, 144
- theirs* 61, 144
- three-way merge
  - binary files and 66
  - defined 144
  - merge file generation 62
  - when scheduled 98
- time zones 48
- timestamps
  - preserving DLLs 53
- tip revision 144
- TMP 135
- tracking
  - defects 11
  - jobs 14
- translation
  - CR/LF 41
- trigger 144

two-way merge 144

typemap  
file types 50

## U

umask (1) 130

unicode 52

## UNIX

line-ending convention 43

linefeed convention 41

unlocked client option 41

## usage notes

integration 97

## user

listing submitted changelists 84

## users

email addresses 126

passwords 72

reporting on 126

## V

-v flag

diff markers 62

## variables

environment, how to set 136

version control 11

## view

branch 137

client 12, 23, 34, 138

defined 144

label 141

## views

branch, creating 95

conflicting mappings 39

examples of mappings 37

exclusionary mappings 38

help on 31

jobviews 106

label 86

mappings and 35

multiple mapping lines 37

renaming client files using mappings 38

wildcards 36

## W

## warnings

# and local shells 46

binary files and delta storage 51

changing client root 40

## white paper

best practices 102

branching 102

*Streamed Lines* 102

## wildcards

%0 .. %n 34, 39

\* 34

... 34

defined 34

escaping on command line 80

jobviews 107

local shell considerations 34

local syntax 45

Perforce syntax 45

views 36

## Windows

and p4 admin 132

installing on 131

line-ending convention 43

linefeed convention 41

server 132

setting variables on a Windows service  
136

third-party DLLs 53

## workaround

spaces in filenames 46

working detached 74, 139

## workspace

client 12, 33, 89, 138

comparing files against depot 119

copying files from depot 24

displaying files 56

refreshing 75

## Y

*yours* 61, 145

