
Perforce 2003.2 User's Guide

December 2003

This manual copyright 1997-2003 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

Preface	About This Manual.....	11
	Administering Perforce?.....	11
	Please Give Us Feedback.....	11
Chapter 1	Product Overview.....	13
	Perforce Server and Perforce Client Programs.....	14
	Moving files between the clients and the server.....	14
	File conflicts.....	15
	Labeling groups of files.....	15
	Branching files.....	15
	Job tracking.....	16
	Change notification.....	16
	Protections.....	17
	Other Perforce Clients.....	17
	P4Win.....	17
	P4Web.....	17
	Merge Tools.....	17
	P4 resolve.....	18
	P4WinMerge.....	18
	Other merge utilities.....	18
	Defect Tracking Systems.....	18
	Perforce jobs.....	18
	P4DTI integrations with third-party defect trackers.....	19
	Plug-Ins, reporting and tool integrations.....	19
	IDE Plug-ins.....	19
	P4Report and P4SQL.....	20
	P4OFC.....	20
Chapter 2	Connecting to the Perforce Server.....	21
	Before you begin.....	21
	Setting up your environment to use Perforce.....	21

Telling Perforce clients where the server is..... 22
Verifying the connection to the Perforce server 23

Chapter 3 **Perforce Basics:
Quick Start..... 25**

Underlying concepts 25
 File configurations used in the examples..... 25
Setting up a client workspace 26
 Naming the client workspace 26
 Describing the client workspace to the Perforce server 26
Copying depot files into your workspace 28
Updating the depot with files from your workspace 29
 Adding files to the depot..... 30
 Editing files in the depot 31
 Deleting files from the depot 32
 Submitting with multiple operations 33
Backing out: reverting files to their unopened states 34
Basic reporting commands 34

Chapter 4 **Perforce Basics:
The Details..... 37**

Description of the Client Workspace..... 37
Wildcards 38
 Wildcards and “p4 add” 38
Mapping the Depot to the Client Workspace 38
 Multiple depots..... 39
 Using views 39
 Wildcards in views 41
 Types of mappings 41
Editing Existing Client Specifications..... 44
 Deleting an existing client specification..... 45
 Client specification options..... 45
 Multiple workspace roots for cross-platform work..... 47
 Line-ending conventions (CR/LF translation)..... 47
Referring to Files on the Command Line 48
 Local syntax..... 48

Perforce syntax.....	48
Providing files as arguments to commands	49
Wildcards and Perforce syntax.....	50
Name and String Limitations.....	50
Illegal characters in filenames and Perforce objects	50
Name and description lengths	51
Specifying Older File Revisions.....	51
Using revision specifications without filenames	54
Revision Ranges	54
File Types	55
Base file types.....	56
Forms and Perforce Commands.....	60
Reading forms from standard input;	
Writing forms to standard output.....	61
General Reporting Commands.....	61

Chapter 5 Perforce Basics: Resolving File Conflicts..... 63

RCS Format: How Perforce Stores File Revisions.....	63
Only the differences between revisions are stored.....	63
Use of “diff” to determine file revision differences.....	65
Scheduling Resolves of Conflicting Files	65
Why “p4 sync” to Schedule a Resolve?	65
How Do I Know When a Resolve is Needed?.....	66
Performing Resolves of Conflicting Files.....	66
File revisions used and generated by “p4 resolve”	67
Types of conflicts between file revisions	67
How the merge file is generated	68
The “p4 resolve” options.....	68
Using Flags with Resolve	
to Automatically Accept Particular Revisions	71
Binary files and “p4 resolve”	72
Locking Files to Minimize File Conflicts.....	72
Preventing multiple resolves with p4 lock	73
Preventing multiple checkouts with +l files.....	73
Resolves and Branching.....	74
Resolve Reporting.....	74

Chapter 6	Perforce Basics: Miscellaneous Topics	77
	Reconfiguring the Perforce Environment with \$P4CONFIG	77
	Perforce Passwords.....	78
	Command-Line Flags Common to All Perforce Commands	79
	Working Detached	81
	Finding changed files with “p4 diff”	81
	Using “p4 diff” to update the depot	81
	Refreshing files	82
	Recommendations for Organizing the Depot.....	82
	Renaming Files.....	83
	Revision histories and renamed files.....	83
Chapter 7	Changelists	85
	Working with the Default Changelist	86
	Creating Numbered Changelists Manually	87
	Working With Numbered Changelists	87
	Automatic Creation and Renumbering of Changelists	88
	When submit of the default changelist fails, the changelist is assigned a number	88
	Perforce May Renumber a Changelist upon Submission	89
	Deleting Changelists.....	89
	Changelist Reporting.....	90
Chapter 8	Labels	91
	Why not just use changelist numbers?	91
	Using labels.....	91
	Creating a new label.....	92
	Labeling all revisions in your workspace	92
	Tagging specific files and revisions with p4 labelsync.....	93
	Untagging files with p4 labelsync	93
	Previewing labelsync’s results.....	93
	Preventing inadvertent tagging and untagging of files	93
	Listing files tagged by a label	94
	Using label views.....	94
	Referring to files using a label	95

	Deleting labels.....	95
	Details: How p4 labelsync works	96
	Label Reporting.....	97
Chapter 9	Branching	99
	What is Branching?.....	99
	When to Create a Branch	99
	Perforce's Branching Mechanisms: Introduction	100
	Branching and Merging, Method 1:	
	Branching with File Specifications	101
	Creating branched files.....	101
	Propagating changes between branched files	102
	Propagating changes from branched files to the original files	103
	Branching and Merging, Method 2:	
	Branching with Branch Specifications.....	103
	Branch Specification Usage Notes.....	105
	Integration Usage Notes	106
	Deleting Branches.....	107
	Advanced Integration Functions.....	107
	Integrating specific file revisions	107
	Re-integrating and re-resolving files	108
	How Integrate Works	108
	The yours, theirs, and base files	108
	The integration algorithm	108
	Integrate's actions.....	109
	Integration Reporting.....	110
	For More Information	110
Chapter 10	Job Tracking	111
	Job Usage Overview	111
	Creating and editing jobs using the default job specification.....	112
	Creating and editing jobs with custom job specifications	113
	Viewing jobs by content with jobviews.....	114
	Finding jobs containing particular words.....	114
	Finding jobs by field values	115
	Using and escaping wildcards in jobviews	115
	Negating the sense of a query	115

- Using dates in jobviews.....116
- Comparison operators and field types.....116
- Linking Jobs to Changelists.....117
 - Linking jobs to changelists with the JobView: field.....117
 - Linking jobs to changelists with p4 fix.....118
 - Linking jobs to changelists from within the submit form118
 - Automatic update of job status119
 - What if there's no status field? 120
- Deleting Jobs..... 120
- Integrating with External Defect Tracking Systems 120
- Job Reporting Commands 120

Chapter 11 Reporting and Data Mining..... 123

- Files 123
 - File metadata 123
 - Relationships between client and depot files 125
 - File contents..... 126
- Changelists..... 130
 - Viewing changelists that meet particular criteria 130
 - Files and jobs affected by changelists 131
- Labels..... 132
- Branch and Integration Reporting..... 133
- Job Reporting..... 133
 - Basic job information 133
 - Jobs, fixes, and changelists 134
- Reporting for Daemons..... 135
- System Configuration..... 135
- Special Reporting Flags..... 136
- Reporting with Scripting 136
 - Comparing the change content of two file sets 137

Appendix A Installing Perforce 139

- Getting Perforce 139
- Installing Perforce on UNIX..... 139
 - Download the files and make them executable 140
 - Creating a Perforce server root directory..... 140
 - Telling the Perforce server which port to listen to..... 140

Starting the Perforce server.....	140
Stopping the Perforce server.....	141
Telling Perforce clients which port to talk to.....	141
Installing Perforce on Windows	141
Terminology note: Windows services and servers	142
Starting and stopping Perforce on Windows	142
Appendix B Environment Variables.....	143
Setting and viewing environment variables.....	144
Appendix C Glossary	145
Index	155

About This Manual

This is the *Perforce 2003.2 User's Guide*. It teaches the use of Perforce's Command-Line Client. Other Perforce clients such as P4Win (the Perforce Windows Client) are not discussed here. If you'd like documentation on other Perforce client programs, 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.

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.
- Perforce supports *distributed development*; users can work with files stored on a central server or with files replicated on a proxy server.
- Some *release management* facilities are offered; Perforce can track the file revisions that are part of a particular release.
- Bugs and system improvement requests can be tracked from entry to fix; this capability is known as *defect tracking* or *change management*.
- 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 functionality 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 open source 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. Perforce can be used to store revisions of a manual, to manage Web pages, or to store old versions of operating system administration files.

Perforce Server and Perforce Client Programs

Perforce has a client/server architecture, in which users at other computers are connected to one central machine, the *server*. Each user works on a client machine; at their command, a Perforce client program transfers files to and from the Perforce server. The client programs communicate with the server using TCP/IP.

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

The following programs do the bulk of Perforce's work:

- The Perforce Server (`p4d`) runs on the Perforce server machine. It manages the shared file repository, and keeps track of users, workspaces, and other Perforce metadata.

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

- Perforce client programs (for instance, `p4`) run on Perforce client machines. Client programs send users' requests to the Perforce Server (`p4d`) for processing, and communicate with `p4d` using 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 you or your system administrator have already installed both `p4` and `p4d`. You'll find installation instructions in the *Perforce System Administrator's Guide*, also available at our Web site.

Moving files between the clients and the server

Users create, edit, and delete files in their own directories on client machines; 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 using 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 can 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 approach allows users to simultaneously update all files related to a bug fix or a new feature.

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.

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

File conflicts

When two users edit the same file, their changes can 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 example, 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 name, such as `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 its revisions into a client workspace.

For details about labels, see Chapter 8, Labels.

Branching files

Thus far, it has been assumed that all changes to 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*[™] 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, what files were modified to implement the fix, who fixed it, and whether the fix has been propagated to other codelines. 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 the functionality that is 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 notification

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

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

To learn how to set up the change review daemon, integrate Perforce with third-party defect tracking systems, or write your own daemons, 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. Because 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*.

Other Perforce Clients

The Perforce Command-Line Client (p4) is not the only Perforce client program. Other Perforce client programs, including P4Win, the Perforce Windows Client, may be downloaded from the Perforce web site.

P4Win

The Perforce Windows Client provides a native Microsoft Windows user interface for all SCM tasks. Using the familiar Windows® Explorer look and feel, it shows you your work in progress at a glance and lets you point, click, drag, and drop your way through Perforce tasks.

For more about P4Win, see the product page at:

<http://www.perforce.com/perforce/products/p4win.html>

P4Web

The Perforce web client turns most any Web browser into a complete SCM tool. P4Web will work with a Perforce Server at Release 99.2 or newer, and runs on Unix, Macintosh, and Windows.

For more about P4Web, see the product page at:

<http://www.perforce.com/perforce/products/p4web.html>

Merge Tools

Interactive merge tools allow you to display the differences between file versions, simplifying the process of resolving conflicts that result from parallel or concurrent development efforts. Merge tools often use color-coding to highlight differences and some even include the option to automatically merge non-conflicting changes.

Perforce offers full support for both parallel and concurrent development environments. In situations where concurrent file check-out is not desirable, Perforce can be configured to restrict this capability to specific file types or file locations (for instance, management of digital assets in environments where concurrent development is not encouraged).

P4 resolve

Perforce's "p4 resolve" command includes built-in merge capability for the console environment.

P4WinMerge

P4WinMerge is Perforce's graphical three-way merge and conflict resolution tool for Windows. P4WinMerge uses the "3-pane" approach to display and edit files during the merge process.

P4WinMerge is a stand-alone Windows application; it does not require a Perforce Server when used by itself. However, when invoked from within a Perforce client program like the Perforce Command-Line Client, P4Win, or P4Web, a Perforce Server is necessary.

For more about P4WinMerge, see:

<http://www.perforce.com/perforce/products/p4winmerge.html>

Other merge utilities

Perforce is easily integrated with third-party merge tools and diff utilities. Users need only change an environment variable (such as P4MERGE or P4DIFF) to point to their merge tool of choice.

For more about using third-party merge tools with Perforce, see:

<http://www.perforce.com/perforce/products/merge.html>

Defect Tracking Systems

Perforce provides a number of options for defect tracking. In addition to providing basic built-in defect tracking, Perforce is integrated with several leading defect tracking systems. Activity performed by Perforce users can be automatically sent to your defect tracking system. Conversely, issues and status entered into your defect tracking system can be accessed by Perforce users.

Perforce jobs

Perforce's built-in defect tracking and reporting features are available to all Perforce users.

P4DTI integrations with third-party defect trackers

Although Perforce provides built-in defect tracking, some companies prefer to use the defect tracking system they've already got in place, or want to install a different defect tracker for use with Perforce.

Perforce Defect Tracking Integration (P4DTI) is an open source project specifically designed to integrate Perforce with other defect tracking systems by replicating Perforce jobs and changelist numbers to their equivalents in the other system.

P4DTI connects your defect tracking system to Perforce, so that you don't have to switch between your defect tracker and SCM tool and enter duplicate information about your work. P4DTI also links changes made in Perforce with defect tracker issues, making it easy to find out why a change was made, find the work that was done to resolve an issue, or generate reports relating issues to files or codelines.

Activity in your Perforce depot - enhancements, bug fixes, propagation of changes into release branches, and so forth - can be automatically entered into your defect tracking system by P4DTI. Conversely, issues and status entered into your defect tracking system - bug reports, change orders, work assignments, and so on, can be converted automatically to Perforce metadata for access by Perforce users. With P4DTI, you can integrate Perforce with any third-party defect tracking or process management software.

P4DTI uses Perforce's built-in "jobs" feature to mirror data in defect tracking systems. While Perforce jobs can be used without additional software for straightforward issue tracking, P4DTI lets you take advantage of third-party user interfaces, reporting tools, databases, and workflow rules to manage complex processes.

P4DTI runs on Unix and Windows. It can be used with a Perforce Server on any platform at Release 2000.2 or newer.

For more about using third-party defect tracking systems with Perforce, including a list of defect tracking systems for which P4DTI integrations have already been built, see:

<http://www.perforce.com/perforce/products/defecttracking.html>

Plug-Ins, reporting and tool integrations

IDE Plug-ins

Perforce IDE Plug-ins allow developers to work with Perforce from within integrated development environments (IDEs) such as Microsoft Developer Studio, Borland JBuilder, and Metrowerks CodeWarrior.

For more about Perforce IDE Plug-ins, see:

<http://www.perforce.com/perforce/products/plugins-ide.html>

P4Report and P4SQL

The Perforce Reporting System (P4Report) offers query and reporting capability for Perforce depots. P4Report also includes the Perforce SQL Command-Line Client (P4SQL). P4SQL can be used to execute SQL statements either interactively or using scripts.

Based on P4ODBC, the Perforce ODBC Data Source, P4Report can be used by ODBC-compliant reporting tools including Crystal Reports®, Microsoft® Access® and Excel®. P4Report can also be integrated with some defect tracking systems.

For more about P4Report and P4SQL, see:

<http://www.perforce.com/perforce/products/p4report.html>

P4OFC

The Perforce Plug-in for Microsoft Office (P4OFC) adds a “Perforce” menu to Microsoft Word, Microsoft Excel, and Microsoft Powerpoint. This menu provide easy access to common Perforce SCM commands, so that users never have to leave familiar office applications to work with documents under Perforce control.

For more about P4OFC, see:

<http://www.perforce.com/perforce/products/plugins-ofc.html>

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 139, 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 144 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 -0700 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.

In the text forms used by Perforce, field names always start in the leftmost column of text, and field values are indented with tabs or spaces. Perforce requires that there be at least one space or a tab prior to the contents of a form field.

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. See the general 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 37.

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 30.

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. To reduce the risk of overwriting changes by accident, you may want to preview a revert by using `p4 revert -n` before running `p4 revert`. The `-n` option reports what files would be reverted by `p4 revert` without actually reverting the files.

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.

Command	Meaning
<code>p4 help usage</code>	Describes command-line flags common to all Perforce commands.
<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:

Command	Meaning
<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.

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 reside.

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 43 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 workspaces spanning multiple drives

To specify a Perforce client workspace that spans multiple Windows drives, use a `Root:` of `null`, and specify the drive letters in the client workspace view. Use uppercase drive letters when specifying workspaces across multiple drives. For example:

```
Client: eds_win
Owner: edk
Description:
    Ed's Windows Workspace
Root: null
Options:      nomodtime noclobber
View:
    //depot/main/...      "//eds_win/C:/Current Release/..."
    //depot/rel1.0/...    //eds_win/D:/old/rel1.0/...
    //depot/rel2.0/...    //eds_win/D:/old/rel2.0/...
```


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_coredumps`, 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 44 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/file.c  //joe/proj1/newfile.c
```

The file //depot/proj1/file.c maps to //joe/proj1/newfile.c, but that same client file //joe/proj1/newfile.c maps back to the depot via the higher-precedence second line to //depot/proj1/file.c. 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 51) 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>	Note: 2000.2 or earlier only! 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>

Option	Choice	Default
[un] locked	<p>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>.)</p> <p>If <code>locked</code>, only the owner is able to use, edit, or delete the client spec. Note that a Perforce administrator is still able to override the lock with the <code>-f</code> (force) flag.</p>	unlocked
[no] modtime	<p>For files <i>without</i> the <code>+m</code> (modtime) file type modifier:</p> <ul style="list-style-type: none"> • For Perforce clients at the 99.2 level or earlier, if <code>modtime</code> 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. • For Perforce clients at the 2000.1 level or higher, if <code>modtime</code> 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. • If <code>nomodtime</code> is set, the modification date is the date and time of sync, regardless of Perforce client version. <p>For files <i>with</i> the <code>+m</code> (modtime) file type modifier:</p> <ul style="list-style-type: none"> • For Perforce clients at the 99.2 level or earlier, the <code>+m</code> modifier is ignored, and the behavior of <code>modtime</code> and <code>nomodtime</code> is as documented above. • 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 <code>modtime</code> or <code>nomodtime</code> on the client. 	<p><code>nomodtime</code> (i.e. date and time of sync) for most files.</p> <p>Ignored for files with the <code>+m</code> file type modifier.</p>
[no] rmdir	<p>Should <code>p4 sync</code> delete empty directories in a client if all files in the directory have been removed?</p>	<code>normdir</code>

Multiple workspace roots for cross-platform work

Users who work on more than one operating system, but who wish to use the same client workspace, may use the `AltRoots:` field in the `p4 client` form.

You may specify up to two alternate client workspace roots. Perforce compares the current working directory against the main `Root:` first, and then against the two `AltRoots:` if they exist. The first root to match the current working directory is used. If no roots match, the main root is used.

Note | If you are using a Windows directory in any of your client roots, you must specify the Windows directory as your main client `Root:` and specify your other workspace root directories in the `AltRoots:` field.

If `edk`'s current working directory is under `/usr/edk/elm`, then Perforce uses the UNIX path as his client workspace root, rather than `e:\porting\edk\elm`. This allows `edk` to use the same client workspace specification for both UNIX and Windows development.

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

If you are using multiple client workspace roots, you can always find out which workspace root is in effect by examining the `Client root:` as reported by `p4 info`.

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: LF only

Option	Meaning
mac	Macintosh-style: CR only
win	Windows-style: CR, LF
share	<p>Shared mode: Line endings are LF with any CR/LF pairs translated to LF-only style before storage or syncing with the depot.</p> <p>When you sync your client workspace, line endings will be LF. If you edit the file on a Windows machine, and your editor inserts CRs before each LF, the extra CRs 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/file.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
J*	Files in the current directory starting with J
*/help	All files called help in current subdirectories
...	All files under the current directory and its subdirectories
.../*.c	All such files ending in .c
/usr/edk/...	All files under /usr/edk
//weasel/...	All files on client (or depot) weasel
//depot/...	All files in the depot named depot
//...	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 wildcards, the pathname component separator (/), and the revision-specifying characters @ and #, are not allowed in file names, label names, or other identifiers:

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

Using spaces in file and path names

Use quotation marks to enclose depot-side or client side mappings of files or directories that contain spaces. For instance, the mapping:

```
"//depot/release 1.2/doc/..." "//eds_ws/1.2 documentation/..."
```

maps all files in `//depot/main/release 1.2/doc` into the `1.2 documentation` subdirectory of client workspace `eds_ws`.

Other Perforce objects, such as branch names, client names, label names, and so on, may be *specified* with spaces, but these spaces are 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. All names given to Perforce objects such as branches, clients, and so on, are 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 client 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>Warning! 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, offset from GMT, 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 can be applied to the base types to enable preservation of timestamps, support for RCS keyword expansion, file compression on the server, and more.

When you add a file, Perforce attempts to determine the type of the file automatically: when a file is opened with `p4 add`, Perforce first determines 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 unicode mode; for details, see your system administrator.)

Once set, a file's type is inherited from one revision to the next, but can be overridden or changed by opening the file with the `-t filetype` 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`. You can determine the type of an existing file by using `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 cannot be merged or compared.

Some file types are compressed to `gzip` format when stored 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.

Base file types

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
<code>apple</code>	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 http://www.perforce.com/perforce/technical.html	full file, compressed, AppleSingle format.
<code>resource</code>	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 <code>apple</code> file type instead. For full details, please see the Mac platform notes at http://www.perforce.com/perforce/technical.html	full file, compressed

Keyword	Description	Comments	Server Storage Type
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

File type modifiers

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 <code>\$Id\$</code> and <code>\$Header\$</code> keywords: This pair of modifiers exists primarily for backwards compatibility with versions of Perforce prior to 2000.1, and corresponds to the <code>+k (ktext)</code> modifier in earlier versions of Perforce.
+k	RCS keyword expansion	Expands RCS (Revision Control System) keywords. RCS keywords are case-sensitive. When using keywords in files, a colon after the keyword (e.g., <code>\$Id:\$</code>) is optional. Supported keywords are: <ul style="list-style-type: none"> • <code>\$Id\$</code> • <code>\$Header\$</code> • <code>\$Date\$</code> • <code>\$DateTime\$</code> • <code>\$Change\$</code> • <code>\$File\$</code> • <code>\$Revision\$</code> • <code>\$Author\$</code>

Modifier	Description	Comments
+l	Exclusive open (locking)	If set, only one user at a time will be able to open a file for editing. Useful for binary file types (e.g., graphics) where merging of changes from multiple authors is meaningless.
+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 purged from the depot upon submission of new revisions. 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.

File type keywords

Versions of Perforce prior to 99.1 used a set of keywords to specify file types. The following table lists the older keywords and their current base file types and modifiers:

Old Keyword	Description	Base Filetype	Modifiers
text	Text file	text	none
xtext	Executable text file	text	+x
ktext	Text file with RCS keyword expansion	text	+k
kxtext	Executable text file with RCS keyword expansion	text	+kx
binary	Non-text file	binary	none
xbinary	Executable binary file	binary	+x
ctext	Compressed text file	text	+C
cxtext	Compressed executable text file	text	+Cx
symlink	Symbolic link	symlink	none
resource	Macintosh resource fork	resource	none

Old Keyword	Description	Base Filetype	Modifiers
uresource	Uncompressed Macintosh resource fork	resource	+F
ltext	Long text file	text	+F
xltext	Executable long text file	text	+Fx
ubinary	Uncompressed binary file	binary	+F
uxbinary	Uncompressed executable binary file	binary	+Fx
tempobj	Temporary object	ubinary	+FSw
ctempobj	Temporary object (compressed)	cbinary	+Sw
xtempobj	Temporary executable object	ubinary	+FSwx
xunicode	Executable unicode	unicode	+x

Overriding file types with the typemap table

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`. Your system administrator can use the `p4 typemap` command to set up a table matching Perforce file types to file name specifications.

When 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.

Preserving timestamps with the modtime (+m) modifier

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 `+m` modifier is useful when developing using 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. If you use the `+m` modifier on a file, Perforce ignores the `modtime` ("file's timestamp at time of submission") or `nomodtime` ("date and time on the client at time of sync") options on the `p4 client` form when syncing the file, and always restore the file's original timestamp at the time of submit.

Expanding RCS keywords with the +k modifier

If you use the +k modifier to activate RCS keyword expansion for a file, RCS keywords are expanded as follows:

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

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. When in doubt about what fields can be modified, see the *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 then 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.

The following Perforce reporting commands generate information on depot files, not files within the client workspace. When files are specified in local or client syntax on the command line, Perforce uses the client workspace view to map the specified files to their locations in the depot.

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 workspace `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 `file.c` for edit.
2. Lisa opens the same file in her client for edit.
3. Ed and Lisa both edit their client workspace versions of `file.c`.
4. Ed submits a changelist containing `file.c`, and the submit succeeds.
5. Lisa submits a changelist with her version of `file.c`; 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 `file`. The head revision (`file#3`) looks like this:

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

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

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

And revision one might be a representation of this:

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

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

```
Reconstructed file#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 file#n
(text from the original version)
==== THEIR VERSION file#m
(text from their file)
==== YOUR VERSION file
(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:*

```
Acme 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:

```
Acme 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 `file` for edit.
2. Lisa opens the same file in her client for edit.
3. Ed and Lisa both edit their client workspace versions of `file`.
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 `file`.
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.

Locked files appear in the output of `p4 opened` with an indication of `*locked*`. On UNIX, you can find all locked files you have open with the following command:

```
p4 opened | grep "*locked*"
```

This lists all open files you have locked with `p4 lock`.

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 `+l` 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>	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 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.
<code>p4 diff2 file1 file2</code>	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. The diff routine used by <code>p4d</code> cannot be overridden.
<code>p4 sync -n [filenames]</code>	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.
<code>p4 resolved</code>	Reports which files have been resolved but not yet submitted.

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:

P4CLIENT	P4DIFF	P4EDITOR	P4USER	P4CHARSET
P4PORT	P4MERGE	P4PASSWD	P4HOST	P4LANGUAGE

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

Some 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`. The flags are as follows:

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/file.c</code> Opens file <code>file.c</code> for editing under client workspace <code>joe</code> .
<code>-C charset</code>	For servers in unicode mode, override the <code>P4CHARSET</code> variable.	<code>p4 -C utf8 print //depot/file</code>
<code>-d directory</code>	Specifies the current directory, overriding the environment variable <code>PWD</code> .	<code>p4 -d ~elm/src edit one two</code> Opens files <code>one</code> and <code>two</code> for edit; these files are found relative to <code>~elm/src</code> .

Flag	Meaning	Example
-G	Cause all output (and batch input for form commands using the <code>-i</code> option) to be formatted as marshalled Python dictionary objects	<code>p4 -G info</code>
-H <i>host</i>	Specify the host name, overriding the environment variable <code>P4HOST</code> .	<code>p4 -H host print //depot/file</code>
-L <i>language</i>	For servers with non-English error messages, override the <code>P4LANGUAGE</code> variable.	Reserved for system integrators.
-p <i>server</i>	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> .
-P <i>password</i>	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.
-s	Prepend a tag to each line of output so as to make output more amenable to scripting.	<code>p4 -s info</code>
-u <i>username</i>	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 works without the <code>-P</code> flag only if <code>bill</code> has not set a Perforce password.
-x <i>filename</i>	Instructs <code>p4</code> to read arguments, one per line, from the named file.	See "Working Detached" on page 81.
-V	Displays the version of the <code>p4</code> executable.	<code>p4 -V</code>

All Perforce commands can take these flags, even commands for which these flag usages are useless (for instance, `p4 -u bill -d /usr/joe help`). Other flags are available as well; these additional flags are command dependent. See 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 for 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 workspace 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 `zeus`, `athena`, and `apollo`, three subtrees might be created within the depot:

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

```
//depot/zeus/...    //joe/...
```

And Sarah, who's working on the athena and apollo projects, might set up her client workspace as:

```
//depot/athena/...    //sarah/athena/...
//depot/apollo/...    //sarah/apollo/...
```

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

Another way of solving the same problem is to have the Perforce system administrator create one depot for each project or branch. See the *Perforce System Administrator's Guide* for details about setting up multiple depots.

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/one` and *to_file* is `d2/two`, then `one` is moved to the `d2` directory, and is renamed `two`. 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 administrators.

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/file1.c`,
 - User B types `p4 edit //depot/file1.c`,

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

User A's submit is rejected, since the file revision of `file1.c` 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 a Perforce administrator 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 set of tagged file revisions. You can use labels to reproduce the state of these files within a client workspace.

Labels provide a method of naming important sets of file revisions for future reference. For example, you might want to tag the file revisions that compose a particular release with the label `release2.0.1`. At a later time, you can retrieve all the tagged revisions into a client workspace by syncing the workspace to the label.

Create a label when you want to:

- keep track of all the file revisions contained in a particular release of software,
- distribute a particular set of file revisions to other users, or
- branch from a known set of file revisions.

Why not just use changelist numbers?

Labels share certain important characteristics with changelist numbers, as both refer to particular sets of file revisions, and both can be used to refer to all the revisions in the set. Labels have some important advantages over change numbers:

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

Using labels

Labeling files is a two-step process.

1. Create the label.
2. Tag the file revisions with the label.

Creating a new label

Use `p4 label labelname` to create a label. This command brings up a form similar to the `p4 client` form.

Label names share the same namespace as client workspaces, branches, and depots. Thus, a new label name must be distinct from any existing client workspace, branch, or depot name.

After you have created the label, you can use `p4 labelsync` to tag files.

For instance, you can create a new label `my_label` by typing:

```
p4 label my_label
```

The following form is displayed:

Label: <code>my_label</code>
Owner: <code>edk</code>
Description: Created by <code>edk</code> .
Options: <code>unlocked</code>
View: <code>//depot/...</code>

(The `View:` field denotes a label view; this field can be used to limit which files can be tagged with the label. You do not normally need to change the `View:` field; the only thing you need to do is update the label's description.)

After you have created the label, you can use it to tag any revision of any file under Perforce control. Only one revision of a given file can be tagged with a given label, but the same file revision may be tagged with multiple labels. Files are tagged with labels using the `p4 labelsync` command.

Labeling all revisions in your workspace

To tag file revisions in your client workspace with a label, use `p4 labelsync` as follows:

```
p4 labelsync -l my_label
```

All revisions in your client workspace are tagged with `my_label`, and any revisions previously tagged by `my_label` which are not present in your client workspace, are untagged.

Note | Anyone can use `p4 labelsync` with your label to tag or untag files. To prevent others (or yourself) from inadvertently changing which files are tagged by a label, `lock` the label. See “Preventing inadvertent tagging and untagging of files” on page 93 for details.

Tagging specific files and revisions with p4 labelsync

To tag a set of files (in addition to any revisions that may already have been tagged), use `p4 labelsync -a`, specifying the label and the desired file revisions.

For example, to tag all files in your client workspace that reside under `//depot/proj/rel1.0/hdrs` with the label `my_label`, use the following:

```
p4 labelsync -a -l my_label //depot/proj/rel1.0/hdrs/...
```

The revisions of files under `//depot/proj/rel1.0/hdrs/..` that also reside in your client workspace are tagged with the name `my_label`.

(This operation can also be thought of as “adding” the label to the files; hence the use of the `-a` flag with `p4 labelsync`.)

Untagging files with p4 labelsync

You can untag revisions with:

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

If you omit the *filepattern*, an implicit `//...` is assumed; all files formerly tagged by the label are untagged.

To untag a subset of tagged files, supply a file specification. For example, if you have previously tagged all revisions under `//depot/proj/rel1.0/...` with `my_label`, you can untag only the header files with:

```
p4 labelsync -d -l my_label //depot/proj/rel1.0/hdrs/*.h
```

Revisions of the `*.h` header files are no longer tagged with `my_label`.

(Just as you can add a label to files with `p4 labelsync -a`, you can delete the label from files with `p4 labelsync -d`.)

Previewing labelsync’s results

You can preview the results of `p4 labelsync` with `p4 labelsync -n`. This command lists the revisions that would be tagged, untagged, or retagged by the `labelsync` command without actually performing the operation.

Preventing inadvertent tagging and untagging of files

Using `p4 labelsync` with no file arguments tags the eligible files in your client workspace and any applicable label view, and untags all other files. This means that 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`.

Listing files tagged by a label

To list the revisions tagged with *labelname*, use `p4 files`, specifying the label name as follows:

```
p4 files @labelname
```

All revisions tagged with *labelname* are listed with their file type, change action, and changelist number. (This command is equivalent to `p4 files //...@labelname`)

Using label views

The `View:` field in the `p4 label` form limits the files that can be tagged with a label. To include files and directories in a label view, specify the files and directories to be included using depot syntax.

The default label view includes the entire depot (`//depot/...`), which allows you to tag any (and every) file in the depot with `p4 labelsync`. To prevent yourself from inadvertently tagging your entire workspace with a label, set the label's `View:` field to a subset of the depot.

Example: *Using a label view to control what files can be tagged*

Ed is working in a large client workspace, and wishes to tag a recently built set of binaries in this workspace with the label `build1.0`. He wants to ensure that he doesn't inadvertently tag the rest of his workspace by calling `p4 labelsync` without specifying a file argument.

He types `p4 label build1.0` and uses the label's `View:` field to restrict the scope of the label as follows:

```
Label: build1.0
Owner:  edk
Description:
    Created by edk.
Options:      unlocked
View:
    //depot/proj/re11.0/bin/...
```

After he saves from the editor, a new empty label `build1.0` is created. This label can tag only files in the `/re11.0/bin` directory.

With the default `View:` of `//depot/...`, he must type:

```
p4 labelsync -a //depot/proj/re11.0/bin/... -l labelname
```

With the new `View:`, Ed can use `p4 labelsync -l build1.0` to tag the desired files.

The `locked/unlocked` option in the `Options:` field can be used to prevent others from inadvertently retagging files with `p4 labelsync`. See "Preventing inadvertent tagging and untagging of files" on page 93).

Referring to files using a label

You can use a label name anywhere you can refer to files by revision (#1, #head), changelist number (@7381), or date (@2003/07/01).

If `p4 sync @labelname` is called with no file parameters, all files in the workspace that are tagged by the label are synced to the revision specified in the label. All files in the workspace that do not have revisions tagged by the label are deleted from the workspace. Open files or files not under Perforce control are unaffected. This command is equivalent to `p4 sync //...@labelname`.

If `p4 sync @labelname` is called with file arguments, as in `p4 sync files@labelname`, files in the user's workspace that are specified on the command line and also tagged by the label are updated to the tagged revisions.

Example: *Retrieving files tagged by a label into a client workspace.*

Lisa wants to retrieve some of the binaries tagged by Ed's build1.0 label into her client workspace. To get all files tagged by build1.0, she could type:

```
p4 sync //depot/...@build1.0
```

or even:

```
p4 sync @build1.0
```

Instead, she's interested in seeing only one platform's build from that label, so she types:

```
p4 sync //depot/proj/rell1.0/bin/osx/*@build1.0
```

and sees:

```
//depot/proj/rell1.0/bin/osx/server#6 - added as /usr/lisa/osx/server#6
//depot/proj/rell1.0/bin/osx/logger#12 - added as /usr/lisa/osx/logger#12
//depot/proj/rell1.0/bin/osx/install#2 - added as /usr/lisa/osx/install#2
<etc>
```

All files under //depot/proj/rell1.0/bin/osx that are tagged with Ed's build1.0 label and are also in Lisa's client workspace view are retrieved into her workspace.

Deleting labels

To delete a label, use the following command:

```
p4 label -d labelname
```

Details: How p4 labelsync works

The full syntax of the `p4 labelsync` command is:

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

The rules followed by `labelsync` to tag files with a label are as follows:

1. You must be the owner of the label 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 Perforce user name in the `p4 label` form. If you are the owner of a label, you may unlock the label by setting the `Options:` field (also in the `p4 label` form) to `unlocked`.

2. All files tagged with a label must be in the label view specified in the `p4 label` form. Any files or directories not included in a label view are ignored by `p4 labelsync`.
3. When `p4 labelsync` is used to tag a file revision with a label, the revision is tagged with the label if it is not already tagged with the label. If a different revision of the file is already tagged with the label, it is untagged and the newly specified revision is tagged. Any given file revision may be tagged by one or more labels, but only one revision of any file can be tagged with a given label at any one time.
4. If `labelsync` is called with no filename arguments, as in:

```
p4 labelsync -l labelname
```

then all the files in both the label view and the client workspace view are tagged with the label. The revisions tagged by the label are those last synced into the client workspace; these revisions (“the `#have` revisions”) can be seen in the `p4 have` list. Calling `p4 labelsync` this way removes the label from revisions it previously tagged unless those revisions are in your workspace.

5. When you call `p4 labelsync` with file pattern arguments, but the arguments contain no revision specifications, the `#have` revision is tagged with the label.
6. If you call `p4 labelsync` with file pattern arguments and those arguments contain revision specifications, the specified file revisions are tagged with the label.

Specifying a revision in this manner overrides all other ways of specifying a revision with a label; whether the client workspace contains a different revision than the one specified, (or doesn't contain the file at all), the revision specified on the command line is tagged with the label.

The following table lists variations of `p4 labelsync` as typed on the command line, their implicit arguments as parsed by the Perforce Server, and the sets of files from which `p4 labelsync` selects the intersection for tagging.

Call <code>p4 labelsync</code> with	Implicit Arguments
<code>-l label</code> (no files specified)	<code>-l label //myworkspace/...#have</code> Tags every file in your client workspace at the revision currently in your client workspace.
<code>-l label files</code> (files specified in local syntax, no revision specified)	<code>-l label [cwd]/files#have</code> Tags only the files in your client workspace that you specify, at the revision currently in your client workspace.
<code>-l label files#rev</code> (files specified in local syntax, specific revision requested)	<code>-l label [cwd]/files#rev</code> Tags only the files in your client workspace that you specify, at the revision you specify. Files must be in your client workspace view. You can use numeric revision specifiers here, or <code>#none</code> to untag files, or <code>#head</code> to tag the latest revision of a file, even if you haven't synced that revision to your workspace.
<code>-l label //files</code> (files specified in depot syntax, no revision specified)	<code>-l label //files#have</code> Tags only the files in the depot that you specify, at the revision currently in your client workspace, whether the files are in your client workspace view or not.
<code>-l label //files#rev</code> (files specified in depot syntax, specific revision requested)	<code>-l label //files#rev</code> Tags only the files in the depot that you specify, at the revision at the revision you specify, whether the files are in your client workspace view or not.

Label Reporting

The commands that list information about labels are:

Command	Description
<code>p4 labels</code>	List the names, dates, and descriptions of all labels known to the server
<code>p4 labels file#revrange</code>	List the names, dates, and descriptions of all labels that tag the specified revision(s) of <i>file</i> .

Command	Description
<code>p4 files @labelname</code>	Lists all files and revisions tagged by <i>labelname</i> .
<code>p4 sync -n @labelname</code>	Lists the revisions tagged by the label that would be brought into your client workspace, (as well as files under Perforce control that would be deleted from your client workspace because they are <i>not</i> tagged by the label), without updating your client workspace.

Chapter 9 **Branching**

Perforce's *Inter-File Branching*[™] 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 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.

Branching not only enables you to more easily track changes, it creates less overhead on the server. When you copy files with `p4 add`, you create two copies of the same file on the server. When you use branching, Perforce performs a “lazy copy” of the file, so that the depot holds only one copy of the original file and a record that a branch was created.

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 108).

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 manually. (In these cases, a `p4 resolve -as` or `p4 resolve -am` to accept the Perforce-recommended revision is usually sufficient. See "Using Flags with Resolve to Automatically Accept Particular Revisions" on page 71 for details.)

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 [tfiles]
```

If you omit the `tfiles` 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 [tfiles]`
- 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
<code>branch</code>	If the target file does not exist, it is opened for <code>branch</code> . The <code>branch</code> 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> .
<code>integrate</code>	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> .
<code>delete</code>	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 numbered 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 administrators.

Job Usage Overview

There are five related but distinct aspects of using jobs.

- The Perforce superuser or administrator 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
Status:   open
User:     edk
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
Status	open, closed, suspended, or new. An open job is one that has been created but has not yet been fixed. A closed job is one that has been completed. A suspended job is an open job that is not currently being worked on. 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.	new; changes to open after job creation form is closed.
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.
Date	The date the job was last modified, displayed as <code>YYYY/MM/DD HH/MM/SS</code>	The date and time at the moment this job was last modified.

Field Name	Description	Default
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 administrator 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:

```
# Custom Job fields:
# Job:      The job name. 'new' generates a sequenced job number.
# Status:  Either 'open', 'closed', or 'suspended'. Can be changed
# User:    The user who created the job. Can be changed.
# Date:    The date this specification was last modified.
# Type:    The type of the job. Acceptable values are
#          'bug', 'sir', 'problem' or 'unknown'
# Priority: How soon should this job be fixed?
#          Values are 'a', 'b', 'c', or 'unknown'
# Subsystem: One of server/gui/doc/mac/misc/unknown
# Owned_by: Who's fixing the bug
# Description: Comments about the job. Required.

Job:      new
Status:   open
User:     setme
Type:     setme
Priority:  unknown
Subsystem: setme
Owned_by: edk
Description:
          <enter description here>
```

Some of the fields have been set by the administrator 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 administrator 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 administrator.

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: edk
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, for example Linda Hopper
select	One of a set of values	A job's status: open/suspended/closed
date	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
word	The equality operator = must match the value in the word field exactly.
	The inequality operators perform comparisons in ASCII order.
text	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<filter'</code> , the job will match the jobview, because <code>bug<filter</code> .
line	See <code>text</code> , above.
select	The equality operator = matches a job if the value of the named field is the specified word. Inequality operators perform comparisons in ASCII order.
date	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.
- manually, by editing them within the `p4 submit` form.

Linking jobs to changelists with the `JobView:` field

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 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.

Linking jobs to changelists with p4 fix

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.*

You can use p4 fix to link a changelist to a job owned by another user.

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 by typing:

```
p4 fix -c 18 options-bug
```

Since changelist 18 has already been submitted, the job's status is changed to closed.

Linking jobs to changelists from within the submit form

You can also add jobs to changelists by editing the Jobs: field (or creating a Jobs: field if none exists) in the p4 submit form.

Any job can be linked to a changelist by adding it to a changelist's change form, or unlinked from a changelist by deleting the job from the 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: eds_ws
User: edk
Status: new
Description:
    Updating "File" I/O 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 job000125 from the form and then quits from the editor. The changelist is submitted without job000125 being associated with the changelist.

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.

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

Ed uses the reporting commands to read the details about job job000125. He fixes this problem, and a number of other bugs; when he next types p4 submit, he sees:

```
Change: new
Client: eds_ws
User: edk
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.

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>

To See a Listing of...	Use This Command:
...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 133, 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 51.

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>

Annotated file contents

Use `p4 annotate` to find out which file revisions or changelists affected lines in a text file.

By default, `p4 annotate` displays the file, each line of which is prepended by a revision number indicating the revision that made the change. The `-a` option displays all lines, including lines no longer present at the head revision, and associated revision ranges. The `-c` option displays changelist numbers, rather than revision numbers.

Example: *Using `p4 annotate` to track changes to a file*

A file is added (`file.txt#1`) to the depot, containing the following lines:

```
This is a text file.
The second line has not been changed.
The third line has not been changed.
```

The third line is deleted and the second line edited so that `file.txt#2` reads:

```
This is a text file.
The second line is new.
```

Finally, a third changelist is submitted, that includes no changes to `file.txt`. After the third changelist, the output of `p4 annotate` and `p4 annotate -c` look like this:

```
$ p4 annotate file.txt
//depot/files/file.txt#3 - edit change 153 (text)
1: This is a text file.
2: The second line is new.
$ p4 annotate -c file.txt
//depot/files/file.txt#3 - edit change 153 (text)
151: This is a text file.
152: The second line is new.
```

The first line of `file.txt` has been present since `file.txt#1`, as submitted in changelist 151. The second line has been present since `file.txt#2`, as submitted in changelist 152.

To show all lines (including deleted lines) in the file, use `p4 annotate -a` as follows:

```
$ p4 annotate -a file.txt
//depot/files/file.txt#3 - edit change 12345 (text)
1-3: This is a text file.
1-1: The second line has not been changed.
1-1: The third line has not been changed.
2-3: The second line is new.
```

The first line of output shows that the first line of the file has been present for revisions 1 through 3. The next two lines of output show lines of `file.txt` present only in revision 1. The last line of output shows that the line added in revision 2 is still present in revision 3.

You can combine the `-a` and `-c` options to display all lines in the file and the changelist numbers (rather than the revision numbers) at which the lines existed.

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>

To See the Differences between...	Use This Command:
...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.

If you want to pass the following flag to an external client diff program:	Then call <code>p4 diff</code> this way:
<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>

To See a List of Changelists...	Use This Command:
...limited to changelists that affect particular files, including only those changelists between revisions <i>m</i> and <i>n</i> of these files	<code>p4 changes filespec#m,#n</code>
...limited to those that affect particular files at each files revisions between labels <i>lab1</i> and <i>lab2</i>	<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 54.

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 diffs 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 diffs	<code>p4 describe -s changenum</code>
...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 133.

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 integrate	<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 resolve	<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 114)	<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>

To See a Listing of...	Use This Command:
...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.

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 `r98.4` is a codeline that was originally branched from `main`:

```
p4 changes //depot/main/... > changes-in-main
p4 changes //depot/r98.4/... > 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 at least one Perforce client program (such as `p4` on UNIX, or `p4.exe` or `p4win.exe` on Windows).

The server and client executables 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, ensure that the `p4d` executable is owned and run by a Perforce user account that has been created for that purpose.

After downloading `p4` and `p4d`, you need to do a few more things before you can start using Perforce:

1. Make the downloaded `p4` and `p4d` files executable.
2. Create a server root directory to hold the Perforce database and versioned files.
3. Provide a TCP/IP port to `p4d` to tell the Perforce server what port to listen to.
4. Provide the name of the Perforce server and the `p4d` port number to the Perforce client program(s).
5. Start the Perforce server (`p4d`).

Download the files and make them executable

On UNIX (or MacOS X), you must also make the Perforce executables (`p4` and `p4d`) executable. After downloading the programs, use the `chmod` command to make them executable, as follows:

```
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*.

To specify a server root, set the environment variable `P4ROOT` to point to the server root, or use the `-r root_dir` flag when invoking `p4d`. Perforce client programs never use the `P4ROOT` directory or environment variable; the `p4d` server is the only process that uses the `P4ROOT` variable.

A Perforce server requires no privileged access; there is no need to run `p4d` as `root` or any other privileged user. See the *System Administrator's Guide* for details.

The server root can be located anywhere, but the account that runs `p4d` must have `read`, `write`, and `execute` permissions on the server root and all directories beneath it. For security purposes, set the `umask(1)` file creation-mode mask of the account that runs `p4d` to a value that denies other users access to the server root directory.

Telling the Perforce server which port to listen to

The `p4d` server and Perforce client programs communicate with each other using TCP/IP. When `p4d` starts, it listens (by default) on port `1666`. The Perforce client assumes (also by default) that its `p4d` server is located on a host named `perforce`, listening on port `1666`.

If `p4d` is to listen on a different port, specify that port with the `-p port_num` flag when starting `p4d` (as in, `p4d -p 1818`), or set the port with the `P4PORT` environment or registry variable.

Unlike `P4ROOT`, the environment variable `P4PORT` is used by both the Perforce server and Perforce client programs, and must be set on both Perforce server machines and Perforce client workstations.

Starting the Perforce server

After setting `p4d`'s `P4PORT` and `P4ROOT` environment variables, start the server by running `p4d` in the background with the command:

```
p4d &
```

Although the example shown is sufficient to run `p4d`, other flags 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

To shut down a Perforce server, use the command:

```
p4 admin stop
```

to gracefully shut down the Perforce server. Only a Perforce superuser can use `p4 admin stop`.

If you are running a release of Perforce from prior to 99.2, you must find the process ID of the `p4d` server and kill the process manually from the UNIX shell. Use `kill -15 (SIGTERM)` instead of `kill -9 (SIGKILL)`, as `p4d` might leave the database in an inconsistent state if `p4d` is in the middle of updating a file when a `SIGKILL` signal is received.

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 21 for information on how to set up your environment to allow Perforce’s client programs to talk to the server.

Installing Perforce on Windows

To install Perforce on Windows, use the Perforce installer (`perforce.exe`) from the Downloads page of the Perforce web site.

The Perforce installer allows you to:

- Install Perforce client software (“User install”).

This option allows you to install `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 you to install Perforce client programs and the Perforce Windows server (`p4d.exe`) and service (`p4s.exe`) executables, or to automatically upgrade an existing Perforce server or service running under Windows.

Under Windows 2000 or higher, you must have Administrator privileges to install Perforce as a service, and Power User privileges to install Perforce as a server.

- **Uninstall Perforce:** remove the Perforce server, service, and client executables, registry keys, and service entries. The Perforce database and the depot files stored under your server root are preserved.

Terminology note: Windows services and servers

The terms “Perforce server” and “p4d” are used interchangeably to refer to “the process which handles requests from Perforce client programs”. 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, the back-end program can be started either as a Windows service (p4s.exe) process that runs at boot time, or as a server (p4d.exe) process that must be invoked from a command 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, the executables use the first three characters of the name with which they were invoked (either p4s or p4d) to determine their behavior. (For example, invoking copies of p4d.exe named p4smyservice.exe or p4dmyserver.exe invoke a service and a server, respectively.)

In most cases, it is preferable to install Perforce as a service, not a server. For a more detailed discussion of the distinction between services and servers, see the *Perforce System Administrator's Guide*.

Starting and stopping Perforce on Windows

If you install Perforce as a service under Windows, the service starts whenever the machine boots. Use the **Services** applet in the **Control Panel** to control the Perforce service's behavior.

If you install Perforce as a server under Windows, invoke p4d.exe from a command prompt. The flags for p4d under Windows are the same as those used under UNIX.

To stop a Perforce service (or server) at Release 99.2 or above, use the command:

```
p4 admin stop
```

Only a Perforce superuser can use p4 admin stop.

For older revisions of Perforce, shut down services manually by using the **Services** applet in the **Control Panel**. Shut down servers running in command prompt windows by typing CTRL-C in the window or by clicking on the icon to Close the command prompt window.

Although these manual shutdown options work with 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, the manual shutdown options are obsolete.

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> .
admin access	An access level that gives the user permission to run Perforce commands that override <i>metadata</i> , but do not affect the state of the server.
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> .

Term	Definition
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.
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.

Term	Definition
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>
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 <i>diffs</i> 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.

Term	Definition
exclusionary access	A permission that denies access to the specified files.
file conflict	<p>In a three-way file merge, a situation in which two revisions of a file differ from each other and from their base file.</p> <p>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.</p>
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> .
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	<p>To compare two sets of files (for example, two codeline branches) and</p> <ul style="list-style-type: none">• determine which changes in one set apply to the other;• determine if the changes have already been propagated,• propagate any outstanding changes.

Term	Definition
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.
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.

Term	Definition
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.
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 Client 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.

Term	Definition
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
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 Perforce depots as part of an AppleSingle file by using Perforce's <code>apple</code> file type.
reverse delta storage	The method that Perforce uses to store revisions of text files. Perforce 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>file#5,7</code> specifies revisions 5 through 7 of file <code>file</code> .

Term	Definition
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 Perforce, the program that executes the commands sent by client programs. The Perforce 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 Perforce server for processing.
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 every Perforce command, including commands that set protections, install triggers, or shut down the server for maintenance.
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.

Term	Definition
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> .
wildcard	A special character used to match other characters in strings. Perforce wildcards are: <ul style="list-style-type: none">• * matches anything except a slash• . . . matches anything including slashes• %d used for parameter substitution in views
write access	A protection level that enables you to run commands that alter the contents of files in the depot. Write access includes <code>read</code> and <code>list</code> accesses.
yours	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.

Index

Symbols

 forbidden in filenames 50
#have
 specifying the revision synced-to 52
#head
 specifying the latest revision 52
#none 52
%
 forbidden in filenames 50
%0 .. %n 38, 43
* (wildcard) 38, 115
+k flag
 keyword expansion 55
... (wildcard) 38
=
 operator in job queries 116
>
 operator in job queries 116
>>>>
 as diff marker 68
@
 forbidden in filenames 50
 syncing to a label's contents 95
^
 operator in job queries 115

A

accepting
 files when resolving 70
access level
 defined 145
adding files to depot 29
administration
 depot configuration 82
allwrite client option 45
annotate 127
architecture of Perforce 14
atomic change transactions 85
 branching and integration 106

 defined 14, 145

 example 33

automating Perforce 61

B

base

 defined 145

 resolving 67

base file types 55

baseless merge 106

basics of Perforce 25

BeOS

 symbolic links 56

binary files

 how files are stored 56

 how revisions are stored 55

 resolving 72

branch specs

 branching with 103

 creating 103

 deleting 107

 example 104

 exclusionary mappings allowed 105

 usage notes 105

 using with `p4 integrate -b` 104

branch views

 creating 103

 defined 145

branches

 comparing files between 129

 defined 145

 deleting 107

 listing files in 133

 propagating changes between files 102,
 105

 when to create 99

branching

 automatic 103

 best practices 110

 branch command 103

- branch spec example 104
- codelines and 99
- copying files vs. 102
- defined 99
- files without common ancestors 106
- introduced 15
- manually, with `p4 integrate` 101
- reporting 133
- reverse integration and 106
- two techniques 100
- when to branch 99
- white paper 110
- bug tracking 13
- build management 13
- C**
- carriage return 45
- change management 13
- change review 13
 - defined 146
- changelist number
 - defined 146
- changelists
 - adding and removing files 86
 - associated jobs 89, 131, 134
 - atomic change transactions and 85
 - automatic renumbering of 88
 - default 30, 85
 - default (defined) 147
 - defined 145
 - deleting 89
 - files open in 86
 - introduced 14
 - jobs 117
 - jobs vs. 16, 111
 - moving files between 87
 - numbered 85, 87, 91
 - `p4 reopen` command 89
 - pending (defined) 150
 - processed atomically 14
 - reporting 90, 130
 - reporting by user 90
 - scheduling resolves 65
 - status of 86
 - submitting 29
- changes
 - conflicting 63, 68
 - how to propagate between codelines 103
 - propagating between codelines 102, 103
 - undoing with `p4 revert` 34
- chunks, diff 67
- client files
 - mapping to depot files 125
- client programs 14, 21
- client root
 - changing 44
 - defined 37, 146
 - null 40
 - specifying 27
- client side (of mapping) 39
- client specification
 - defining 26
 - deleting 45
 - editing 44
- client syntax 49
- client view
 - changing 44
 - defined 146
 - exclusionary mappings 42
 - introduced 15
 - specifying 27, 38
- client workspace
 - changing client root 44
 - changing view 44
 - comparing files against depot 74, 129
 - defined 25, 37, 146
 - displaying files 62
 - listing 135
 - moving files between client and server 14
 - options 45
 - `p4 have` command 35
 - populating 29, 95
 - refreshing 82
 - spanning multiple drives 40
 - specifying 26
 - state of 37
 - switching between 26

- user directories 14
- client/server architecture 14
- `clobber` client option 45
- codelines
 - branching and 99
 - comparing files between 129
 - defined 146
 - listing files in 133
 - propagating changes between 103
 - resolving differences between 103
 - when to branch 99
- command-line
 - common flags and `p4 help usage` 35
 - flags, common to all commands 79
 - specifying files on 48
- commands
 - applying to multiple revisions at once 54
 - forms and 60
 - reporting 34, 61, 74, 120, 123
- comparing
 - files 128
- `compress` client option 45
- concurrent development 13
- configuration
 - changing 77
- configuration files 77
- conflicting changes 63
- conflicts
 - file 15, 72
 - file (defined) 148
- counter
 - defined 147
- CR/LF translation 45, 47
- creating
 - jobs 112
- `crlf` client option 45
- cross-platform development
 - line endings 47
- customizing
 - job specification 113
- D**
- daemons
 - reporting 135
- default changelist
 - defined 147
 - introduced 29
 - using 85
- default depot
 - defined 147
- default job specification 112
- defect tracking
 - interfacing with third-party products 120
 - jobs and 16
 - using jobs 111
- deleting
 - branch specs 107
 - client specifications 45
 - files 29, 32
 - jobs 120
 - labels 95
- delta
 - defined 147
- depot
 - adding files from workspace 29
 - changelists and 14
 - comparing against files in workspace 129
 - compressing files 56
 - copying files to workspace 28
 - default (defined) 147
 - defined 14, 25, 147
 - listing 135
 - local (defined) 149
 - mapping to client files 125
 - multiple 39
 - organizing 82
 - remote (defined) 151
 - side of mapping 39
 - syntax 49
 - syntax (defined) 147
 - updating after working offline 81
- detached
 - defined 147
- development
 - concurrent 13
 - distributed 13
- diff

- chunks 67
- differences between revisions 65
- excluding 90
- markers 68
- suppressing display 132
- two algorithms used by Perforce 129
- diffs
 - annotated 127
- directories
 - and spaces 51
 - client workspace 14
 - removing empty 46
- distributed development 13
- donor file
 - defined 147
- E**
- editing
 - client specifications 44
 - files 29, 31
- email notification 146
- environment variables
 - P4CHARSET 143
 - P4CLIENT 26, 143
 - P4CONFIG 77, 143
 - P4DIFF 143
 - P4EDITOR 26, 143
 - P4HOST 143
 - P4JOURNAL 143
 - P4LOG 143
 - P4MERGE 143
 - P4PAGER 143
 - P4PASSWD 78, 143
 - P4PORT 22, 143
 - P4ROOT 143
 - P4USER 143
 - PWD 143
 - setting 144
 - TMP 143
- error messages 23
- example set
 - for this manual 25
- examples
 - adding files to depot 30
 - advanced integration 107
 - branch spec 104
 - combining file type modifiers 55
 - creating a label 92
 - filling forms with `-i` and `-o` 61
 - linking jobs and changelists 117, 118
 - `p4 job` 112
 - propagating changes to branches 105
 - RCS keyword expansion 60
 - reporting and scripting 137
 - reporting on jobs 120
 - resolving file conflicts 70
 - use of `%0 .. %n` wildcard 43
- exclusionary mappings
 - branch specs and 105
 - client views and 42
 - defined 147
- exclusive-open
 - locking vs. 72
- F**
- fields
 - jobviews and 116
- file conflict
 - defined 148
 - introduced 15
 - resolving 70
- file format
 - RCS (defined) 151
- file repository
 - defined 148
- file revision
 - defined 148
- file specifications
 - branching with 101
- file type modifiers
 - combining 55
 - listed 57
- file types
 - +l 72
 - apple 56
 - binary 56
 - compressed in depot 56
 - determined by Perforce 55

- determining 55
 - explained 55
 - keywords 58
 - listed 56, 58
 - overview 55
 - resource 56
 - specifying 55, 57
 - symlink 56
 - text 56
- filenames
- and spaces 51
 - forbidden characters 50
 - spaces in 50
- files
- adding to depot 29, 30
 - annotated 127
 - binary 55, 72
 - changelist revision specifier 52
 - changelists and 14
 - changing type 55
 - client workspace 38
 - client workspace revision specifier 52
 - command line syntax 48, 49
 - commands for reporting 123
 - comparing 74, 128
 - conflicting 63, 74
 - copying vs. branching 102
 - deleting from depot 29, 32
 - deleting from labels 95
 - displaying branch contents 133
 - displaying contents 126
 - displaying integrated and submitted 133
 - displaying label contents 98, 132
 - displaying mappings 62
 - displaying opened 62, 125
 - displaying resolved but not submitted 74, 133
 - displaying revision history 62
 - displaying workspace contents 62
 - donor (defined) 147
 - editing 29, 31
 - have revision 52
 - head revision 52
 - integrating 107
 - label revision specifier 52
 - listing with `p4 files` 62
 - locked 73
 - locking 72
 - managed by Perforce 37
 - merging 67, 68
 - merging (defined) 150
 - modifying a changelist 86
 - moving between changelists 87
 - moving between workspace and server 14, 28
 - multi-forked 56
 - nonexistent revision 52
 - opening 29, 86
 - permissions 45, 57
 - propagating changes between branches 102
 - removing from workspace 52
 - renaming 83
 - reopening in other changelist 89
 - re-resolving 108
 - resolving 66, 70
 - result 67
 - specifying revision 51, 52, 53, 54
 - specifying type 55, 57
 - stored in RCS format 63
 - submitting changes to depot 29
 - target (defined) 152
 - text 55
 - text (defined) 152
 - theirs (defined) 152
 - types of 55
 - undoing changes 34
 - wildcards 38
 - working offline 81
 - yours (defined) 153
- fix
- defined 148
 - jobs and changelists 134
- flags
- common to all commands 35, 79
 - i flag 61

- n flag 136
 - o flag 136
- forms
 - automating processing of 61
 - P4EDITOR 26
 - standard input/output and 61
 - using 60
- G**
- getting started with Perforce 25
- group
 - defined 148
- H**
- have list
 - defined 148
- have revision
 - defined 52
- head revision
 - defined 52, 148
 - resolving conflicts 65
- help
 - displaying command help 34
 - displaying view syntax 35
 - p4 help command 34
- history
 - displaying revision history 124
- host
 - Perforce server 21
- I**
- i flag
 - filling forms with standard input 61
- installation
 - UNIX 139
 - Windows 141
- installing
 - on Windows 141
- integration
 - advanced functions 107
 - defined 148
 - displaying integrated files 133
 - displaying submitted integrated files 133
 - files without common ancestors 106
 - forcing 108
 - lazy copy (defined) 149
 - previewing 133
 - reporting commands 110
 - reverse 106
 - specific file revisions 107
 - specifying direction of 106
 - technical explanation 108
- Inter-File Branching
 - defined 16, 149
 - use of 99
- J**
- job specification
 - customizing 113
 - default 112
 - defined 149
- job tracking 13, 16, 111
- jobs 111
 - * wildcard 115
 - changelists 117
 - changelists associated 120, 131, 134
 - changelists vs. 16, 111
 - creating 112
 - defined 111
 - deleting 120
 - editing 112
 - jobviews 114, 116, 117
 - reporting 120, 133
 - searching 114, 116
 - third-party defect trackers and 120
- K**
- keywords
 - expansion 55, 57
 - RCS examples 60
 - specifying Perforce file types 58
- L**
- label specifier
 - without filenames 54
- label view 94
 - defined 149
- labels
 - changelist numbers vs. 91
 - changing owner of 96
 - client workspaces and 95
 - defined 149

- deleting 95
- deleting files from 95
- displaying contents 97, 98, 132
- introduced 15
- locking 93
- reporting 97, 132
- unlocking 96
- labelsync
 - ownership required 96
 - syntax 96
- lazy copy 102
 - defined 149
- lifecycle management 13
- limitations
 - description lengths 50
 - valid filenames 50
- line endings 47
- LineEnd 47
- linefeed convention 45
- listing
 - file contents 126
 - files in a branch 133
 - files in a label 132
 - files resolved but not submitted 133
 - integrated and submitted files 133
 - jobs in system 133
 - opened files 125
- local syntax
 - defined 48, 149
 - wildcards and 50
- locked client option 46
- locked files
 - finding 73
- locking files 72
 - defined 149
 - p4 lock vs. +l 72
- M**
- Macintosh
 - file types 56
 - line-ending convention 48
 - linefeed convention 45
 - resource fork 56
 - resource fork (defined) 151
- mappings
 - client-side (defined) 146
 - conflicting 44
 - defined 150
 - depot and client sides 39
 - depot-side (defined) 147
 - directories with spaces 51
 - displaying 62, 125
 - examples 41
 - exclusionary 105
 - exclusionary (defined) 147
 - multiple 105
 - renaming client files 43
 - views and 39
- markers, difference 68
- merge
 - baseless 106
 - defined 67, 150
 - three-way 15, 68
 - three-way (defined) 152
 - two-way (defined) 153
- merging
 - conflicting changes 68
 - files (defined) 150
- metadata 124
 - defined 150
- mode
 - files in workspace 45, 57
- modtime client option 46
- moving files
 - between changelists 87
- multi-forked file 56
- multiple depots 39
- multiple mappings 105
- N**
- n flag
 - previewing commands 34, 35, 136
- namespace
 - shared for labels, clients, branches, and depots 92
- network
 - data compression 45
- new changelist 86

noallwrite client option 45
noclobber client option 45
nocompress client option 45
nocrlf client option 45
nomodtime client option 46
normdir client option 46
numbered changelists
 creating 87

O

-o flag
 scripting 61, 136
offline
 working with Perforce 81
older revisions 51
opened files
 listing 62
operators
 job queries 116
options
 client workspace 45
 p4 resolve command 68
overriding 59
owner
 changing label 96

P

p4 admin
 and Windows 142
 stopping server with 141
p4 annotate 127
p4 commands
 add command 29, 30
 branch command 103, 107
 branches command 133
 change command 89
 changes command 90, 130
 client command 26, 38, 44
 common flags 79
 counters command 135
 delete command 29, 32
 depots command 136
 describe command 90, 131
 diff command 74, 81, 128
 diff2 command 74, 128

edit command 29, 31
filelog command 62, 124, 133
files command 55, 62, 98, 124
fix command 89, 117, 118
fixes command 134
have command 35, 62, 126
help command 34
info command 23, 34
integrate command 83, 101, 103, 105, 106,
 108, 133
job command 112
jobs command 133
label command 95
labels command 97, 132
labelsync command 95
lock command 72
opened command 55, 62, 125
passwd command 78
print command 62, 126
rename command 83
reopen command 55, 89
resolve command 66, 68, 72, 103, 133
resolved command 74, 133
revert command 34
review command 135
reviews command 135
submit command 29, 66, 85, 86, 88
sync command 29, 62, 65, 74, 82, 95, 126
typemap command 59
user command 79, 117
users command 135
where command 62, 126

P4CHARSET 143

P4CLIENT 26, 143

P4CONFIG 77, 143

p4d

 host 21

 port 21

 purpose of 14, 21

p4d.exe 142

P4DIFF 143

P4DTI 120

P4EDITOR 26, 143

- P4HOST 143
- P4JOURNAL 143
- P4LOG 143
- P4MERGE 143
- P4PAGER 143
- P4PASSWD 78, 143
- P4PORT 22, 143
 - and server 140
- P4ROOT 140, 143
- p4s.exe 142
- P4USER 143
- parametric substitution 38, 43
- passwords 78, 79
- pending changelist
 - defined 150
 - deleting 89
 - submitting 86
- Perforce client programs
 - connecting to server 21
 - purpose 14, 21
- Perforce server
 - and P4PORT 140
 - connecting to 21
 - defined 151
 - host 21
 - port 21
 - purpose of 14, 21
 - tracks state of client workspace 37
 - vs. service 142
 - working when disconnected from 81
- Perforce service
 - vs. server 142
- Perforce syntax
 - defined 48
 - wildcards 50
- perforce.exe 141
- permissions
 - files in workspace 45, 57
 - user (defined) 145
- port
 - for server 140
 - Perforce server 21
- pre-submit trigger
 - defined 153
- previewing
 - integration results 133
 - label contents 93
 - n flag 136
 - resolve results 133
 - revert results 34
 - sync results 35, 62, 74
- propagating changes
 - branches 105
- proxy 13
- PWD 143
- R**
- RCS format
 - defined 151
 - files 63
- RCS keyword expansion 57
 - +k modifier 55
 - examples 60
- recent changelists
 - p4 changes command 90
- release management 13
- remote depot
 - defined 151
- removing files
 - from depot 32
 - from workspace 52
- renaming files 83
- renumbering
 - changelists, automatic 88
- reporting
 - basic commands 34, 61
 - branches 133
 - changelists 90, 130, 131
 - daemons and 135
 - file metadata 124
 - files 123
 - integration 110, 133
 - jobs 120, 133
 - labels 97, 132
 - overview 123
 - resolves 74
 - scripting 136

- repository
 - file (defined) 148
 - resolve
 - between codelines 103
 - branching and 74
 - conflicting changes 63
 - default 70
 - defined 15, 151
 - detecting 66
 - diff chunks 67
 - displaying files before submission 133
 - multiple 108
 - performing 66
 - preventing multiple 73
 - previewing 133
 - reporting 74
 - scheduling 65
 - resource fork 56
 - defined 151
 - result
 - resolving 67
 - reverse delta storage
 - defined 151
 - reverse integration 106
 - revert
 - defined 151
 - example 34
 - revision
 - base (defined) 145
 - diffs and 65
 - file (defined) 148
 - have 52
 - head 52
 - head (defined) 148
 - history 62, 124
 - number (defined) 151
 - range 54
 - range (defined) 151
 - specification 51, 54
 - specification (defined) 152
 - tip (defined) 153
 - `rmdir` client option 46
 - root
 - client 37
- S**
- SCM 13
 - scripting
 - examples 61
 - o flag 136
 - reporting 136
 - searching
 - jobs 114
 - server
 - connecting to 21
 - Perforce 14, 21
 - Perforce (defined) 151
 - port 140
 - stopping with `p4 admin` 141
 - verifying connection 23
 - vs. service 142
 - Windows 142
 - server root
 - and `P4ROOT` 140
 - creating 140
 - defined 140
 - setting environment variables 144
 - setting up
 - client workspaces 26
 - environment 21
 - `p4 info` 23
 - shell
 - parsing wildcards 38
 - software configuration management 13
 - spaces
 - filenames 50
 - within filenames 51
 - special characters
 - filenames 50
 - specification
 - revision (defined) 152
 - standard input
 - filling forms with 61
 - standard output
 - generating forms with 61
 - `p4 print` command 62
 - stopping server

- with `p4 admin` 141
- storage
 - full-file vs. delta (defined) 148
 - reverse delta (defined) 151
- submit
 - defined 152
- submitted changelist 86
- submitting
 - multiple changes at once 33
- subscribe
 - defined 152
- symbolic links 56
 - file types and 55
 - non-UNIX systems 56
- sync
 - forcing 82
 - preview 62, 74
- syntax
 - client 49
 - depot 49
 - depot (defined) 147
 - local 48
 - local (defined) 149
 - Perforce syntax 48
 - specifying files 49
- system administration
 - checkpoint (defined) 146
 - groups (defined) 148
 - journal (defined) 149
 - reporting 135
- T**
- target files
 - defined 152
- TCP/IP 14
 - and port number 140
- text files 55, 56
 - defined 152
- theirs* 67
 - defined 152
- three-way merge
 - binary files and 72
 - defined 152
 - merge file generation 68
 - when scheduled 106
- time zones 53
- timestamps
 - preserving DLLs 59
- tip revision
 - defined 153
- TMP 143
- tracking
 - defects 13
 - jobs 16
- translation
 - CR/LF 45
- trigger
 - defined 153
- two-way merge
 - defined 153
- typemap
 - file types 59
- U**
- `umask (1)` 140
- unicode 57
- UNIX
 - line-ending convention 47
 - linefeed convention 45
- unlocked client option 46
- usage notes
 - integration 106
- users
 - email addresses 135
 - listing submitted changelists 90
 - passwords 78
 - reporting on 135
- V**
- `-v` flag
 - diff markers 68
- variables
 - environment, how to set 144
- version control 13
- views
 - branch (defined) 145
 - branch, creating 103
 - client 15, 27, 38
 - client (defined) 146

- conflicting mappings 44
- defined 153
- examples of mappings 41
- exclusionary mappings 42
- help on 35
- jobviews 114
- label 94
- label (defined) 149
- mappings and 39
- multiple mapping lines 42
- renaming client files using mappings 43
- wildcards 41

W

warnings

- # and local shells 51
- binary files and delta storage 56
- changing client root 44

white paper

- best practices 110
- branching 110
- Streamed Lines* 110

wildcards

- %0 .. %n 38, 43
- * 38
- ... 38
- defined 38
- escaping on command line 86
- jobviews 115
- local shell considerations 38
- local syntax 50
- Perforce syntax 50
- views 41

Windows

- and p4 admin 142
- installer 141
- installing on 141
- line-ending convention 48
- linefeed convention 45
- multiple drives 40
- server 142
- setting variables on a Windows service 144
- third-party DLLs 59

- working detached 81
- working detached (defined) 147
- workspace
 - client 14, 37, 95
 - client (defined) 146
 - comparing files against depot 129
 - copying files from depot 29
 - displaying files 62
 - refreshing 82
 - spanning multiple drives 40

Y

- yours* 67
 - defined 153