
Perforce 2002.2 System Administrator's Guide

December 2002

This manual copyright 1997-2002 Perforce Software.

All rights reserved.

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

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

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

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

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

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

Table of Contents

Preface	About This Manual.....	9
	Please Give Us Feedback	9
Chapter 1	Welcome to Perforce: Installing and Upgrading.....	11
	Getting Perforce	11
	Installing Perforce on UNIX.....	11
	Download the files and make them executable	12
	Create a Perforce server root directory	12
	Telling the Perforce server which port to listen to.....	13
	Telling Perforce client programs which port to talk to	13
	Starting the Perforce server.....	14
	Stopping the Perforce server.....	14
	Installing Perforce on Windows	14
	Terminology note: Windows services and servers	15
	Starting and stopping Perforce on Windows	16
	Upgrading a Perforce Server.....	16
	Using old client programs with a new server	16
	Important Notes for 2001.1 and later.....	16
	UNIX upgrades.....	17
	Windows upgrades	18
	Installation and Administration Tips.....	18
	Release and license information.....	18
	Observe proper backup procedures	19
	Use separate physical drives for server root and journal.....	19
	Use protections and passwords.....	19
	Allocate disk space for anticipated growth	20
	Managing disk space after installation	20
	Large filesystem support.....	21
	UNIX and NFS support.....	21
	Windows: Username and password required for network drives.....	22
	UNIX: Run p4d as a non-privileged user	22
	Logging errors.....	23
	Case sensitivity issues.....	23

	Tune for performance.....	23
Chapter 2	Supporting Perforce: Backup and Recovery	25
	Backup and Recovery Concepts	25
	Checkpoint files	26
	Journal files.....	28
	Versioned files	30
	Backup Procedures	31
	Recovery Procedures.....	33
	Database corruption, versioned files unaffected.....	33
	Both database and versioned files lost or damaged	35
	Ensuring system integrity after any restoration	37
 Chapter 3	 Administering Perforce: Superuser Tasks	 39
	Basic Perforce Administration	39
	Resetting user passwords.....	39
	Creating new users.....	39
	Preventing creation of new users	40
	Deleting obsolete users	41
	Reverting files left open by obsolete users.....	41
	Reclaiming disk space by obliterating files.....	42
	Deleting changelists and editing changelist descriptions	43
	File verification by signature	43
	Defining filetypes with p4 typemap	44
	Forcing operations with the -f flag.....	45
	Advanced Perforce Administration	46
	Running Perforce through a firewall	46
	Specifying IP addresses in P4PORT	49
	Running from inetd on UNIX.....	50
	Case sensitivity and multi-platform development.....	50
	Perforce server trace flags	52
	Migrating to a new machine	53
	Moving your versioned files and Perforce database	53
	Changing the IP address of your server.....	55

	Changing the hostname of your server.....	55
	Using Multiple Depots.....	56
	Distributed development notes.....	57
	Defining new depots.....	58
	Other depot operations	59
	Limiting access from other servers	60
	Users working with multiple depots.....	60
Chapter 4	Administering Perforce: Protections.....	61
	When Should Protections Be Set?.....	61
	Setting Protections with “p4 protect”	61
	The permission lines’ five fields.....	61
	Access levels.....	62
	Which users should receive which permissions?	63
	Default protections.....	64
	Interpreting multiple permission lines	64
	Exclusionary protections	65
	Granting Access to Groups of Users.....	66
	Creating and editing groups.....	66
	Groups and protections	66
	Deleting groups	67
	How Protections are Implemented	67
	Access Levels Required by Perforce Commands.....	68
Chapter 5	Customizing Perforce: Job Specifications.....	71
	The Default Perforce Job Template.....	71
	The Job Template’s Fields	72
	The Fields: field.....	73
	The Presets: field.....	75
	The Values: fields.....	75
	The Comments: field.....	76
	Caveats, Warnings, and Recommendations	77
	Example: A Custom Template	78
	Working with third-party defect tracking systems.....	79

	Using P4DTI - Perforce Defect Tracking Integration.....	80
	Building your own integration.....	80
	Getting more information	80
Chapter 6	Scripting Perforce: Daemons and Triggers.....	83
	Triggers.....	83
	Using triggers.....	85
	Triggers and security.....	86
	Triggers and Windows.....	86
	Daemons.....	87
	Perforce's change review daemon	87
	Creating other daemons	88
	Commands used by daemons	89
	Daemons and counters	90
	Scripting and buffering.....	90
Chapter 7	Tuning Perforce for Performance.....	91
	Tuning for Performance	91
	Memory.....	91
	Filesystem performance.....	91
	Disk space allocation.....	92
	Network	93
	CPU.....	93
	Diagnosing Slow Response Times.....	94
	Hostname vs. IP address	94
	Try p4 info vs. P4Win	94
	Windows wildcards.....	95
	DNS lookups and the hosts file	95
	Location of the "p4" executable	95
	Preventing Server Swamp	96
	Using tight views.....	96
	Assigning protections	97
	Limiting database queries	98
	Scripting efficiently	100
	Checkpoints for Database Tree Rebalancing.....	102

Chapter 8	Perforce and Windows	103
	Using the Perforce installer	103
	Upgrade notes.....	103
	Installation options.....	103
	Windows services vs. Windows servers.....	105
	Starting and stopping the Perforce service.....	106
	Starting and stopping the Perforce server	106
	Installing the Perforce service on a network drive.....	107
	Multiple Perforce services under Windows.....	107
	Windows configuration parameter precedence	108
	Resolving Windows-related instabilities.....	109
	Users having trouble with P4EDITOR or P4DIFF	110
Chapter 9	Perforce Proxy	111
	System Requirements.....	112
	Installing P4P.....	112
	UNIX	112
	Windows.....	112
	Running P4P.....	112
	Running as a Windows service	112
	P4P flags.....	113
	Administering P4P	114
	No backups required	114
	Managing disk space consumption.....	114
	Determining if your Perforce client is using the proxy	114
	P4P and protections	115
	Determining if specific files are being delivered from the proxy	115
	Maximizing performance improvement	116
	Network topologies versus P4P	116
	Pre-loading the cache directory for optimal initial performance	117
	Distributing disk space consumption.....	117
	Reducing server CPU usage by disabling file compression.....	117
Appendix A	Perforce Server (p4d) Reference.....	119
	Synopsis.....	119
	Syntax.....	119

Description	119
Exit Status	119
Options.....	119
Usage Notes.....	120
Related Commands	121
 Index.....	 123

This is the *Perforce 2002.2 System Administrator's Guide*.

It describes the installation, configuration, and operation of a Perforce server from an administrator's perspective. This includes the set of tasks typically performed by a "system administrator" (for instance, installation and configuration of the software, as well as ensuring uptime and data integrity), as well as those performed by a "Perforce administrator", including setting up Perforce users, setting Perforce depot access controls, resetting Perforce user passwords, and so on.

Since Perforce requires no special system permissions, the Perforce administrator typically does not require root-level access. Depending on your site's needs, your Perforce administrator need not be your system administrator.

Both the UNIX and Windows versions of the Perforce server are administered from the command line. If you are unfamiliar with the Perforce Command-Line Client, you will find the *Perforce Command Reference*, which describes the Perforce Command-Line Client in detail, to be a good companion to this guide.

Although this guide will teach you how to administer a Perforce server, it won't teach you the basics of actually using Perforce. You may also wish to consult the *Perforce User's Guide* for more information on Perforce from a user's perspective.

All our documentation is available from our web site at <http://www.perforce.com>.

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.

Welcome to Perforce: Installing and Upgrading

This chapter describes how to install a Perforce server or upgrade an existing installation.

Warning! If you're upgrading an existing installation to Release 2001.1 or later, see the notes in "Upgrading a Perforce Server" on page 16 before proceeding.

A short checklist of things to consider at installation time is offered, along with some basic security and administration tips. More detailed notes on administrative tasks are found in later chapters.

Windows | Windows administrators will note that many of the examples in this book are based on the UNIX version of the Perforce server. In almost all cases, they are common to both Windows and UNIX installations; Perforce's behavior is generally the same regardless of whether executed in a UNIX shell or from the MS-DOS command line.

Where the UNIX and Windows versions of Perforce differ, a note to that effect will be made. For Windows-specific information, see "Perforce and Windows" on page 103.

OS X | The material for UNIX also applies to Mac OS X.

Getting Perforce

Perforce requires at least two executables: the server (`p4d`), and any of the Perforce client programs (for instance, `p4` on UNIX, `p4.exe` or `p4win.exe` on Windows).

The programs are available from the Downloads page on the Perforce web site:

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

Go to the web page, select the files for your platform, and save the files to disk.

Installing Perforce on UNIX

Although `p4` and `p4d` can be installed in any directory, on UNIX the Perforce client programs typically reside in `/usr/local/bin`, and the Perforce server is usually located either in `/usr/local/bin` or in its own server root directory. Perforce client programs can be installed on any machine that has TCP/IP access to the `p4d` host.

To limit access to the Perforce server files, we recommend that `p4d` be owned and run by a Perforce user account that has been created for that purpose.

Once you’ve downloaded `p4` and `p4d`, you need to do a few more things before you can use Perforce. Briefly:

- Download the files and make them executable,
- Create a root directory for the Perforce files,
- Provide a TCP/IP port to `p4d`,
- Provide the name of the Perforce server and the `p4d` port number to the Perforce client program(s), and
- Start the Perforce server (`p4d`).

Download the files and make them executable

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

```
chmod +x p4
chmod +x p4d
```

Create a Perforce server root directory

Perforce stores all of its data in files and subdirectories of its own root directory, which can reside anywhere on the server system. This directory is called the *server root*. This directory should be owned by the account that runs `p4d`, and can be named anything at all. The only necessary permissions are `read`, `write`, and `execute` for the user who invokes `p4d`.

Since Perforce will store all submitted files under this directory, the size of the directory can become quite large. Disk space management is reviewed in “Installation and Administration Tips” on page 18 and described in more detail in “Disk space allocation” on page 92.

For security purposes, `read` and `write` access to the server root should be restricted to prevent anyone but the account owner from reading, modifying or even listing the actual depot files. To ensure that temporary files cannot be read by unauthorized users, set the `umask(1)` file creation-mode mask of the account owner to a value that will not permit other users to read the contents of the server root directory or its files.

You are strongly advised not to run `p4d` as `root` or any other privileged user. For more information, see the section entitled “UNIX: Run `p4d` as a non-privileged user” on page 22.

The environment variable `P4ROOT` should be set to point to the server root. Alternatively, the `-r root_dir` flag can be provided when `p4d` is started to specify a server root directory. The Perforce client programs never use this directory directly, and do not need

to know the value of `P4ROOT`; the `p4d` server is the only process which uses the `P4ROOT` environment variable.

Unlike `P4ROOT`, the environment variable `P4PORT` is used by both the Perforce server and Perforce client programs, and should be set on both. Its use is discussed in the next two sections.

Telling the Perforce server which port to listen to

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

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

Telling Perforce client programs which port to talk to

Perforce client programs need to know the TCP/IP port on which the `p4d` server program is listening. The easiest way to do this under UNIX is to set each Perforce user's `P4PORT` environment variable to `host:port#`, where `host` is the name of the machine on which `p4d` is running, and `port#` is the port on which `p4d` is listening.

Examples:

If <code>P4PORT</code> is...	Then...
<code>dogs:3435</code>	The client program uses the <code>p4d</code> server on host <code>dogs</code> listening at port 3435.
<code>x.com:1818</code>	The client program uses the <code>p4d</code> server on host <code>x.com</code> listening at port 1818.

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

Examples:

If <code>P4PORT</code> is...	Then...
<code>3435</code>	The client program will use the <code>p4d</code> server on its local host listening at port 3435.
<code><not set></code>	The client program will use the <code>p4d</code> server on the host named or aliased <code>perforce</code> listening on port 1666.

If your `p4d` host is not named `perforce`, you can choose to simplify life somewhat for your Perforce users by setting `perforce` as an alias to the true host name in their workstations' `/etc/hosts` files, or by doing so via Sun's NIS or Internet DNS.

Starting the Perforce server

After `p4d`'s `P4PORT` and `P4ROOT` environment variables have been set, `p4d` can be run in the background with the command:

```
p4d &
```

Although this command is sufficient to run `p4d`, other flags, for instance, those that control such things as error logging, checkpointing, and journaling, can be provided.

Example: *Starting a Perforce server.*

`P4PORT` can be overridden by starting `p4d` with the `-p` flag, and `P4ROOT` can be overridden by starting `p4d` with the `-r` flag. A journal file may be specified with the `-J` flag, and errors may be logged to a file specified with a `-L` flag. The startup command would then have this form:

```
p4d -r /usr/local/p4root -J /var/log/journal -L /var/log/p4err -p 1818 &
```

These flags (and others) are discussed in “Supporting Perforce: Backup and Recovery” on page 25. A complete list of server flags appears in the “Perforce Server (p4d) Reference” on page 119.

Stopping the Perforce server

If you are running Perforce 99.2 or above, use the command

```
p4 admin stop
```

to shut down the Perforce server. Only a Perforce superuser may use this command.

If you are running an earlier version of Perforce, you'll have to find the process ID of the `p4d` server and kill it manually from the UNIX shell. The use of `kill -15 (SIGTERM)` is preferable to `kill -9 (SIGKILL)`, as the database could be left in an inconsistent state if `p4d` happened to be in the middle of updating a file when a `SIGKILL` signal was received.

Installing Perforce on Windows

Installation of Perforce on Windows is handled by the installer. You can get the installer by downloading it from the Downloads page of the Perforce web site.

The Perforce installer (`perforce.exe`) allows you to:

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

This option allows for the installation of `p4.exe` (the Perforce Command-Line Client), `p4win.exe` (P4Win, the Perforce Windows Client), and `p4scc.dll` (Perforce’s implementation of the Microsoft common SCM interface).

- Install Perforce as either a Windows server or service as appropriate. (“Administrator typical” and “Administrator custom” install).

These options allow for the installation of both the Perforce client software as well as the Perforce Windows server (`p4d.exe`) and service (`p4s.exe`) executables.

You can also use either of these options to automatically upgrade an existing Perforce server or service running under Windows.

Under Windows 2000 or higher, you will need 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 database and depot files in your server root, however, will be preserved.

For more about installing on Windows, see “Using the Perforce installer” on page 103.

Terminology note: Windows services and servers

In most cases, it makes no difference whether the Perforce server program was installed on UNIX, as an NT service, or as an NT server. Consequently, the terms “Perforce server” and “`p4d`” are used interchangeably to refer to “the process which handles requests from Perforce clients”. In cases where the distinction between an NT server and an NT service are important, the distinction is made.

On UNIX systems, there is only one Perforce “server” program (`p4d`) responsible for this back-end task. On Windows, however, this back-end program can be started either as an NT service (`p4s.exe`), which can be set to run at boot time, or as an NT server (`p4d.exe`), which is invoked from an MS-DOS prompt.

The Perforce service (`p4s.exe`) and the Perforce server (`p4d.exe`) executables are copies of each other; they are identical apart from their filenames. When run, they use the first three characters of the name with which they were invoked (either `p4s` or `p4d`) to determine their behavior. (For example, invoking copies named `p4smyservice.exe` or `p4dmyserver.exe` will invoke a service and a server, respectively.)

In most cases, you will want to install Perforce as a service, not a server. For a more detailed discussion of the distinction between the two options, see “Windows services vs. Windows servers” on page 105.

Starting and stopping Perforce on Windows

If you're running Perforce as a service under Windows, it will be started when the machine boots. You can configure it within the **Services** applet in the **Control Panel**.

If you're running Perforce as a server under Windows, invoke `p4d.exe` from an MS-DOS command prompt. The flags under Windows are the same as those under UNIX.

If you are running Perforce 99.2 or above, whether as a service or a server, use the command

```
p4 admin stop
```

to shut down the service or server. Only a Perforce superuser may use this command.

For older revisions of Perforce, you'll have to shut down services by using the **Services** applet in the **Control Panel**, and servers running in MS-DOS windows by typing CTRL-C in the window or clicking on the icon to **Close** the window.

While these options will work with both Release 99.2 and earlier versions of Perforce, they are not necessarily "clean", in the sense that the server or service is shut down abruptly. With the availability of the `p4 admin stop` command in 99.2, their use is no longer recommended.

Upgrading a Perforce Server

Whether your server is running on Windows or UNIX, you *must* back up your server (see "Backup Procedures" on page 31) as part of any upgrade process.

Warning! If you are upgrading to 2001.1 or later, it is imperative that you read the notes pertaining to the 2001.1 upgrade.

Using old client programs with a new server

Although pre-98.1 Perforce client programs (`p4`, `p4.exe`, `p4win.exe`, and `p4scc.dll`) generally work with newer server versions with no trouble, some features in new server releases require client upgrades. In general, users with older client programs will still be able to use features available from the Perforce server at the client's release level, but will not be able to use the new server features offered by subsequent server upgrades.

Perforce's remote depot support is an exception: remote depot support is not guaranteed to work unless *all* of your Perforce servers are at or above Release 98.2.

Important Notes for 2001.1 and later

On small installations (installations with fewer than 1000 submitted changelists), the 2001.1 server automatically upgrades the underlying database for versions 98.2 and up.

On larger installations, you will have to upgrade the database manually. Although the new database will likely be smaller than the pre-2001.1 database, the upgrade process can be disk space-intensive. You will need approximately three times the size of the existing database to store the temporary files required by the upgrade.

Note | If you are disk space-constrained, see the Release Notes for a more precise estimate of the amount of disk space required.

By turning off journaling during the upgrade (by setting `P4JOURNAL` to `off`), it may be possible to reduce the amount of disk space required for the upgrade. (Remember to turn journaling back on when the upgrade is complete!)

If you are upgrading from Release 97.3 or earlier to 2001.1, you will likely have to make an intermediate checkpoint. Please contact Perforce technical support for assistance before upgrading your server.

UNIX upgrades

To upgrade your current Perforce server to a newer version, your Perforce license file must be current. Expired licenses will not work with upgraded servers. (This is not a problem if you are running a two-user installation with no license.)

You *must* back up your server as described on “Backup Procedures” on page 31 as part of any upgrade process.

It is a good idea to run `p4 verify` as part of your upgrade. See “Verifying during server upgrades” on page 44 for details.

Warning! Upgrading to Release 2001.1 or later will require an upgrade of your database files. Downgrading thereafter will require restoration from backups.

If you wish to keep your pre-2001.1 server available as a fallback option when performing your 2001.1 upgrade, you should back up your entire server root (including the `db.*` files) after stopping the server.

Upgrading from UNIX Release 98.2 or later

If you have a valid license (or require no license) and are upgrading from Release 98.2 or later:

1. Download the new `p4d` executable for your platform
2. Stop the current instance of `p4d`
3. Make a checkpoint and back up your old installation

4. Install the new `p4d` in the desired location
5. Run `p4d -xu` to upgrade the database.

Note | If your server has fewer than 1000 changes, the upgrade will run automatically. Larger installations will require that `p4d -xu` be run manually.

Either way, you must have sufficient disk space to complete the upgrade. The required amount is typically two to three times the size of the larger of the `db.have` or `db.integ` files.

The `db.have` and `db.integ` files reside in your `P4ROOT` directory.

6. Restart the new `p4d` with your site's usual parameters.

Your users should then be able to use the new server.

Windows upgrades

On Windows, download the installer (`perforce.exe`) and follow the installation dialog.

The upgrade process on Windows is extremely conservative; if any error condition occurs during the upgrade, you will always be able to revert to using your pre-upgrade Perforce server or service.

Note | If your server has fewer than 1000 changes, the upgrade will run automatically. Larger installations will require that `p4d -xu` be run manually.

Either way, you must have sufficient disk space to complete the upgrade. The required amount is typically two to three times the size of the larger of the `db.have` or `db.integ` files.

The `db.have` and `db.integ` files reside in your `P4ROOT` directory.

If you have any questions or difficulties during an upgrade, contact Perforce technical support.

Installation and Administration Tips

Release and license information

Perforce servers are licensed according to how many users they will support.

This licensing information lives in a file called `license` in the server root directory. It is a plain text file supplied by Perforce Software. Without the `license` file, the Perforce server will limit itself to two users and two client workspaces.

Current licensing information may be viewed by invoking `p4d -v` from the server root directory or by specifying the server root directory either on the command line (`p4d -v -r server_root`) or in the `P4ROOT` environment variable.

When the server is running, the license information may also be viewed with `p4 info`.

Version information will be displayed when invoking `p4d -v` or `p4 -v`.

Observe proper backup procedures

Regular backups of your Perforce data are vital. The key concepts are:

- Make sure journaling is active,
- Create checkpoints regularly, and
- Use `p4 verify` regularly.

See “Supporting Perforce: Backup and Recovery” on page 25 for a full discussion of backup and restoration procedures.

Use separate physical drives for server root and journal

Whether installing on UNIX or Windows, it is usually advisable to have your `P4ROOT` directory (that is, the directory containing your database and versioned files) on a different physical drive than your journal file.

By storing the journal on a separate drive, you can be reasonably sure that if a disk failure corrupts the drive containing `P4ROOT`, it will *not* affect your journal file. The journal file can then be used to restore any lost or damaged metadata.

Further details are available in “Supporting Perforce: Backup and Recovery” on page 25.

Use protections and passwords

Until you define a Perforce superuser, every Perforce user is a Perforce superuser, and can run any Perforce command on any file. The administrator who installs Perforce should use:

```
p4 protect
```

to define a Perforce superuser as soon as possible after installing the Perforce server. For more information, see “Administering Perforce: Protections” on page 61

Furthermore, until your users have passwords defined, any user will be able to impersonate any other Perforce user, either with the `-u` flag or by setting `P4USER` to a

Perforce user's username. Proper use of Perforce passwords (as described in the *Perforce User's Guide*) can protect against this. See the *Perforce User's Guide* for details.

To set (or reset) a user's password, use `p4 passwd username` (as a Perforce superuser), and enter the new password for the user, or invoke `p4 user -f username` (also while as a Perforce superuser) and enter the new password into the user specification form. The former command will only work in release 99.1 or later; the latter command will work under all releases from 97.3 onwards.

The security-conscious Perforce superuser will use `p4 protect` to make sure that no access higher than `list` is granted to non-privileged users, and require that each user have a Perforce password.

Allocate disk space for anticipated growth

In general, you'll need sufficient space in your `P4ROOT` directory to hold your depot files (that is, the files created by your users), and an additional 0.5K per user per file to hold the data describing the files, their status, and their histories. As a rule of thumb, you also probably want at least enough disk space to hold three times the size of your present collection of versioned files.

For a more detailed example of a disk sizing estimate, see "Disk space allocation" on page 92.

Managing disk space after installation

All of Perforce's versioned files reside in subdirectories of the server root, as do its database files, and (by default) the checkpoints and journals. The stored versioned file depots are grow-only, and this can clearly present disk space problems on high use systems. The following approaches may be used to remedy this:

- Tell Perforce to store the journal file on a separate physical disk. Use the `P4JOURNAL` environment variable or `p4d -J` to specify the location of the journal file.
- Checkpoint on a daily basis to keep the journal file short.
- Compress checkpoints, or use the `-z` option to tell `p4d` to compress them while creating them.
- Use the `-jc prefix` option with the `p4d` command to write the checkpoint to a different disk. Alternately, use the default checkpoint files, but back up your checkpoints and then delete them from the root directory. Old checkpoints are needed when recovering from a crash, and if your checkpoint and journal files reside on the same disk as your depot, a hardware failure could leave you without the ability to restore your database.
- On UNIX systems, some or all of the depot directories may be relocated to other disks by using symbolic links. Creation of symbolic links and movement of depot files to other volumes should be done only while the Perforce server is not running.

- Due to the nature of the implementation of their access methods, the database files themselves may become internally unbalanced, resulting in them taking up more space than necessary. The database files can sometimes be reduced in size by recreating them from a checkpoint. This should be done only if the database files are more than about 10 times the size of the checkpoint in total. See “Checkpoints for Database Tree Rebalancing” on page 102.

Large filesystem support

Earlier versions of the Perforce server, as well as some operating systems, limit Perforce database files (the `db.*` files in the `P4ROOT` directory which contain your metadata) to 2GB in size.

The `db.have` file holds label contents and the list of files currently opened in client workspaces, and tends to grow the most quickly. If you anticipate any of your Perforce database files growing beyond the 2GB level, you should install the Perforce server on a platform with support for large files.

As of this writing, the following combinations of operating system and Perforce server revision will support database files larger than 2GB:

Operating System	OS version:	Perforce Server Revision
Tru64 UNIX (a.k.a. Digital UNIX, OSF/1)	All versions	98.2/5713 or higher
FreeBSD	All versions	98.2/5713 or higher
Windows NT, 2000	All versions, SP6 recommended for NT	98.2/8127 or higher
SGI IRIX 6.2	All versions	98.2/5713 or higher
SGI IRIX 5.3	Only with the SGI- supplied <code>xfs</code> upgrade	98.2/5713 or higher <code>xfs</code> OS upgrade required
Solaris	2.6 and higher	98.2/7488 <i>compiled for 2.6 or higher</i>

UNIX and NFS support

The Perforce server process has been tested and is supported on the Solaris 2.6 implementation of NFS. Because Perforce client programs never directly access the files in `P4ROOT`, the only process needing access to `P4ROOT` is the `p4d` server itself.

Consequently, under Solaris 2.6 or higher, you can store your journal, log, depot, and `db.*` files on NFS-mounted filesystems.

Some issues still remain regarding file locking on non-commercial implementations of NFS (for instance, Linux and FreeBSD). On these platforms, we recommend that you store your journal, log, depot, and `db.*` files on a drive local to the server machine, not on an NFS-mounted volume.

These issues affect only the Perforce Server process (`p4d`). Perforce client programs, (such as `p4`, the Perforce Command-Line Client) have always been able to work with client workspaces on NFS-mounted drives, such as workspaces located in users' home directories.

Windows: Username and password required for network drives

By default, the Perforce service runs under the Windows local `System` account. Because Windows requires a real account name and password to access files on a network drive, if Perforce is installed as a service under Windows with `P4ROOT` pointing to a network drive, the installer will query for an account name and a password. The Perforce service will be configured with the supplied data and run as the specified user instead of `System`. (This account must have `Administrator` privileges on the machine.)

Although Perforce operates reliably with its root directory on a network drive, does so at a substantial performance penalty, as all writes to the database have to be performed over the network. For optimal performance, it is still best to install the Windows service to use local drives rather than networked drives.

For more information, see “Installing the Perforce service on a network drive” on page 107.

UNIX: Run `p4d` as a non-privileged user

While it is possible to run the Perforce server as `root`, it is highly inadvisable to do so. Sound administration practice demands that processes which don't require `root` access should never be run as `root`. For Perforce, this means that the owner of the `p4d` process should never be a privileged account.

Windows	On Windows, directory permissions are set securely by default; when running as a server, the Perforce server root is accessible only to the user who invoked the server from the MS-DOS command line. When installed as a service, the files are owned by the <code>LocalSystem</code> account, and are accessible only to those with <code>Administrator</code> access.
----------------	--

A good way to manage a Perforce installation on UNIX is to create a UNIX userid for it (for example, “`perforce`”) and (optionally) a UNIX group for it (for example, “`p4admin`”). The `umask(1)` command can be used to ensure that the server root (`P4ROOT`) and all files and directories beneath it are created as writable only by the UNIX user `perforce`, and (optionally) readable by members of the UNIX group `p4admin`.

The Perforce server (`p4d`), running as UNIX user `perforce`, can write to files in the server root, but none of your users will be able to overwrite its files. Access to read the files created by `p4d` (that is, the depot files, checkpoints, journals, and so on) can be granted to trusted users by making them members of the UNIX group `p4admin`.

Logging errors

The Perforce server's error output file can be specified with the `-L` flag to `p4d`, or can be defined in the environment variable `P4LOG`. If no error output file is defined, errors are dumped to `p4d`'s standard error.

Although `p4d` tries to ensure that all error messages reach the user, if an error occurs and the client program disconnects before the error is received, `p4d` will also log these errors to its error output.

The Perforce server also has trace flags used for debugging purposes. See “Perforce server trace flags” on page 52 for details.

Case sensitivity issues

Whether your Perforce server is running on Windows or UNIX, if your site is involved in cross-platform development (i.e. Perforce clients on both Windows and UNIX machines), your users will still need to be made aware of certain details regarding case sensitivity issues. See “Case sensitivity and multi-platform development” on page 50 for details.

Tune for performance

Perforce is a relatively light consumer of network traffic and CPU power. The most important variables determining performance will be the efficiency of your server's disk I/O subsystem and the number of files referenced in any given user-originated Perforce operation.

For more detailed performance tuning information, see “Tuning Perforce for Performance” on page 91.

Supporting Perforce: Backup and Recovery

The Perforce server stores two kinds of data: *versioned files* and *metadata*. Both are stored in the server's root directory.

- *Versioned files* are files submitted by Perforce users. Versioned files are stored in directory trees called *depots*. There is one subdirectory under the server's root directory for each depot in your Perforce installation. The versioned files for a given depot are stored in a tree of directories beneath this subdirectory.
- *Database files* store *metadata*, including changelists, opened files, client specs, branch specs, and other data concerning the history and present state of the versioned files. Database files appear as `db.*` files in the top level of the server root directory. Each `db.*` file contains a single, binary-encoded database table.

Backup and Recovery Concepts

Disk space shortages, hardware failures, and system crashes can corrupt any of the Perforce server's files. That's why the entire Perforce root directory structure (your versioned files and your database) should be backed up regularly.

As mentioned earlier, versioned files are stored in subdirectories beneath your Perforce server root, and can be restored directly from backups without any loss of integrity.

The files making up the Perforce database, on the other hand, may not have been in a state of transactional integrity at the moment they were copied to the system backups. Restoring the `db.*` files from system backups may result in an inconsistent database. The only way to guarantee the integrity of the database after it's been damaged is to reconstruct the `db.*` files from Perforce checkpoint and journal files.

- A *checkpoint* is a snapshot or copy of the database at a particular moment in time.
- A *journal* is a log that records updates made to the database since the last snapshot was taken.

The checkpoint file is often much smaller than the original database, and can be made smaller still by compressing it. The journal file, on the other hand, can grow quite large; it is truncated whenever a checkpoint is made, and the older journal is renamed. The older journal files can then be backed up offline, freeing up more space locally.

Both the checkpoint and journal are text files, and have the same format. A checkpoint and, if available, its subsequent journal, can restore the Perforce database.

Warning! Checkpoints and journals archive only the Perforce database files, *not* the files in the depot directories! You must always back up the depot files (your versioned files) with the standard OS backup commands after checkpointing.

Because the information stored in the Perforce database is as irreplaceable as your versioned files, checkpointing and journaling are an integral part of administering a Perforce server, and should be performed regularly.

Checkpoint files

A *checkpoint* is a file that contains all information necessary to recreate the metadata in the Perforce database. When you create a checkpoint, the Perforce database is locked, allowing you to take an internally consistent snapshot of that database.

Versioned files are backed up separately from checkpoints. This means that a checkpoint does *not* contain the contents of versioned files, and as such, ***you cannot restore any versioned files from a checkpoint.*** You can, however, restore all changelists, labels, jobs, and so on, from a checkpoint.

To guarantee database integrity upon restoration, the checkpoint must be as old as, or older than, the versioned files in the depot. This means that the database should be checkpointed, and the checkpoint generation must be complete, before the backup of the versioned files starts.

Regular checkpointing is important to keep the journal from getting too long. Making a checkpoint immediately before backing up your system is good practice.

Creating a checkpoint

Checkpoints are not created automatically; someone or something must run the checkpoint command on the Perforce server machine. You can create a checkpoint by invoking the `p4d` program with the `-jc` (journal-create) flag:

```
p4d -r root -jc
```

This can be run while the Perforce server (`p4d`) is running.

To make the checkpoint, `p4d` locks the database and then dumps its contents to a file named `checkpoint.n`, where `n` is a sequence number. Before it unlocks the database, `p4d` also copies the journal file to a file named `journal.n-1`, and then truncates the current journal. This guarantees that the last checkpoint (`checkpoint.n`) combined with the current journal (`journal`) will always reflect the full contents of the database at the time the checkpoint was created.

(The sequence numbers reflect the roll-forward nature of the journal; to restore databases to older checkpoints, match the sequence numbers. That is, the database reflected by `checkpoint.6` can be restored by restoring the database stored in `checkpoint.5` and rolling forward the changes recorded in `journal.5`. In most cases, you're only interested in restoring the current database, which is reflected by the highest-numbered `checkpoint.n` rolled forward with the changes in the current `journal`.)

You can specify a prefix for the checkpoint and journal filename by using the `-jc` option. That is, if you create a checkpoint with:

```
p4d -jc prefix
```

your checkpoint and journal files will be named `prefix.ckp.n`, or `prefix.jnl.n` respectively, where *prefix* is as specified on the command line and *n* is a sequence number. If no *prefix* is specified, the default filenames `checkpoint.n` and `journal.n` will be used.

Note	The meaning of the argument to <code>-jc</code> changed in Release 99.2. Prior to Release 99.2, the files created with <code>p4d -jc prefix</code> would have been <code>prefix.n</code> (for the checkpoint) and <code>journal.n</code> (for the old journal). The behavior in 99.2 is a change from that in previous releases; if you have scripts which rely on the old behavior, you may have to modify them.
-------------	---

As of Release 99.2, if you need to take a checkpoint but are not on the machine running the Perforce server, you can create a checkpoint remotely with the `p4 admin` command. Use:

```
p4 admin checkpoint [prefix]
```

to take the checkpoint and optionally specify a *prefix* to the checkpoint and journal files. (You must be a Perforce superuser to use `p4 admin`.)

A checkpoint file may be compressed, archived, or moved onto another disk. At that time or shortly thereafter, the files in the depot subdirectories should be archived as well.

When recovering, *the checkpoint must be at least as old as the files in the depots*. (that is, the versioned files can be newer than the checkpoint, but not the other way around.) As you might expect, the shorter this time gap, the better.

You can set up an automated program to create your checkpoints on a regular schedule. Be sure to always check the program's output to ensure that the checkpoint creation was successful. The first time you need a checkpoint is not a good time to discover your checkpoint program wasn't working.

If the checkpoint command itself fails, contact Perforce Technical Support immediately. Checkpoint failure is usually a symptom of a resource problem (disk space, permissions, etc.) that can put your database at risk if not handled correctly.

Journal files

The *journal* is the running transaction log that keeps track of all database modifications since the last checkpoint. It's the bridge between two checkpoints.

If you have Monday's checkpoint and the journal that was collected from then until Wednesday, those two files (Monday's checkpoint plus the accumulated journal) contain the same information as a checkpoint made Wednesday. If a disk crash were to cause corruption in your Perforce database on Wednesday at noon, for instance, you could still restore the database even though Wednesday's checkpoint hadn't yet been made.

Warning! By default, the current journal file name is `journal` and it resides in the `P4ROOT` directory. However, if a disk failure corrupts that root directory, your journal file will be inaccessible too.

We strongly recommend that you set up your system so that the journal is written to a filesystem other than the `P4ROOT` filesystem. You can specify this from the command line, or set `P4JOURNAL` before starting the Perforce server to tell it where to write the journal.

To restore your database, you only need to keep the most recent journal file accessible, but it doesn't hurt to archive old journals with old checkpoints, should you ever need to restore to an older checkpoint.

Enabling journaling on Windows

For Windows installations, if you used the installer (`perforce.exe`) to install a Perforce server or service, journaling is turned on for you.

If you installed Perforce without the installer (for an example of when you might do this, see "Multiple Perforce services under Windows" on page 107), you do not have to create an empty file named `journal` in order to enable journaling under a manual installation on Windows.

Enabling journaling on UNIX

For UNIX installations, journaling is also automatically enabled.

If `P4JOURNAL` is left unset (and no location is specified on the command line), the default location for the journal is `$P4ROOT/journal`.

After enabling journaling

Be sure to create a new checkpoint with `p4d -jc` (and `-J journalfile` if required) immediately after enabling journaling. Once journaling is enabled, you'll need make regular checkpoints to control the size of the journal file. An extremely large current journal is a sign that a checkpoint is needed.

Every checkpoint after your first checkpoint starts a new journal file and renames the old one. The old journal is renamed to `journal.n`, (or `prefix.jnl.n` for Release 99.2 or later) where `n` is a sequence number, and a new journal file is created.

By default, the journal is written to the file `journal` in the server root directory (`P4ROOT`). Since there is no sure protection against disk crashes, the journal file and the Perforce server root should be located on different filesystems, ideally on different physical disk drives. The name and location of the journal can be changed by specifying the name of the journal file in the environment variable `P4JOURNAL`, or by providing the `-J filename` flag to `p4d`.

Warning! If you create a journal file with the `-J filename` flag, make sure that subsequent checkpoints use the same file, or the journal will not be properly renamed.

Whether you use `P4JOURNAL` or the `-J journalfile` option to `p4d`, the journal file name can be provided either as an absolute path, or as a path relative to the server root.

Example: Specifying journal files

Starting the server with:

```
$ p4d -r $P4ROOT -p 1666 -J /usr/local/perforce/journalfile
Perforce Server starting...
```

requires that you either checkpoint with:

```
$ p4d -r $P4ROOT -J /usr/local/perforce/journalfile -jc
Checkpointing to checkpoint.19...
Saving journal to journal.18...
Truncating /usr/local/perforce/journalfile...
```

or set `P4JOURNAL` to `/usr/local/perforce/journal` and use

```
$ p4d -r $P4ROOT -jc
Checkpointing to checkpoint.19...
Saving journal to journal.18...
Truncating /usr/local/perforce/journalfile...
```

If your `P4JOURNAL` environment variable (or command-line specification) doesn't match the setting used when you started the Perforce server, the checkpoint is still created, but the journal is neither saved nor truncated. This is highly undesirable!

Disabling journaling

To disable journaling, stop the server, remove the existing journal file (if it exists), set the environment variable `P4JOURNAL` to `off`, and restart `p4d` without the `-J` flag.

Versioned files

Your checkpoint and journal files are used to reconstruct the Perforce database files only. Your versioned files are stored in directories under the Perforce server root, and must be backed up separately.

Versioned file formats

Versioned files are stored in subdirectories beneath your server root. Text files are stored in RCS format, with filenames of the form *filename,v*. There is generally one RCS-format (*,v*) file per text file. Binary files are stored in full in their own directories named *filename,d*. Depending on the Perforce file type selected by the user storing the file, there may be one or more archived binary files in each *filename,d* directory. If more than one file resides in a *filename,d* directory, each one refers to a different revision of the binary file, and is named *1.n*, where *n* is the revision number.

As of Release 99.2, Perforce also supports the AppleSingle file format for Macintosh. On the server, these files are stored in full, compressed, just like other binary files. They are stored in the Mac's AppleSingle file format; if need be, these files can be copied directly from the server root, uncompressed, and used as-is on a Macintosh.

Because Perforce uses compression in the depot files, a system administrator should not rely on the compressibility of the data when sizing backup media. Both text and binary files are either compressed by the Perforce server (denoted by the *.gz* suffix) before storage, or are stored uncompressed. At most installations, if any binary files in the depot subdirectories are being stored uncompressed, they were probably incompressible to begin with. (For example, many image, music, and video file formats are incompressible.)

Back up after checkpointing

In order to ensure that the versioned files reflect all the information in the database after a post-crash restoration, the *db.** files must be restored from a checkpoint that is at least as old as (or older than) your versioned files. For this reason, you should create the checkpoint before backing up the versioned files in the depot directory or directories.

While your versioned files can be newer than the data stored in your checkpoint, it is in your best interest to keep this difference to a minimum; in general, you'll want your backup script to back up your versioned files immediately after successfully completing a checkpoint.

Backup Procedures

To back up your Perforce server, perform the following steps as part of your nightly backup procedure:

1. Verify the integrity of your server and add file signatures to any new files:

```
p4 verify //...  
p4 verify -u //...
```

You may wish to pass the `-q` (quiet) option to `p4 verify`. If called with the `-q` option, `p4 verify` will produce output only when errors are detected.

The first command (`p4 verify`) recomputes the MD5 signatures of all of your archived files and compares them with those stored when `p4 verify -u` was first run on them. It also ensures that all files known to Perforce actually exist in the depot subdirectories; a disk-full condition that results in corruption of the database or archived files during the day can be detected by examining the output of these commands.

The second command (`p4 verify -u`) updates the database with MD5 signatures for any new file revisions for which checksums have not yet been computed.

By running `p4 verify -u` before the backup, you ensure that you create and store checksums for any files new to the depot since your last backup, and that these checksums are stored as part of the backup you're about to take.

The use of `p4 verify` is optional, but is good practice not only because it allows you to spot any server corruption before a backup is made, but it also gives you the ability, following a crash, to detect whether or not the files restored from your backups are in good condition.

Note | If your site is very large, `p4 verify` may take some time to run, and you may wish to perform this step on a weekly basis rather than on a daily basis. For more about the `p4 verify` command, see “File verification by signature” on page 43.

2. Make a checkpoint by invoking `p4d` with the `-jc` (journal-create) flag, or by using the `p4 admin` command. Use one of:

```
p4d -jc
```

or (as of Release 99.2 or higher):

```
p4 admin checkpoint
```

Because `p4d` locks the entire database when making the checkpoint, you do not generally have to stop your Perforce server during any part of the backup procedure.

Note | If your site is very large (say, several GB of `db.*` files), creating a checkpoint may take a considerable length of time.

Under such circumstances, you may wish to defer checkpoint creation and journal truncation until times of low system activity. You might, for instance, archive only the `journal` file in your nightly backup, and only create checkpoints and roll the journal file on a weekly basis.

If you are using the `-z` flag to create a gzip-compressed checkpoint, the checkpoint file will be named as specified. If you want the compressed checkpoint file to end in `.gz`, you should explicitly specify the `.gz` on the command line.

3. Ensure that the checkpoint has been created successfully before backing up any files. (After a disk crash, the last thing you want to discover is that the checkpoints you've been backing up for the past three weeks were incomplete!)

You can tell that the checkpoint command has completed successfully by examining the error code returned from `p4d -jc`, or by observing the truncation of the current journal file.

4. Once the checkpoint has been created successfully, back up the checkpoint file, the old journal file, and your versioned files.

(If you don't require an audit trail, you don't actually need to back up the journal. It is, however, usually good practice to do so.)

Note | There are rare instances (for instance, users obliterating files during backup, or submitting files on Windows during the file backup portion of the process) in which your depot files may change during the interval between the time the checkpoint was taken and the time at which the depot files get backed up by the backup utility.

Most sites are affected by these issues; having the Perforce server available on a 24/7 basis is generally a benefit worth this minor risk, especially if backups are being performed at times of low system activity.

If, however, the reliability of every backup is of paramount importance, consider stopping the Perforce server before checkpointing, and restarting it after the backup process has completed. Doing so will eliminate all risk of the system state changing during the backup process.

You never need to back up the `db.*` files. Your latest checkpoint and journal contain all the information necessary to re-create them. More significantly, a database

restored from `db.*` files is not guaranteed to be in a state of transactional integrity. A database restored from a checkpoint is.

Windows | On Windows, if you make your system backup while the Perforce server is running, you must ensure that your backup program doesn't attempt to back up the `db.*` files.

If you try to back up the `db.*` files with a running server, Windows locks them while the backup program backs them up. During this brief period, the Perforce server is unable to access the files; if a user attempts to perform an operation that would update the file, the server may fail.

If your backup software doesn't allow you to exclude the `db.*` files from the backup process, you should stop the server with `p4 admin stop` before backing up, and restart the server after the backup process is complete.

Recovery Procedures

If the database files become corrupted or lost, either because of disk errors, a hardware failure such as a disk crash, the database can be recreated with your stored checkpoint and journal.

There are many ways in which systems can fail; while this guide cannot address all of them, it can at least provide a general guideline for recovery from the two most common situations, specifically:

- corruption of your Perforce database only, without damage to your versioned files, and
- corruption to both your database and versioned files.

The recovery procedures for each failure are slightly different, and are discussed separately in the following two sections.

If you suspect corruption in either your database or versioned files, contact Perforce technical support.

Database corruption, versioned files unaffected

If only your database has been corrupted, (that is, your `db.*` files were on a drive that crashed, but you were using symbolic links to store your versioned files on a separate physical drive), you need only re-create your database.

You *will* need:

- The last checkpoint file, which should be available from the latest `P4ROOT` directory backup.

- The current journal file, which should be on a separate filesystem from your `P4ROOT` directory, and which should therefore have been unaffected by any damage to the filesystem where your `P4ROOT` directory was held.

You will *not* need:

- Your backup of your versioned files; if they weren't affected by the crash, they're already up to date.

To recover the database

1. Stop the current instance of `p4d`:

```
p4 admin stop
```

(You must be a Perforce superuser to use `p4 admin`.)

2. Rename (or move) the corrupt database ("`db.*`") files:

```
mv your_root_dir/db.* /tmp
```

The corrupt `db.*` files aren't actually used in the restoration process, but it's safe practice not to delete them until you're certain your restoration was successful.

3. Invoke `p4d` with the `-jr` (journal-restore) flag, specifying your most recent checkpoint and current journal. If you explicitly specify the server root (`$P4ROOT`), the `-r $P4ROOT` argument must precede the `-jr` flag:

```
p4d -r $P4ROOT -jr checkpoint_file journal_file
```

This recovers the database as it existed when the last checkpoint was taken, and then apply the changes recorded in the journal file since the checkpoint was taken.

Note If you're using the `-z` (compress) option to compress your checkpoints upon creation, you'll have to restore the uncompressed journal file separately from the compressed checkpoint.

That is, instead of using:

```
p4d -r $P4ROOT -jr checkpoint_file journal_file
```

you'll use two commands:

```
p4d -r $P4ROOT -z -jr checkpoint_file.gz
```

```
p4d -r $P4ROOT -jr journal_file
```

You must explicitly specify the `.gz` extension yourself when using the `-z` flag, and ensure that the `-r $P4ROOT` argument precedes the `-jr` flag.

Check your system

Your restoration is complete. See "Ensuring system integrity after any restoration" on page 37 to make sure your restoration was successful.

Your system state

The database recovered from your most recent checkpoint, after you've applied the accumulated changes stored in the current journal file, is up to date as of the time of failure.

After recovery, both your database and versioned files should reflect all changes made up to the time of the crash, and no data should have been lost.

Both database and versioned files lost or damaged

If both your database and your versioned files were corrupted, you need to restore both the database and your versioned files, and you'll need to ensure that the versioned files are no older than the restored database.

You *will* need:

- The last checkpoint file, which should be available from the latest P4ROOT directory backup.
- Your versioned files, which should be available from the latest P4ROOT directory backup.

You will *not* need:

- Your current journal file. The journal contains a record of changes to the metadata and versioned files that occurred between the last backup and the crash; because you'll be restoring a set of versioned files from a backup taken *before* that crash, the checkpoint alone contains the metadata useful for the recovery, and the information in the journal is of limited or no use.

To recover the database

1. Stop the current instance of p4d:

```
p4 admin stop
```

(You must be a Perforce superuser to use `p4 admin`.)

2. Rename (or move) the corrupt database ("db.") files:

```
mv your_root_dir/db.* /tmp
```

The corrupt `db.*` files aren't actually used in the restoration process, but it's safe practice not to delete them until you're certain your restoration was successful.

3. Invoke `p4d` with the `-jr` (journal-restore) flag, specifying *only* your most recent checkpoint:

```
p4d -r $P4ROOT -jr checkpoint_file
```

This recovers the database as it existed when the last checkpoint was taken, but does not apply any of the changes in the journal file. (The `-r $P4ROOT` argument must precede the `-jr` flag.)

The database recovery without the roll-forward of changes in the journal file brings the database up to date as of the time of your last backup. In this scenario, you do not want to apply the changes in the journal file, because the versioned files you restored reflect only the depot as it existed as of the last checkpoint.

To recover your versioned files

4. After recovering the database, you then need to restore the versioned files according to your system's restoration procedures (for instance, the UNIX `restore(1)` command) to ensure that they are as new as the database.

Check your system

Your restoration is complete. See “Ensuring system integrity after any restoration” on page 37 to make sure your restoration was successful.

Note that files submitted to the depot between the time of the last system backup and the disk crash will not be present in the depot.

Note Although “new” files (submitted to the depot but not yet backed up) will not appear in the depot after restoration, it's possible (indeed, highly probable!) that at one or more of your users will have up-to-date copies of such files present in their client workspaces.

Your users can find such files by using Perforce to examine how files in their client workspaces differ from those in the depot. If they run:

```
p4 diff -se
```

...they'll be provided with a list of files in their workspace which differ from the files Perforce believes them to have. After verifying that these files are indeed the files you wish to restore, you may wish to have one of your users open these files for `edit` and submit them to the depot in a changelist.

Your system state

After recovery, your depot directories may not contain the newest versioned files. That is, files submitted after the last system backup but before the disk crash may have been lost.

- In most cases, the latest revisions of such files can be restored from the copies still residing in your users' client workspaces.
- In a case where *only* your versioned files (but *not* the database, which may have resided on a separate disk and remained unaffected by the crash) were lost, you may also be able to make a separate copy of your database and apply your journal to it in order to

examine recent changelists to track down files submitted between the last backup and the disk crash.

In either case, contact Perforce technical support for further assistance.

Ensuring system integrity after any restoration

After any restoration, it's wise to run `p4 verify` to ensure the versioned files are at least as new as the database:

```
p4 verify -q ///
```

This command verifies the integrity of the versioned files. The `-q` (quiet) option tells the command to only produce output on error conditions. Ideally, this command should produce no output.

If any versioned files are reported as `MISSING` by the `p4 verify` command, you'll know that there is information in the database concerning files that didn't get restored. The usual cause is that you restored from a checkpoint and journal made after the backup of your versioned files. (that is, that your backup of the versioned files was older than the database.)

If (as recommended) you've been using `p4 verify -u` to generate and store MD5 signatures for your versioned files as part of your backup routine, you can run `p4 verify` on the server after restoration to reassure yourself that your restoration was successful.

If you have any difficulties restoring your system after a crash, contact Perforce Technical Support for assistance.

Administering Perforce: Superuser Tasks

This chapter describes both basic tasks associated with day-to-day Perforce administration as well as setting up advanced Perforce configurations, dealing with cross-platform development issues, migrating Perforce servers from one machine to another, and setting up remote and local depots.

Each of the tasks described in this chapter requires that you be a Perforce superuser (access level `super`) or administrator (access level `admin`) as defined in the Perforce protections table. For more about controlling Perforce superuser access, and protections in general, see “Administering Perforce: Protections” on page 61.

Basic Perforce Administration

The following tasks commonly performed by Perforce administrators and superusers are:

- User operations, including resetting passwords, creating new users, disabling the automatic creation of new users, and cleaning up files left open by former users,
- Administrative operations, including obliterating files to reclaim disk space, editing previously submitted changelists, verifying server integrity, defining filetypes to fine-tune Perforce’s file type detection mechanism, and the use of the `-f` flag to force operations.

Resetting user passwords

From time to time, it is inevitable that users forget Perforce passwords. A Perforce superuser can set a new password for any user with:

- Release 99.1 and later:

```
p4 passwd username
```

(Perforce prompts the superuser for a new password for user *username*.)

- Pre-99.1 releases:

```
p4 user -f username
```

(The user specification form for *username* appears, and the new password may be entered.)

Creating new users

By default, Perforce creates a new user in its database whenever a command is issued by a username it hasn’t seen before.

Superusers can use the `-f` (force) flag to create a new user as follows:

```
p4 user -f username
```

Fill in the form fields with the information for the user you wish to create.

The `p4 user` command also has an option (`-i`) to take its input from the standard input instead of the forms editor. If you wish to create a large number of new users at once, you can write a script which creates output in the same format as that used by the forms editor and then pipes each pre-generated “form” to `p4 users -i -f username`, where the `username` is also specified by a variable within your calling script.

Preventing creation of new users

As mentioned, Perforce’s default behavior is to create a new user in its database whenever a command is issued by a username it hasn’t seen before.

To prevent Perforce from automatically creating new users, all users must be explicitly listed in the protections table. The easiest way to ensure that this is the case is to put all users into a Perforce group, and to configure Perforce to only permit access to members of that group.

Example: *Setting up users in a group.*

A Perforce superuser wants to prevent the server from creating new users. He starts by setting up a group called `p4users` for the three users currently at his site. He types:

```
p4 group p4users
```

and fills in the form as follows:

```
# A Perforce Group Specification.
#
# Group:      The name of the group.
# MaxResults: A limit on the data size of operations for users in
#             this group, or 'unlimited'.
# Subgroups:  Other groups automatically included in this group.
# Users:      The users in the group. One per line.
Group:  p4users
MaxResults:  unlimited
MaxScanRows:  unlimited
Subgroups:
Users:
    edk
    lisag
    sarahm
```

He then uses `p4 protect` to edit the protections table. The relevant line of the default protections table looks like this:

```
write user * * //...
```


*This grants write permission to any user matching * (that is, to all users) from any host (the second *) in all areas of the depot (that is, to files in //...).*

After using `p4 group p4users` to create the Perforce group `p4users`, he uses `p4 protect` to change this line in the protections table to read:

```
write group p4users * //...
```

The replacement protection grants only write access to users whose group matches `p4users`. Members of `p4users` may use Perforce from any host () and have write access to all areas of the depot (//...).*

As long as no other lines in the protections table grant permission to all users, all users are now defined within `p4 protect`, and the server will no longer automatically create new user entries when new users attempt to access Perforce.

For a more in-depth description of Perforce protections, see “Administering Perforce: Protections” on page 61.

Deleting obsolete users

Each user on the system consumes one Perforce license. A Perforce administrator can free up licenses from unused users by deleting them.

```
p4 user -d username
```

You must first revert (or submit) any open files opened by a user before deleting that user. If you attempt to delete a user who has opened files, Perforce will display an error message to that effect.

Reverting files left open by obsolete users

If files have been left open by a nonexistent or obsolete user (for instance, a departing employee), a Perforce administrator can revert the files by deleting the client spec in which they were opened.

For example, if the output of `p4 opened` shows:

```
//depot/main/code/file.c#8 - edit default change (txt) by jim@stlouis
```

the “stlouis” client spec can be deleted with:

```
p4 client -d -f stlouis
```

Deleting a user’s client spec automatically reverts all files opened by that client, and also removes that client’s “have list”. Note that it does *not* affect any files in the workspace actually used by that client; the files can still be accessed by other employees.

Reclaiming disk space by obliterating files

Warning! Use `p4 obliterate` with caution. This is the only command in Perforce that actually removes file data.

The depot is always growing, and this is not always desirable: a submit might have been performed incorrectly, creating hundreds of unneeded files, or perhaps there are simply a lot of old files around that are no longer being used. Because `p4 delete` merely marks files as deleted in their head revisions, it can't be used to free up disk space on the server. This is where `p4 obliterate` can be useful.

Perforce administrators can use `p4 obliterate filename` to remove all traces of a file from a depot, making the file indistinguishable from one that never existed in the first place.

Note | The purpose of a software configuration management system is to allow your site to maintain a history of which operations were performed on which files.

The `p4 obliterate` command defeats this purpose; as such, it is only intended to be used when cleaning up messes in the depot, and not as part of your normal software development process.

By default, `p4 obliterate filename` does nothing; it merely reports on what it would do. To actually destroy the files, use `p4 obliterate -y filename`.

If you need to destroy only one revision of a file (perhaps someone inadvertently stored some line art as a 20-megabyte uncompressed TIFF in place of its 500K-long compressed equivalent), specify only the desired revision number on the command line. For instance, to destroy revision #5 of a file, use:

```
p4 obliterate -y file#5
```

Revision ranges are also acceptable: To destroy revisions 5 through 7 of a file:

```
p4 obliterate -y file#5,7
```

Warning! If you mean to obliterate a revision range, be certain you've specified it properly. If you omit the specify revision range, **all** revisions of the file will be obliterated!

The safest way to use `p4 obliterate` is to use it **without** the `-y` (confirmation) flag until you're certain you've specified the files and revisions correctly.

The `p4 obliterate` command has one more flag: `-z`. When you branch a file from one area of the depot into another, a "lazy copy" is created - the file itself isn't copied, only a

record that the branch has occurred. If, for some reason, you wish to undo the “lazy copy” and create a new copy of the branched file’s contents in your depot subdirectories, you could “obliterate” the lazy copy and create a new one by using `p4 obliterate -y -z filename`.

Removing lazy copies by using the `-z` flag typically *increases* disk space usage. The only practical use of `p4 obliterate -y -z` is to undo lazy copies in order to allow you to manually remove archive files without breaking any linked metadata pointing to the deleted files.

If a user sees the following error message while trying to access files:

```
Operation:user-sync
Librarian checkout path failed
```

where *path* is the path of a previously-obliterated file, the user has probably encountered a problem that resulted from an earlier use of `p4 obliterate` from an older (pre-98.2/10314) Perforce server. Contact Perforce technical support for a workaround.

Deleting changelists and editing changelist descriptions

Perforce administrators can use the `-f` (force) flag with `p4 change` to change the description or username of a submitted changelist. The syntax is `p4 change -f changenumber`. This presents the standard changelist form, but allows you to edit the change time, description, and/or username.

You can also use the `-f` flag to delete any submitted changelists that have been emptied of files with `p4 obliterate`. The full syntax is `p4 change -d -f changenumber`.

Example: *Updating changelist 123 and deleting changelist 124*

Use `p4 change` with the `-f` (force) flag:

```
p4 change -f 123
p4 change -d -f 124
```

The User: and Description: fields for change 123 are edited, and change 124 is deleted.

File verification by signature

Perforce administrators can use the `p4 verify filenames` command to generate 128-bit MD5 signatures of each revision of the named files. The signatures created by `p4 verify -u` can later be used to confirm proper recovery in case of a crash: if the signatures of the recovered files match the previously saved signatures, the files were recovered accurately.

To generate signatures and store them in the Perforce database, use the `-u` flag. Subsequent verifications will be compared against the stored signatures; if a new signature does not match the signature in the Perforce database for that file revision, Perforce adds the characters `BAD!` after the signature.

If you ever see a `BAD!` signature during a `p4 verify` command, your database or versioned files may have been corrupted, and you should contact Perforce Technical Support.

Because subsequent verifications can only be performed against previously stored signatures, the `p4 verify -u` command should be used regularly.

A good strategy, for instance, might be to run `p4 verify` on a nightly basis before performing your system backups, proceeding with the backup only if the `p4 verify` reports no corruption. Generation and storage of new checksums (`p4 verify -u`) following a successful `p4 verify` could be performed nightly, or even weekly.

Verifying during server upgrades

It is also good practice to use `p4 verify` during server upgrades:

1. Before the upgrade, run:

```
p4 verify -qu //...
```

to generate the new checksums.
2. Take a checkpoint and copy the checkpoint and your versioned files to a safe place.
3. Perform the server upgrade.
4. After the upgrade, run:

```
p4 verify -q //...
```

to verify the integrity of your system.

Defining filetypes with `p4 typemap`

By default, Perforce automatically determines if a file is of type `text` or `binary` based on an analysis of the first 1024 bytes of a file. If the high bit is clear in each of the first 1024 bytes, Perforce assumes it to be `text`; otherwise, it is assumed to be `binary`. Although this default behavior can be overridden by the use of the `-t filetype` flag, it's easy for users to overlook this, particularly in cases where files' types are usually (but not always) detected correctly.

The `p4 typemap` command solves this problem by allowing system administrators to set up a table that links Perforce file types with file name specifications. If an entry in the `typemap` table matches an entry in the table, it overrides the file type that would otherwise be assigned by the Perforce client.

One common use for `p4 typemap` is for users dealing with Adobe PDF (Portable Document Format) files. Some PDF files start with a series of comment fields and textual data, and if the comments are sufficiently long, the files will be erroneously detected by Perforce as being of type `text`. Similarly, files in RTF (Rich Text Format) format may sometimes be erroneously detected as `text`.

Perforce administrators may use `p4 typemap` to tell the Perforce server to regard all such files as `binary` by modifying the `typemap` table as follows:

```
Typemap:
    binary //....pdf
    binary //....rtf
```

The first three periods (“...”) in the specification are a Perforce wildcard specifying that all files beneath the root directory are to be included as part of the mapping. The fourth period and the file extension specify that the specification applies to files ending in “.pdf” (or “.rtf”).

For more information, see the `p4 typemap` page in the *Perforce Command Reference*.

Forcing operations with the -f flag

Certain commands allow Perforce administrators and superusers to use the `-f` flag to force certain operations unavailable to ordinary users.

Perforce administrators can use this flag with `p4 branch`, `p4 change`, `p4 client`, `p4 job`, `p4 label`, and `p4 unlock`. Perforce superusers can also use it to override the `p4 user` command.

The usages and meanings of this flag are as follows:

Command	Syntax	Function
p4 branch	p4 branch -f <i>branchname</i>	Allows the modification date to be changed while editing the branch specification
	p4 branch -f -d <i>branchname</i>	Deletes the branch, ignoring ownership
p4 change	p4 change -f [<i>changelist#</i>]	Allows the modification date to be changed while editing the changelist specification
	p4 change -f <i>changelist#</i>	Allows the description field and username in a committed changelist to be edited
	p4 change -f -d <i>changelist#</i>	Deletes empty, committed changelists
p4 client	p4 client -f <i>clientname</i>	Allows the modification date to be changed while editing the client specification


Command	Syntax	Function
	<code>p4 client -f -d <i>clientname</i></code>	Deletes the client, ignoring ownership, even if the client has opened files
<code>p4 job</code>	<code>p4 job -f [<i>jobname</i>]</code>	Allows the manual update of read-only fields
<code>p4 label</code>	<code>p4 label -f <i>labelname</i></code>	Allows the modification date to be changed while editing the label specification
	<code>p4 label -f -d <i>labelname</i></code>	Deletes the label, ignoring ownership
<code>p4 unlock</code>	<code>p4 unlock -c <i>changelist</i> -f <i>file</i></code>	Releases a lock (set with <code>p4 lock</code>) on an open file in a pending numbered changelist, ignoring ownership.
<code>p4 user</code>	<code>p4 user -f <i>username</i></code>	Allows the update of all fields, ignoring ownership. This command requires <code>super</code> access.
	<code>p4 user -f -d <i>username</i></code>	Deletes the user, ignoring ownership. This command requires <code>super</code> access.

Advanced Perforce Administration

Running Perforce through a firewall


Perforce clients communicate with a Perforce server using TCP/IP. The server listens for connections at a specified port on the machine on which it's running, and clients make connections to that port.

The port on which the server listens is specified when the server is started. The number is arbitrary, so long as it does not conflict with any other networking services and is greater than 1024. The port number on the client machine is dynamically allocated.



A *firewall* is a network element which prevents any packets from outside a local (trusted) network from reaching that local network. This is done at a low level in the network protocol; any packets not coming from a trusted IP address are simply ignored.

In the following diagram, the Perforce client is on an untrusted part of the network. None of its connection requests reach the machine with the Perforce server. Consequently, the user running the client through the firewall is unable to use Perforce.



Secure shell

To solve this problem, you have to make the connection to the Perforce server from within the trusted network. This can be done securely using a package called *secure shell* (`ssh`).

Secure shell (`ssh`) is meant to be a replacement for the UNIX `rsh` (remote shell) command, which allows you to log into a remote system and execute commands on it. The “secure” part of “secure shell” comes from the fact that the connection is encrypted, so none of the data is visible while it passes through the untrusted network. With simple utilities like `rsh`, all traffic - even passwords - is unencrypted and visible to all intermediate hosts, creating an unacceptable security hazard.

Secure shell is available for free in source form for a multitude of UNIX platforms from <http://www.openssh.com>. This page also links to ports of `ssh` for OS/2 and Amiga, as well as commercial implementations for Windows and Macintosh from Data Fellows (<http://www.datafellows.com>) and SSH (<http://www.ssh.com>).


The OpenSSH FAQ (Frequently Asked Questions) can also be found online at the main site (<http://www.openssh.com/faq.html>).

Solving the problem

Once you have `ssh` up and running, the simplest thing to do is to use it to log into the firewall machine and run the Perforce client from the firewall. While it has the advantage of simplicity, it's a poor solution: you typically want your client files accessible on your local machine, and of course, there's no guarantee that your firewall machine will match your development platform.

A good solution takes advantage of `ssh`'s ability to *forward* arbitrary TCP/IP connections. By using `ssh`, you can make your Perforce client appear as though it's connecting from the firewall machine over the local (trusted) network. In reality, your client remains on your local machine, but all packets from your local machine are first sent to the firewall through the secure channel set up by `ssh`.

Suppose the Perforce server is on `p4dbox.bigcorp.com`, and the firewall machine is called `firewall.bigcorp.com`. In our example, we'll arbitrarily choose local port 4242, and assume that the Perforce server is listening on port 3710.



Packets ultimately destined for your client's port 4242 are first sent to the firewall, and `ssh` forwards them securely to your client. Likewise, connections made to port 4242 of the firewall machine will end up being routed to port 3710 of the Perforce server.

On UNIX, the `ssh` command on your own machine to set up and forward the TCP/IP connection would be:

```
ssh -L 4242:p4dbox.bigcorp.com:3710 firewall.bigcorp.com
```

At this point, it may be necessary to provide a password to log into `firewall.bigcorp.com`. Once the connection is established, `ssh` listens at port 4242 on the local machine, and forwards packets over its encrypted connection to `firewall.bigcorp.com`; the firewall then forwards them by normal channels to port 3710 on `p4dbox.bigcorp.com`.

All that remains is to tell the Perforce client to use port 4242 by setting the environment variable `P4PORT` to 4242.

Normally, setting `P4PORT=4242` would normally indicate that we are trying to connect to a Perforce server on the local machine listening at port 4242. In this case, `ssh` takes the role of the Perforce server. Anything a client sends to port 4242 of the local machine is forwarded by `ssh` to the firewall, which passes it to the real Perforce server at `p4dbbox.bigcorp.com`. Since all of this is transparent to the Perforce client, it doesn't matter whether the client is talking to an instance of `ssh` that's forwarding traffic from port 4242 of the local machine, or if it's talking to a real Perforce server residing on the local machine.

The only glitch is that there's a login session you don't normally want on the firewall machine.

This can be solved by running

```
ssh -L 4242:p4dbbox.bigcorp.com:3710 firewall.bigcorp.com -f sleep 9999999 -f
```

on the remote system.

This tells `ssh` on `firewall.bigcorp.com` to fork a long-running `sleep` command in the background after the password prompt. Effectively, this sets up the `ssh` link and keeps it up; there is no login session to terminate.

Finally, `ssh` can be configured to “do the right thing” so that it is unnecessary to type such a long command with each session. The Windows version of `ssh`, for instance, has a GUI to configure this.

One final concern: with port 4242 on the local machine now forwarded to a supposedly secure server, your local machine is part of the trusted network; it is prudent to make sure the local machine really *is* secure. The Windows version of `ssh` has an option to *only* allow local connections to the forwarded port, which is a wise precaution; your machine will be able to use port 4242, but a third party's machine will be ignored.

Specifying IP addresses in P4PORT

Under most circumstances, your Perforce server's `P4PORT` setting consists solely of a port number.

If, however, you specify both an IP address *and* a port number in `P4PORT` when starting `p4d`, the Perforce server takes the IP address into account, and ignores requests from any IP addresses other than the one specified in `P4PORT`.

Although this isn't the default behavior, it can be useful. For instance, if you want to tell `p4d` to listen only to a specific network interface or IP address, you can make your Perforce server ignore all non-local connection requests by setting `P4PORT=localhost:port`.

Running from inetd on UNIX

Under a normal installation, the Perforce server is run on UNIX as a background process which waits for connections from clients. It is possible, however, to have `p4d` start up only when connections are made to it, using `inetd` and `p4d -i`.

If you wish to do this, add the following line to `/etc/inetd.conf`:

```
p4dservice stream tcp nowait username /usr/local/bin/p4d p4d -i -rp4droot
```

and add the following to `/etc/services`:

```
p4dservice nnnn/tcp
```

where:

- `p4dservice` is the service name you choose for this Perforce server
- `/usr/local/bin` is the directory holding your `p4d` binary
- `p4droot` is the root directory (`P4DROOT`) to use for this Perforce server (for example, `/usr/local/p4d`)
- `username` is the UNIX user name to use for running this Perforce server
- `nnnn` is the port number for this Perforce server to use

Note the “extra” `p4d` on the `/etc/inetd.conf` line must be there; `inetd` passes this to the OS as `argv[0]`. The first argument, then, is the `-i` flag, which causes `p4d` not to run in the background as a daemon, but rather to serve the single client connected to it on `stdin/stdout`. (This is the convention used for services started by `inetd`.)

This method is an alternative to running `p4d` from a startup script. It can also be useful for providing special services; for example, at Perforce, we have a number of test servers running on UNIX, each defined as an `inetd` service with its own port number.

There are caveats with this method:

- `inetd` may disallow excessive connections, so a script which invokes several thousand `p4` commands, each of which spawns an instance of `p4d` via `inetd` may cause `inetd` to temporarily disable the service. Depending on your system, you may need to configure `inetd` to ignore or raise this limit.
- There is no easy way to disable the server, since the `p4d` executable is run each time; disabling the server requires modifying `/etc/inetd.conf` and restarting `inetd`.

Case sensitivity and multi-platform development

Early (pre-97.2) releases of the Perforce server treated all filenames, pathnames, and database entity names with case significance, whether the server was running on UNIX or Windows.

For example, `//depot/main/foo.c` and `//depot/MAIN/FOO.C` were treated as two completely different files. This caused problems where users on UNIX were connecting to a Perforce server running on Windows, because the filesystem underlying the server could not store files with the case-variant names submitted by UNIX users.

If you are running a pre-97.2 server on Windows, please contact support@perforce.com to discuss upgrading your server and database. In release 97.3, the behavior was changed, and only the UNIX server supports case-sensitive names. However, there are still some case-sensitivity problems which users can run into when sharing development projects across UNIX and Windows.

To summarize:

- The Perforce server on UNIX supports case-sensitive names.
- The Perforce server on Windows ignores case differences.
- Case is always ignored in keyword-based job searches, regardless of platform

Case-sensitive	UNIX server	Windows server
Pathnames and filenames	Yes	No
Database entities (clients, labels, etc.)	Yes	No
Job search keywords	No	No

To find out what platform your Perforce server runs on, use `p4 info`.

Perforce server on UNIX

If your Perforce server is on UNIX, and you have users on both UNIX and Windows, your UNIX users must be very careful not to submit files whose names differ only by case. Although the UNIX server can support these files, when Windows users sync their workspaces, they'll find files overwriting each other.

Conversely, Windows users will have to be careful to use case consistently in file and path names when adding new files. They may not realize that files added as `//depot/main/foo.c` and `//depot/MAIN/bar.c` will appear in two different directories in a UNIX user's workspace.

The UNIX Perforce server always respects case in client names, label names, branch view names, and so on. Windows users connecting to a UNIX server should be aware that the lowercased workstation names are used as the default names for new client workspaces. For examples, if a new user creates a client spec on a Windows machine named `ROCKET`, his client workspace is named `rocket` by default. If he later sets `P4CLIENT` to `ROCKET` (or `Rocket`), Perforce will tell him his client is undefined. He must set `P4CLIENT` to `rocket` (or unset it) to use the client workspace he defined.

Perforce server on Windows

If your Perforce server is running on Windows, your UNIX users must be aware that their Perforce server will store case-variant files in the same namespace.

For example, users who try something like this:

```
p4 add foo/file1
p4 add foo/file2
p4 add FOO/file3
```

should be aware that all three files will be stored in the same depot directory. The depot path and filenames assigned to the Windows server will be those first referenced. (In this case, the depot path name would be `foo`, and not `FOO`.)

Perforce server trace flags

You can turn on command tracing in the Perforce server by adding the `-v server=1` flag to the `p4d` startup command. Use `P4LOG` or the `-L logfile` flag to name a log file. For example:

```
p4d -r /usr/perforce -v server=1 -p 1666 -L /usr/perforce/logfile
```

Trace output appears in the specified log file, and shows the date, time, username, IP address, and command for each request processed by the server. Before turning on logging, you should make sure that you have adequate disk space.

Windows	<p>Prior to Release 98.1, you could not set this trace flag when running Perforce as a service; you could set this flag (on Windows only) when running <code>p4d.exe</code> a server process from the MS-DOS command line.</p> <p>As of Release 98.1, you can use the <code>p4 set</code> command to set <code>P4DEBUG</code> as a registry variable to “<code>server=1</code>” and thereby use this trace flag with Perforce installed as a service on Windows.</p> <p>Prior to Release 97.3, the server trace flag was unavailable on any server platform.</p>
----------------	--

The server trace flags and their meanings are as follows:

Trace flag	Meaning
<code>server=1</code>	Logs server commands to the server log file. (Requires server at release 98.1 or higher)

Trace flag	Meaning
<code>server=2</code>	In addition to data logged at level 1, logs server command completion. (Requires server at release 2001.1 or higher)
<code>server=3</code>	In addition to data logged at level 2, adds basic information on server CPU time used. (Requires server at release 2001.2 or higher)

In most cases, the Perforce server trace flags are useful only to administrators working with Perforce Technical Support to diagnose or investigate a problem.

Migrating to a new machine

The procedure for moving an existing Perforce installation from one machine to another depends on whether or not you're moving between machines

- of identical architecture,
- of different architectures using the same text file (CR/LF) format, or
- of different architecture *and* different text file format.

There are also additional considerations if the new machine has a different IP address or hostname.

The Perforce server stores two types of data under the Perforce root directory: *versioned files* and a *database* containing *metadata* describing those files. Your versioned files are the ones created and maintained by your users, and your database is a set of Perforce-maintained binary files holding the history and present state of the versioned files. In order to move a Perforce server to a new machine, both the versioned files and the database must be successfully migrated from the old machine to the new machine.

For more about the distinction between versioned files and database, as well as for an overview of backup and restore procedures in general, see “Backup and Recovery Concepts” on page 25.

Moving your versioned files and Perforce database

Between machines of the same architecture

If the architecture of the two machines is the same (e.g., SPARC/SPARC, x86/x86), the versioned files and database can be copied directly between the machines, and you only need to move the server root directory tree to the new machine. You can use `tar`, `cp`, `xcopy.exe`, or any other method. Copy everything in and under the `P4ROOT` directory - the `db.*` files (your database) as well as the depot subdirectories (your versioned files).

1. Back up your server (including a `p4 verify` before the backup) and take a checkpoint.
2. On the old machine, stop `p4d`.
3. Copy the contents of your old server root (`P4ROOT`) and all its subdirectories on the old machine into the new server root directory on the new machine.
4. Start `p4d` on the new machine with the desired flags.
5. Run `p4 verify` on the new machine to ensure that the database and your versioned files were transferred correctly to the new machine.

(Although the backup, checkpoint, and subsequent `p4 verify` are not strictly necessary in this case, it's always good practice to verify, checkpoint, and back up your system before any migration, and likewise to perform a subsequent verification after migration.)

Between different architectures using the same text format

If the internal data representation (big-endian vs. little-endian) convention differs between the two machines (e.g., Linux-on-x86/SPARC, NT-on-Alpha/NT-on-x86), but their operating systems use the same CR/LF text file conventions, you can still simply move the server root directory tree to the new machine.

Although the versioned files are portable across architectures, the database, as stored in the `db.*` files, is not. To transfer the database, you will need to create a checkpoint of your Perforce server on the old machine and use that checkpoint to recreate the database on the new machine. The checkpoint is a text file which can be read by a Perforce server on any architecture. For more details, see “Creating a checkpoint” on page 26.

After creating the checkpoint, you can use `tar`, `cp`, `xcopy.exe`, or any other method to copy the checkpoint file and the depot directories to the new machine. (You don't need to copy the `db.*` files, because they will be recreated from the checkpoint you took.)

1. On the old machine, use `p4 verify` to ensure that the database is in a consistent state.
2. On the old machine, stop `p4d`.
3. On the old machine, create a checkpoint:

```
p4d -jc checkpointfile
```
4. Copy the contents of your old server root (`P4ROOT`) and all its subdirectories on the old machine into the new server root directory on the new machine.

(To be precise, you don't need to copy the `db.*` files, just the checkpoint and the depot subdirectories. The `db.*` files will be recreated from the checkpoint. If it's more convenient to copy everything, then copy everything.)

5. On the new machine, if you copied the `db.*` files, be sure to remove them from the new `P4ROOT` before continuing.
6. Recreate a new set of `db.*` files suitable for your new machine's architecture from the checkpoint you created:

```
p4d -jr checkpointfile
```
7. Start `p4d` on the new machine with the desired flags.
8. Run `p4 verify` on the new machine to ensure that the database and your versioned files were transferred correctly to the new machine.

Between Windows and UNIX

In this case, both the architecture of the system *and* the CR/LF text file convention may be different. You still have to create a checkpoint, copy it, and recreate the database on the new platform, but when you move the depot subdirectories containing your versioned files, you will also have to address the issue of the differing linefeed convention between the two platforms.

Depot subdirectories can contain both text and binary files. The text files (in RCS format, ending with “,v”) and binary files (directories of individual binary files, each directory ending with “,d”) will need to be transferred differently in order to translate the line endings on the text files while leaving the binary files unchanged.

As with all other migrations, be sure to run `p4 verify` after your migration.

Contact Perforce Technical Support for assistance when migrating a Perforce server from Windows to UNIX or vice-versa.

Changing the IP address of your server

If the IP address of the new machine is not the same as that of the old machine, you may need to update any IP-address-based protections in your protections table. See “Administering Perforce: Protections” on page 61 for information on setting protections for your Perforce server.

If you are a licensed Perforce customer, you will also need a new license file to reflect the new IP address. Contact Perforce technical support to obtain an updated license.

Changing the hostname of your server

If the hostname of the new machine serving Perforce is different from that of its predecessor, your users will need to change their `P4PORT` settings. If the old machine is being retired or renamed, consider setting an alias for the new machine to match that of the old machine, so that your users won't have to change their `P4PORT` settings.

Using Multiple Depots


Just as Perforce servers can host multiple depots, Perforce client programs can access files from multiple depots. These other depots may reside within the Perforce server normally accessed by the Perforce client, or they may reside within other, *remote*, Perforce servers.

When using local depots, the user's Perforce client program communicates with the Perforce server specified by the user's `P4PORT` environment variable or equivalent setting.

When using remote depots, the user's Perforce client uses its default Perforce server as a means to access a second, *remote*, Perforce server. Because of this behavior, the client doesn't need to know where the files are actually stored, and doesn't need direct access to the remote Perforce server.

The use of depots on remote servers ("remote depots") is limited to read-only operations; a Perforce client may not add, edit, integrate into, or delete files that reside in depots on other servers. Depots sharing the same Perforce server as the client ("local depots") are not subject to this limitation.

The following diagram illustrates how Perforce clients use a user's default Perforce server to access files on another Perforce server.



Distributed development notes

Remote depots are designed to support shared *code*, not shared *development*. They enable independent organizations with their own Perforce installations to view files and integrate changes from depots in other installations. Remote depots are not a generalized solution for load-balancing or network access problems.

To support shared development, or deal with load-balancing or network access problems, use Perforce Proxy. For details, see “Perforce Proxy” on page 111.

When and when not to use remote depots

The term “remote depot” is actually somewhat misleading. New Perforce customers tend to assume that if their users are geographically distributed, they need to set up separate Perforce installations (servers running `p4d`) and interconnect them with remote depot support.

This is not the case. Perforce is designed to cope with the latencies of large networks and inherently supports users with client workspaces at remote sites. A single Perforce installation is ready, out of the box, to support a shared development project, regardless of the geographic distribution of its contributors.

If you’re doing distributed development, you probably want to use a single Perforce installation, with all code in depots managed by one Perforce server, and (optionally) cached with one or more instances of Perforce Proxy. Partitioning joint development projects into separate Perforce installations will not improve throughput, and usually only complicates administration.

If, however, you regularly import works from other organizations as part of your own organization’s body of software, you may wish to consider using Perforce’s remote depot features. This is what remote depots were designed for: facilitating the sharing of code, not development, across organizations.

Restrictions on remote depots

Users of remote depots encounter several restrictions:

- Access is restricted to read-only operations. You cannot submit files into a remote depot. You cannot edit files in remote depots. You can, however, create a branch in a local depot from files in a remote depot, and then integrate changes from the remote depot into the local branch. This integration is a one-way operation, as you cannot make changes in the local branch and integrate them back into the remote depot.
- Remote depots may be accessed only by Perforce servers running at the same release levels.
- A Perforce server’s metadata (information about client workspaces, changelists, labels, and so on) cannot be accessed using remote depots. Commands like `p4 client only` return specifications of entities defined in the local server’s metadata.

In general, most of the restrictions associated with remote depots are either insignificant (lack of remote access to metadata) or beneficial (read-only access) if you're using remote depots for their intended purpose, namely the sharing of code, not development, between separate organizations.

If you're sharing development between separate organizations, use local depots from a single Perforce server, and, if necessary, cache them at each organization with Perforce Proxy.

Defining new depots

New depots (local or remote) in a server namespace are defined with the command `p4 depot depotname`. Depots may be defined as either *local* or *remote* depots.

Defining local depots

To define a new local depot (that is, a new depot in the current Perforce server namespace), call `p4 depot` with the new depot name, and edit only the `Map:` field in the resulting form. For example, to create a new depot called `book` with the files stored in the local Perforce server namespace in a root subdirectory called `book` (that is, `$P4ROOT/book`), enter the command `p4 depot book`, and fill in the resulting form as follows:

Depot:	book
Type:	local
Address:	subdir
Map:	book/...

Although you can set the `Map:` field to point to a depot directory other than one matching the depot name, there is rarely any advantage to be had by doing so, and it can make life confusing if you (as an administrator) ever need to work with the directories in the server root.

Defining remote depots

Defining a new depot on a remote Perforce server is only slightly more complicated than defining a local depot. Set the `Type:` to `remote`, provide the server's address in the `Address:` field, and set the `Map:` field to map into the remote depot namespace.

Example: Defining a remote depot

Lisa is working on a GUI porting project. She and Ed are using different Perforce servers; his is on host `pine`, and it's listening on port `1818`. Lisa wants to grab Ed's GUI routines for her own use; she knows that Ed's color routine files are located on his Perforce server's single depot under the subdirectory `graphics/GUI`.

Lisa's first step towards accessing Ed's files is to create a new depot. She'll call this depot `gui`; she'd type `p4 depot GUI` and fill in the form as follows:

Depot:	<code>gui</code>
Type:	<code>remote</code>
Address:	<code>pine:1818</code>
Map:	<code>//depot/graphics/gui/...</code>

This creates a remote depot called `gui` on Lisa's Perforce server; this depot (`gui`) maps to Ed's depot's namespace under its `graphics/gui` subdirectory.

The Map: field

The `Map:` field is analogous to a client's view, except that the view may contain multiple mappings and the `Map:` field always contains a single mapping. This single client mapping format changes depending on whether the depot being defined is `local` or `remote`:

- For local depots, the mapping should contain a subdirectory relative to the file space of the Perforce server root directory. For example, `graphics/gui/...` maps to the `graphics/gui` subdirectory of the server root.
- For remote depots, the mapping should contain a subdirectory relative to the remote depot namespace. For example, `//depot/graphic/gui/...` would map to the `graphic/gui` subdirectory of the remote server depot named `depot`.

Note that the mapping subdirectory must always contains the “...” wildcard on its right side.

If you are unfamiliar with client views and mappings, please consult the *Perforce User's Guide* for general information about how Perforce mappings work.

Other depot operations

Naming depots

Depot names share the same namespace as branches, clients, and labels. For example, `//foo` refers uniquely to one of the depot `foo`, the client `foo`, the branch `foo`, or the label `foo`; you can't simultaneously have both a depot and a label named `foo`.

Listing depots

You can list all depots known to the current Perforce server with the `p4 depots` command.

Deleting depots

You can delete depots with `p4 depot -d depotname`.

To delete a depot, it must be empty; you must first obliterate all files in the depot with `p4 obliterate`.

For local depots, `p4 obliterate` deletes the versioned files as well as all their associated metadata. For remote depots, `p4 obliterate` erases *only* the locally held client and label records; the files and metadata still residing on the remote server remain intact.

Before using `p4 obliterate`, and *especially* if you're about to use it to obliterate all files in a depot, read and understand the warnings in "Reclaiming disk space by obliterating files" on page 42.

Limiting access from other servers

Remote depots are always accessed by a virtual user named `remote`, which does not consume a Perforce license. By default, all the files on any Perforce server may be accessed remotely.

To limit or eliminate remote access to a particular server, use `p4 protect` to set permissions for user `remote` on that server.

For example, to eliminate remote access to all files in all depots on a particular server, set the following permission on that server:

```
read user remote * -//...
```

Since `remote` depots can only be used for `read` access, it is not necessary to remove `write` or `super` access.

Users working with multiple depots

If your users will be using remote depots as well as local depots, they should be aware of certain caveats regarding the behavior of files in remote depots as opposed to local depots, as described in the preceding sections.

The *Perforce User's Guide* contains detailed information for users who will be working with more than one depot.

Administering Perforce: Protections

Perforce provides a protection scheme to prevent unauthorized or inadvertent access to the depot. The protections determine which Perforce commands can be run, on which files, by whom, and from which host. Protections are set with the `p4 protect` command.

When Should Protections Be Set?

Before `p4 protect` is run, every Perforce user is a superuser, and can access and change anything in the depot. The first time `protect` is invoked, a protections table is created that gives the invoking user superuser access from all hosts, and lowers everyone else's access level to write permission on all files from all hosts. Therefore, `protect` should be run as the concluding step of all new Perforce installations; the superuser can change the access levels as needed at any time.

The Perforce protections are stored in the `db.protect` file in the server root directory; if `p4 protect` is first run by an unauthorized user (or if you accidentally lock yourself out!) the depot can be brought back to its unprotected state by removing this file.

Setting Protections with “p4 protect”

The `p4 protect` form contains a single form field called `Protections:` that consists of multiple lines. Each line in `Protections:` contains subfields, and the table looks like this:

Example: *A sample protections table:*

Protections:				
read	user	emily	*	//depot/elm_proj/...
write	group	devgrp	*	//...
write	user	*	195.3.24.*	-//...
write	user	joe	*	-//...
write	user	lisag	*	-//depot/...
write	user	lisag	*	//depot/doc/...
super	user	edk	*	//...

(The five fields may not line up vertically on your screen; they are aligned here for readability.)

The permission lines' five fields

Each line specifies a particular permission; each permission is defined by five fields.

The meanings of these fields are:

Field	Meaning
Access Level	Which access level is being granted: <code>list</code> , <code>read</code> , <code>open</code> , <code>write</code> , <code>review</code> , <code>admin</code> , or <code>super</code> . These are described below.
User/Group	Does this protection apply to a user or a group? The value must be <code>user</code> or <code>group</code> .
Name	The user or group whose protection level is being defined. This field may contain the “*” wildcard. A “*” by itself grants this protection to everyone, “*e” grants this protection to every user (or group) whose username ends with an “e”.
Host	<p>The TCP/IP address of the host being granted access. This must be provided as the numeric address of the host in dotted quad notation (for instance, <code>192.168.41.2</code>).</p> <p>This field may contain the “*” wildcard. A “*” by itself means that this protection is being granted for all hosts. The wildcard can be used as in any string, so “<code>192.168.41.*</code>” would define access to any subnet within <code>192.168.41</code>, and “<code>*3*</code>” would refer to any IP address with a “3” in it.</p> <p>Because the client’s IP address is provided by the Internet Protocol itself, this field provides as much security as is provided by the network.</p> <p>If you are using Perforce Proxy, see “P4P and protections” on page 115 to find out how to use host-based protections for users connecting to a Perforce Server from behind a Perforce Proxy.</p>
Files	<p>A file specification representing the files in the depot on which permissions are being granted. Perforce wildcards can be used in the specification.</p> <p>“<code>//...</code>” means all files in all depots.</p>

Access levels

The access level is described by the first value on each line. The seven access levels are:

Level	Meaning
<code>list</code>	Permission is granted to run Perforce commands that display file metadata, such as <code>p4 filelog</code> . No permission is granted to view or change the contents of the files.
<code>read</code>	The user(s) can run those Perforce commands that are needed to read files, such as <code>p4 client</code> and <code>p4 sync</code> . The <code>read</code> permission includes <code>list</code> access.

Level	Meaning
open	Grants permission to read files from the depot into the client workspace, and gives permission to open and edit those files. This permission does not allow the user to write the files back to the depot. <code>open</code> is similar to <code>write</code> , except that with <code>open</code> permission, users are not allowed to run <code>p4 submit</code> or <code>p4 lock</code> . The <code>open</code> permission includes <code>read</code> and <code>list</code> access.
write	Permission is granted to run those commands that edit, delete, or add files. The <code>write</code> permission includes <code>read</code> , <code>list</code> , and <code>open</code> access. This permission allows use of all Perforce commands except <code>protect</code> , <code>depot</code> , <code>obliterate</code> , and <code>verify</code> .
review	A special permission granted to review daemons. It includes <code>list</code> and <code>read</code> access, plus use of the <code>p4 review</code> command. Only review daemons require this permission.
admin	For Perforce administrators; grants permission to run Perforce commands that affect metadata, but not server operation. Provides <code>write</code> and <code>review</code> access plus the added ability to override other users' branch specifications, client specifications, jobs, labels, and change descriptions, as well as to update the <code>typemap</code> table, <code>verify</code> and <code>obliterate</code> files, and to customize job specifications.
super	For Perforce superusers; grants permission to run all Perforce commands. Provides <code>write</code> , <code>review</code> , and <code>admin</code> access plus the added ability to create depots and triggers, edit protections and user groups, delete users, reset passwords, and to shut down the server.

Each Perforce command is associated with a particular minimum access level. For example, to run `p4 sync` on a particular file, the user must have been granted at least `read` access on that file.

The access level required to run a particular command can usually be reasoned from knowledge of what the command does. For example, it is somewhat obvious that `p4 print` would require `read` access. For a full list of the minimum access levels required to run each Perforce command, see “How Protections are Implemented” on page 67.

Which users should receive which permissions?

The simplest method of granting permissions is to give `write` permission to all users who don't need to manage the Perforce system, and give `super` access to those who do, but there are times when this simple solution isn't sufficient.

Read access to particular files should be granted to users who don't ever need to edit those files. For example, an engineer might have `write` permission for source files, but have only `read` access to the documentation, while managers might be granted only `read` access to all files.

Because `open` access allows local editing of files, but doesn't allow these files to be written to the depot, `open` access is usually granted only in unusual circumstances. You might choose `open` access over `write` access when users are testing their changes locally, but when these changes should not be seen by other users. For instance, bug testers may want to change code in order to test theories as to why particular bugs occur, but these changes would be for their own use, and would not be written to the depot. Perhaps a codeline has been frozen, and local changes are to be submitted to the depot only after careful review by the development team. In these cases, `open` access would be granted until the code changes have been approved, after which time the protection level would be upgraded to `write` and the changes submitted.

Default protections

Before `p4 protect` is invoked, every user has superuser privileges. When `p4 protect` is first run, two permissions are set by default. The default protections table looks like this:

<code>write</code>	<code>user</code>	<code>*</code>	<code>*</code>	<code>//...</code>
<code>super</code>	<code>user</code>	<code>edk</code>	<code>*</code>	<code>//...</code>

This indicates that `write` access is granted to all users, on all hosts, to all files. Additionally, the user who first invoked `p4 protect` (in this case, `edk`) is granted superuser privileges.

Interpreting multiple permission lines

The access rights granted to any user are defined by the union of mappings in the protection lines that match her user name and client IP address. (This behavior is slightly different when exclusionary protections are provided and is described in the next section.)

Example: Multiple Permission Lines

Lisa, whose Perforce username is `lisag`, is using a client with the IP address `195.42.39.17`. The protections file reads as follows:

<code>read</code>	<code>user</code>	<code>*</code>	<code>195.42.39.17</code>	<code>//...</code>
<code>write</code>	<code>user</code>	<code>lisag</code>	<code>195.42.39.17</code>	<code>//depot/elm_proj/doc/...</code>
<code>read</code>	<code>user</code>	<code>lisag</code>	<code>*</code>	<code>//...</code>
<code>super</code>	<code>user</code>	<code>edk</code>	<code>*</code>	<code>//...</code>

The union of the first three permissions apply to Lisa. Her username is `lisag`, and she's currently using a client workspace on the host specified in lines 1 and 2. Thus, she can write

files located in the depot's doc subdirectory, but can only read other files. Lisa tries the following:

She types `p4 edit //lisag/doc/elm-help.1`, and is successful.

She types `p4 edit //lisag/READ.ME`, and is told that she doesn't have the proper permission. She is trying to write to a file to which has only read access. She types `p4 sync //lisag/READ.ME`, and this command succeeds, as only read access is needed, and this is granted to her on line 1.

Lisa later switches to another machine with IP address 195.42.39.13. She types `p4 edit //lisag/doc/elm-help.1`, and the command fails; when she's using this host, only the third permission applies to her, and she only has read privileges.

Exclusionary protections

A user can be denied access to particular files by prefacing the fifth field in a permission line with a minus sign (" - "). This is useful for giving most users access to a particular set of files, while denying access to the same files to only a few users.

To use exclusionary mappings properly, it is necessary to understand some of their peculiarities:

- When an exclusionary protection is included in the protections table, the order of the protections is relevant: the exclusionary protection is used to remove any matching protections above it in the table.
- No matter what access level is provided in an exclusionary protection, all access levels for the matching files and IP addresses are denied. The access levels provided in exclusionary protections are irrelevant. The reasons for this counterintuitive behavior are described in the section "How Protections are Implemented" on page 67.

Example: *Exclusionary protections.*

Ed has used `p4 protect` to set up protections as follows:

write	user	*	*	//...
read	user	emily	*	//depot/elm_proj/...
super	user	joe	*	-//...
list	user	lisag	*	-//...
write	user	lisag	*	//depot/elm_proj/doc/...

The first permission looks like it grants write access to all users to all files in all depots, but this is overruled by later exclusionary protections for certain users.

The third permission denies Joe permission to access any file from any host. No subsequent lines grant Joe any further permissions; thus, Joe has been effectively locked out of Perforce.

The fourth permission denies Lisa all access to all files on all hosts, but the fifth permission gives her back write access on all files within a specific directory. If the fourth and fifth lines were switched, Lisa would be unable to run any Perforce command.

Granting Access to Groups of Users

Perforce *groups* simplify maintenance of the protections table. The names of users with identical access requirements can be stored in a single group; the group name can then be entered in the table, and all the users in that group receive the specified permissions.

Groups are maintained with `p4 group` and their protections assigned with `p4 protect`. Only Perforce superusers may use these commands.

Creating and editing groups

If `p4 group groupname` is called with a non-existent *groupname*, a new group named *groupname* is created. Calling `p4 group` with an existing *groupname* allows editing of the user list for this group.

The command `p4 group groupname` displays a form with two fields: `Group:` and `Users:`. The `Group:` field stores the group name, and cannot be edited; `Users:` is empty when the group is first created, and must be filled in. User names are entered under the `Users:` field header; each user name must be typed on its own line, and should be indented. A single user may be listed in any number of groups.

As of Release 99.2, groups can contain other groups, not just users. To add all users in a previously defined group to the group you're presently working with, include the group name in the `Subgroups:` field of the `p4 group` form. User and group names occupy separate namespaces, so groups and users can have the same names.

Groups and protections

To use a group with the `p4 protect` form, specify a group name instead of a user name in any line in the protections table, and set the value of the second field on the line to `group` instead of `user`. All the users in that group will be granted the specified access.

Example: *Granting access to Perforce groups.*

This protections table grants list access to all members of the group devgrp, and super access to user edk:

list	group	devgrp	*	//...
super	user	edk	*	//...

If a user belongs to multiple groups, one permission may override another, but the actual permissions granted to a specific user can be determined by replacing the names of all

groups that a particular user belongs to with the user's name within the protections table, and applying the rules described earlier in this chapter.

Deleting groups

To delete a group, invoke

```
p4 group -d groupname
```

Alternately, invoke `p4 group groupname` and delete all the users from the group in the resulting editor form. The group will be deleted when the form is closed.

How Protections are Implemented

This section describes the algorithm that Perforce follows to implement its protection scheme. Protections can be used properly without reading this section, as the material here is provided to explain the logic behind the behavior described above.

Users' access to files is determined by the following steps:

- The command is looked up in the command access level table shown in “Access Levels Required by Perforce Commands” on page 68 to determine the minimum access level needed to run that command. In our example, `p4 print` is the command, and the minimum access level required to run that command is `read`.
- Perforce makes the first of two passes through the protections table. Both passes move up the protections table, bottom to top, looking for the first relevant line.

The first pass determines whether or not the user is allowed to know whether or not the file exists. This search simply looks for the first line encountered that matches the user name, host IP address, and file argument. If the first matching line found is an inclusionary protection, then the user has permission to at least list the file, and Perforce proceeds to the second pass. Otherwise, if the first matching protection found is an exclusionary mapping, or if the top of the protections table is reached without a matching protection being found, then the user has no permission to even list the file, and will receive a message like `File not on client`.

Example: *Interpreting the order of mappings in the protections table.*

Suppose that our protections table is set as follows:

write	user	*	*	//...
read	user	edk	*	-//...
read	user	edk	*	//depot/elm_proj/...

If Ed runs `p4 print //depot/foo`, Perforce examines the protections table bottom to top, and first encounters the last line. The files specified there don't match the file that Ed wants to print, so this line is irrelevant. The second-to-last line is examined next; this line matches Ed's

user name, his IP address, and the file he wants to print; since this line is an exclusionary mapping, Ed isn't allowed to even list the file.

- If the first pass is successful, a second pass is made at the protections table; this pass is the same as the first, except that access level is now taken into account.

If an inclusionary protection line is the first line encountered that matches the user name, IP address, file argument, and has an access level greater than or equal to the access level required by the given command, then the user is given permission to run the command.

If an exclusionary mapping is the first line encountered that matches according to the above criteria, or if the top of the protections table is reached without finding a matching protection, then the user has no permission to run the command, and will receive the message "You don't have permission for this operation".

Access Levels Required by Perforce Commands

The following table lists the minimum access level required to run each command. For example, since `p4 add` requires at least `open` access, `p4 add` can be run if `open`, `write`, `admin`, or `super` protections are granted.

Command	Access Level	Command	Access Level
<code>add</code>	<code>open</code>	<code>integrate</code> ^d	<code>open</code>
<code>admin</code>	<code>super</code>	<code>integrated</code>	<code>list</code>
<code>annotate</code>	<code>read</code>	<code>job</code> ^{b e}	<code>open</code>
<code>branch</code> ^e	<code>open</code>	<code>jobs</code> ^a	<code>list</code>
<code>branches</code>	<code>list</code>	<code>jobspec</code> ^b	<code>admin</code>
<code>change</code> ^e	<code>open</code>	<code>label</code> ^{a e}	<code>open</code>
<code>changes</code> ^a	<code>list</code>	<code>labels</code> ^{a b}	<code>list</code>
<code>client</code> ^e	<code>list</code>	<code>labelsync</code>	<code>open</code>
<code>clients</code>	<code>list</code>	<code>lock</code>	<code>write</code>
<code>counter</code> ^c	<code>review</code>	<code>obliterate</code>	<code>admin</code>
<code>counters</code>	<code>list</code>	<code>opened</code>	<code>list</code>
<code>delete</code>	<code>open</code>	<code>passwd</code>	<code>list</code>
<code>depot</code> ^{a b}	<code>super</code>	<code>print</code>	<code>read</code>
<code>depots</code> ^a	<code>list</code>	<code>protect</code> ^a	<code>super</code>

Command	Access Level	Command	Access Level
describe	read	reopen	open
describe -s	list	resolve	open
diff	read	resolved	open
diff2	read	revert	open
dirs	list	review ^a	review
edit	open	reviews ^a	list
filelog	list	set	list
files	list	submit	write
fix ^a	open	sync	read
fixes ^a	list	triggers	super
fstat	list	typemap	admin
group ^{a b}	super	unlock ^e	open
groups ^a	list	user ^{a b}	list
have	list	users ^a	list
help	none	verify	admin
info	none	where ^a	none

^a This command doesn't operate on specific files. Thus, permission is granted to run the command if the user has the specified access to at least one file in the depot.

^b The `-o` flag to this command, which allows the form to be read but not edited, requires only list access.

^c list access is required to view an existing counter's value; review access is required to change a counter's value or create a new counter.

^d To run `p4 integrate`, the user needs open access on the target files and read access on the donor files.

^e The `-f` flag to override existing metadata or other users' data requires admin access.

Those commands that list files, such as `p4 describe`, will only list those files to which the user has at least list access.

Some of these commands (for instance, `p4 change`, when editing a previously submitted changelist) take a `-f` flag which can only be used by Perforce superusers. See "Forcing operations with the `-f` flag" on page 45 for details.

Customizing Perforce: Job Specifications

Perforce's jobs feature allows changelists to be linked to enhancement requests, problem reports, and other user-defined tasks. Perforce also offers P4DTI (Perforce Defect Tracking Integration) as a way to integrate third-party defect tracking tools with Perforce. See "Working with third-party defect tracking systems" on page 79 for details.

The Perforce user's use of `p4 job` is discussed in the *Perforce User's Guide*. This chapter covers administrator modification of the jobs system.

Perforce's default jobs template has five fields for tracking jobs. These fields are sufficient for small-scale operations, but as projects managed by Perforce grow, the information stored in these fields may be insufficient. To modify the job template, use the `p4 jobspec` command. You must be a Perforce administrator to use `p4 jobspec`.

This chapter discusses the mechanics of altering the Perforce job template. Certain changes to the template are forbidden. Others are permissible, but are not recommended.

Warning! Improper modifications to the Perforce job template can lead to corruption of your server's database. Recommendations, caveats, and warnings about changes to job templates are summarized at the end of this chapter.

The Default Perforce Job Template

To understand how Perforce jobs are specified, we will examine the default Perforce job template. The examples that follow in this chapter are based upon modifications to the default Perforce job template.

A job created with the default Perforce job template has this format:

```
# A Perforce Job Specification.
#
# 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.
# Description:   Comments about the job. Required.
Job: new
Status: open
User: edk
Date: 1998/06/03 23:16:43
Description:
    <enter description here>
```

The template from which this job was created can be viewed and edited with `p4 jobspec`. The default job specification template looks like this:

```
# A Perforce Job Specification.
#
# Updating this form can be dangerous!
# See 'p4 help jobspec' for proper directions.
Fields:
    101 Job word 32 required
    102 Status select 10 required
    103 User word 32 required
    104 Date date 20 once
    105 Description text 0 required
Comments:
    # A Perforce Job Specification.
    #
    # 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.
    # Description: Comments about the job. Required.
Values:
    Status open/suspended/closed
Presets:
    Status open
    User $user
    Date $now
    Description $blank
```

The Job Template's Fields

There are six fields and field types in the `p4 jobspec` form. These fields define the template for all Perforce jobs stored on your server. The fields and field types are:

Field / Field Type	Meaning
Fields:	A list of fields to be included in each job. Each field consists of an ID#, a name, a datatype, a length, and a setting.
Required:	A list of fields for which values <i>must</i> be entered.
Readonly:	A list of fields, the default values of which cannot be changed by the user. Each field in this list requires a corresponding <code>Presets:</code> entry in the job specification.

Field / Field Type	Meaning
Values:	<p>A list of fields whose datatype is <code>select</code>.</p> <p>For each <code>select</code> field, you must add a line containing the field's name, a space, and its list of acceptable values, separated by slashes.</p>
Presets:	<p>A list of fields and their default values.</p> <p>Values can be either literal strings or variables supported by Perforce.</p>
Comments:	<p>The comments that appear at the top of the <code>p4 job</code> form.</p> <p>These comments are also used by P4Win, the Perforce Windows Client.</p>

The Fields: field

The `p4 jobspec` field `Fields:` lists the fields to be tracked by your jobs, and specifies the order in which they appear on the `p4 job` form.

The default `Fields:` field lists these fields:

```
Fields:
    101 Job word 32 required
    102 Status select 10 required
    103 User word 32 required
    104 Date date 20 once
    105 Description text 0 required
```

Each field must be listed on a separate line, and is comprised of five separate descriptors:

Field Descriptor	Meaning
ID#	<p>A unique integer identifier by which this field is indexed. After a field has been created and jobs entered into the system, the name of this field can change, but the data becomes inaccessible if the ID number changes.</p> <p>ID numbers must be between 101 and 199.</p>
Name	The name of the field as it should appear on the <code>p4 job</code> form.
Data Type	One of five datatypes (<code>word</code> , <code>text</code> , <code>line</code> , <code>select</code> , or <code>date</code>), as described in the next table.

Field Descriptor	Meaning
Length	<p>The recommended size of the field's text box as displayed in P4Win, the Perforce Windows Client. To display a text box with room for multiple lines of input, use a length of 0; to display a single line, enter the <code>Length</code> as the maximum number of characters in the line.</p> <p>The value of this field has no effect on jobs edited from the Perforce command line, and is not related to the actual length of the values stored by the server.</p>
Persistence	<p>Determines whether a field is read-only, contains default values, is required, and so on. The valid values for this field are:</p> <ul style="list-style-type: none">• <code>optional</code>: field can take any value or can be deleted.• <code>default</code>: a default value is provided, but it can be changed or erased.• <code>required</code>: a default is given; it can be changed but the field can't be left empty.• <code>once</code>: read-only; the field is set once to a default value and is never changed.• <code>always</code>: read-only; the field value is reset to the default value when the job is saved. Useful only with the <code>\$now</code> variable to change job modification dates, and with the <code>\$user</code> variable to change the name of the user who last modified the job. <p>In version 98.2 of Perforce, a field's persistence was specified in a very different way. If you have upgraded from 98.2, no conversion need be done; the old persistences will appear in the <code>p4 jobspec</code> form in the new template.</p>

The five field datatypes are:

Field Type	Explanation	Example
word	A single word.	A <code>userid</code> : <code>edk</code>
text	A block of text that can span multiple lines.	A job's description
line	One line of text.	A user's real name: <code>Ed K.</code>

Field Type	Explanation	Example
select	One of a set of values. Each field with datatype <code>select</code> must have a corresponding line in the <code>Values:</code> field entered into the job specification.	A job's status. One of: open/suspended/closed
date	A date value.	The date and time of job creation: 1998/07/15:13:21:46

The Presets: field

All fields with a persistence of anything other than `optional` require default values. To assign a default value to a field, create a line in the jobspec form under `Presets:`, consisting of the field name to which you're assigning the default value. Any single-line string can be used as a default value.

Three variables are available for use as default values:

Variable	Value
<code>\$user</code>	The Perforce user creating the job, as specified by the <code>P4USER</code> environment or registry variable, or as overridden with <code>p4 -u username</code> job.
<code>\$now</code>	The date and time at the moment the job is saved.
<code>\$blank</code>	The text <code><enter description here></code> . When users enter jobs, any fields in your jobspec with a preset of <code>\$blank</code> must be filled in by the user before the job is added to the system.

The lines in the `Presets:` field for the standard jobs template are:

```
Presets:
    Status open
    User $user
    Date $now
    Description $blank
```

The Values: fields

You specify the set of possible values for any field of datatype `select` by entering lines in the `Values:` field. Each line should contain the name of the field, a space, and the list of possible values, separated by slashes.

In the default Perforce job specification, the `Status:` field is the only `select` field, and its possible values are defined as follows:

```
Values:
    Status open/suspended/closed
```

Note Prior to version 2000.1 of Perforce, the `Values:` and `Presets:` fields were specified differently.

If you have scripts that rely upon the old style of jobspecs, you might have to modify them. (Scripts that manipulate jobs, but not jobspecs, do not require modification.)

The Comments: field

The `Comments:` field supplies the comments that appear at the top of the `p4 job` form. Because `p4 job` does not automatically tell your users the valid values of `select` fields, which fields are required, and so on, your comments must tell your users everything they need to know about each field.

Each line of the `Comments:` field must be indented by at least one tab stop from the left margin, and must begin with the comment character `#`.

The comments for the default `p4 job` template appear as:

```
Comments:
    # A Perforce Job Specification.
    # 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.
    # Description: Comments about the job. Required.
```

If you administer a Perforce server and your users use P4Win, the Perforce Windows Client, you must take extra care when writing your comments.

P4Win displays these comments in two ways:

- When the P4Win user creates or edits a job and presses the **Form Info** button in the job dialog box, a popup window displays the comments.
- As the (Windows) cursor moves over each field, the first line of the comment following the colon after the field name in the jobspec is displayed in a ToolTip. The remainder of each of these lines is displayed as the ToolTip for the field that matches the first word of the line. Only the first line of the comment is displayed.

For instance, the ToolTip for the `Status:` field in the preceding jobspec reads:

```
Either 'open', 'closed', or 'suspended'. Can be changed
```

Caveats, Warnings, and Recommendations

Although the material in this section has already been presented elsewhere in this chapter, it is important enough to bear repeating. Please follow the guidelines presented here when editing job specifications with `p4 jobspec`.

Warning! Please read and understand the material in this section before attempting to edit a job specification.

- After a field has been created and jobs have been entered, do not change the field's ID#. Any data entered in that field through `p4 job` will be inaccessible.
- Field names can be changed at any time. When changing a field's name, be sure to also change the fieldname in other `p4 jobspec` fields that reference this fieldname. For example, if you create a new field 106 named `severity` and subsequently rename it to `bug-severity`, then the corresponding line in the jobspec's `Presets:` field must be changed to `bug-severity` to reflect the change.
- The comments that you write in the `Comments:` field are the only way to let your users know the requirements for each field. Make these comments understandable and complete. These comments are treated specially in P4Win, the Perforce Windows Client. For P4Win ToolTip compatibility, the first line of each field's comment should be understandable if read on its own.
- Leave the default fields 101 to 105 in the jobs system, and use `p4 jobspec` only to work with new fields. Do not change the names or types of fields 101 to 105 for any reason.
 - Field ID# 101, the job's name, is required by Perforce and must not be deleted.
 - Field ID# 102, the job's status, can be deleted; however, Perforce will subsequently be unable to update the status of jobs linked to particular changelists. If this field is present, Perforce will always set the value of this field to `closed` when a changelist containing this job is submitted, even if `closed` has been deleted from the list of possible values in the jobspec. This is highly undesirable behavior; do not change the name of field 102.
 - Field ID# 105 is assumed to be a job description by Perforce clients. If present, it is also used by `p4 change` and `p4 submit` to describe the jobs fixed by the changelist.

Example: A Custom Template

The following example shows a more complicated jobspec and the resulting job form:

```
# A Custom Job Specification.
#
# Updating this form can be dangerous!
# See 'p4 help jobspec' for proper directions.
Fields:
    101 Job word 32 required
    102 Status select 10 required
    111 Type select 10 required
    112 Priority select 10 required
    113 Subsystem select 10 required
    114 Reported_by word 32 required
    115 Reported_date date 20 once
    105 Description text 0 required
Values:
    Type bug/sir/problem/unknown
    Status open/closed/suspended
    Priority A/B/C/unknown
    Subsystem server/gui/doc/mac/misc/unknown
Presets:
    Status open
    Type setme
    Priority unknown
    Subsystem setme
    Reported_by $user
    Reported_date $now
    Description $blank
Comments:
    # Custom Job fields:
    # Job:                Job number
    # Status:             Has the job been fixed: Acceptable values are
    #                     'open', 'closed', or 'suspended'
    # 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
    # Reported_by:       Who's fixing the bug
    # Reported_date:     When the bug was first entered
    # Description:       Textual description of the bug
```

Running `p4 job` against this jobspec displays the following job form:

```
# Custom Job fields:
# Job:           Job number
# Status:        Has the job been fixed: Acceptable values are
#                'open', 'closed', or 'suspended'
# 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
# Reported_by:   Who's fixing the bug
# Reported_date: When the bug was first entered
# Description:   Textual description of the bug

Job:      new

Status:   open

Type:     setme

Priority:      unknown

Subsystem:    setme

Reported_by:  edk

Description:
    <enter description here>
```

The order of the listing under `Fields:` in the `p4 jobspec` form determines the order in which the fields appear to users in job forms; fields need not be ordered by numeric identifier.

Working with third-party defect tracking systems

With P4DTI, you can integrate Perforce with any third-party defect tracking or process management software.

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 (such as bug reports, change orders, work assignments) can be converted automatically to Perforce metadata for access by Perforce users.

P4DTI can be easily extended to other products; TeamShare and Bugzilla are the first two integrations published.

P4DTI is open source and available under a FreeBSD-like license.

Using P4DTI - Perforce Defect Tracking Integration

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.

Building your own integration

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. Depending on the application, the interface you set up will consist of one or more of the following:

- A trigger or script on the defect tracking system side that adds, updates, or deletes a job in Perforce every time a bug is added, updated, or deleted in the defect tracking system.

The third-party system should generate the data and pass it to a script which reformats the data to resemble the form used by a manual (interactive) invocation of `p4 job`. The script can then pipe the generated form to the standard input of a `p4 job -i` command.

(The `-i` flag to `p4 job` allows `p4 job` to read a job form directly from the standard input, rather than using the interactive “form-and-editor” approach typical of user operations. Further information on automating Perforce with the `-i` option is available in the *Perforce Command Reference*.)

- A trigger on the Perforce side that checks changelists being submitted for any necessary bug fix information.
- A Perforce review daemon that checks successfully-submitted changelists for job fixes and issues the appropriate commands to update the corresponding data in your defect tracking system.

For more about triggers and review daemons, including examples, see “Scripting Perforce: Daemons and Triggers” on page 83.

Getting more information

In addition to the P4DTI-based TeamTrack and Bugzilla integrations, Perforce customers currently integrate Perforce with their own home-grown defect tracking systems, and with third-party systems such as Remedy, Scopus, and ClearTrack.

If you are interested in seeing what other Perforce users have done, visit the Perforce web site and examine the `perforce-user` mailing list archives, which are available under our Documentation page.

You may also wish to consider posting to `perforce-user` to ask if anyone has integrated Perforce with the third-party tool you're interested in, as someone may have already done all the setup work required to work with your system.

Scripting Perforce: Daemons and Triggers

User-written scripts can enhance Perforce's functionality. There are three primary methods of scripting PERFORCE:

- *Wrappers* are scripts that call Perforce commands. Wrappers can be written in any scripting language, are usually tailored to your own site's needs, and are not discussed here.
- *Daemons* run at predetermined times, looking for changes to the Perforce metadata. When a daemon determines that the state of the depot has changed in some useful way, it runs other commands. For example, a daemon might look for newly submitted changelists and send email to users who have previously updated the files that were submitted in those changelists. Perforce provides a number of tools that make daemon-writing easier.
- *Pre-submit triggers* are scripts that Perforce runs whenever users attempt to submit files. A trigger script returns a value to Perforce that determines whether or not the submit should succeed. For example, you might write a script that watches for a particular executable file to be submitted; when this file is submitted, the trigger script might tell the submit to fail if the release notes file has not been updated within the same changelist.

This chapter assumes that you know how to write scripts.

Triggers

Warning! Although Perforce commands that only read data from the depot can be called in a trigger script, ***running PERFORCE commands that write data to the depot is dangerous and should be avoided.*** In particular, ***do not run the*** `p4 submit` ***command from within a trigger script.***

A *pre-submit trigger* is a script called by Perforce when files that you've specified have been submitted. If the script exits with status 0, the submit continues. If it exits with a nonzero status, the submit fails. In the event of failure, the script's standard output is displayed as part of the error message returned to the user by the failed submit.

Triggers can be useful in many situations. Consider the following common uses:

- To validate submits above and beyond the Perforce protections mechanism. For example, it might ensure that Ed isn't allowed to submit file `foo` until Courtney has submitted file `bar`.

- To tell the submit to fail if file `foo` is not submitted in the same changelist as file `bar`.
- To ensure that every submit to a particular codeline fixes at least one job.

Triggers are created and edited with `p4 triggers`. Only a Perforce superuser can run this command. The `p4 triggers` form looks like this:

Triggers:		
<code>relnotes_check</code>	<code>//depot/bld/...</code>	<code>"perl relcheck.pl %user%"</code>
<code>verify_jobs</code>	<code>//depot/...</code>	<code>"python /usr/bin/job.py %change%"</code>

Each line in the trigger table has three fields:

Field	Meaning
Trigger Name	The name of the trigger. The name can be any arbitrary string.
File Specification	<p>A file specification in depot syntax. If a changelist contains any files that match this specification, the script will be run.</p> <p>Multiple file specifications can be linked to the same trigger by listing the trigger multiple times in the trigger table.</p>
Script	<p>The command that Perforce should run when a matching file is submitted.</p> <p>The submit will continue if the trigger script exits with 0, and fail otherwise.</p> <p>The script must be specified in a way that allows the Perforce server account to find the file. You can place the directory of the script in the <code>PATH</code> of the environment in which <code>p4d</code> is running, or specify the full path name of the script within the trigger table.</p> <p>This command should be quoted, and can take any or all of a number of variables as parameters. The most useful of these variables are:</p> <ul style="list-style-type: none">• <code>%user%</code>, which provides the Perforce name of the user who submitted the changelist;• <code>%changelist%</code>, (also abbreviated as <code>%change%</code>), the number of the changelist that's being submitted; and• <code>%client%</code>, the name of the client workspace from which the submit was run.

Example: Creating a trigger

The development group wants to make sure that whenever a .exe file is submitted to the depot, the release notes for the program are submitted at the same time.

You write a trigger script that takes a changelist number as its only argument, does a p4 opened on the changelist, parses the results to find the files included in the changelist, and ensures that for every executable file that's been submitted, a RELNOTES file in the same directory has been submitted. If the changelist includes a RELNOTES file, the script terminates with an exit status of 0; otherwise the exit status is set to 1.

The script written, you add it to the trigger table by editing the form displayed by p4 trigger:

```
Triggers:
    rnotes //depot/...exe    "/usr/bin/rnotetest.pl %changelist%"
```

Whenever an .exe file is submitted, this trigger is run. If the script fails, it returns a nonzero exit status, and Perforce aborts the submit.

Using triggers

Triggers are run in the order entered in the triggers table.

If you have multiple triggers associated with a file pattern, each will be run in the order in which it appears in the triggers table. If one of these triggers fails, no further triggers are executed.

Example: Multiple triggers on the same file:

In the next few examples, we're using %change% as an abbreviation for %changelist%. Either form is acceptable, as they are interchangeable.

*All *.c files must pass through the scripts check1.sh, check2.sh, and check3.sh:*

```
Triggers:
    check1    //depot/src/*.c    "/usr/bin/check1.sh %change%"
    check2    //depot/src/*.c    "/usr/bin/check2.sh %change%"
    check3    //depot/src/*.c    "/usr/bin/check3.sh %change%"
```

If any trigger fails (say, check1.sh), the submit fails immediately and none of the subsequent triggers (that is, check2.sh and check3.sh) are called. Each time a trigger succeeds, the next matching trigger is run.

Multiple file specifications can be linked to the same trigger by listing the trigger multiple times in the trigger table.

Example: *Activating the same trigger for multiple file specs:*

Triggers:			
bugcheck	//depot/*.c	"/usr/bin/checkit.pl	%change%"
bugcheck	//depot/*.h	"/usr/bin/checkit.pl	%change%"
bugcheck	//depot/*.cpp	"/usr/bin/checkit.pl	%change%"

*In this case, the bugcheck trigger runs on the *.c files, the *.h files, and the *.cpp files.*

Some trigger scripts need to know the files that are included in the changelist. Since `p4d` can only pass 1K of data to a trigger script, the file list can't be passed via the trigger. Use `p4 opened -ac changelist#` in your trigger scripts to get the list of files for the changelist number provided as an argument. The actual contents of the files are not accessible from within the trigger script, since the files are not stored in the depot until the submit completes.

Before `p4 submit` runs the trigger script, it creates the changelist, assigns it a number, and locks the files in the changelist. After a trigger script completes, `p4 submit` may run subsequent trigger scripts that cause the submit to fail. For this reason, trigger scripts should not take any actions that assume the submit will succeed. Trigger scripts are meant primarily for changelist validation; if you need to take particular actions based on the success of a submit, use a daemon.

Note	In order to use triggers, the server (<code>p4d</code>) must be able to “fork”, or spawn off processes to run the triggers. This is the default configuration of Perforce. If you invoke <code>p4d</code> with the <code>-f</code> (run in foreground without forking) option, however, you will not be able to use triggers until you restart the server without the <code>-f</code> option.
-------------	--

Triggers and security

Warning! Because triggers are spawned by the `p4d` process, `p4d` should never be run as `root` on UNIX systems.

Triggers and Windows

By default, the Perforce service runs under the Windows local `System` account.

Because Windows requires a real account name and password to access files on a network drive, if the trigger script resides on a network drive, you must configure the service to use a real userid and password to access the script.

For details, see “Installing the Perforce service on a network drive” on page 107.

Daemons

Daemons are processes that are called periodically or run continuously in the background. Daemons that use Perforce usually work by examining the server metadata as often as needed and taking action as often as necessary. Typical daemon applications include:

- A change review daemon that wakes up every ten minutes to see if any changelists have been submitted to the production depot. If any changelists have been submitted, the daemon sends email to those users who have “subscribed” to any of the files included in those changelists. The message informs them that the files they’re interested in have changed.
- A jobs daemon that generates a report at the end of each day to create a report on open jobs. It shows the number of jobs in each category, the severity each job, and more. The report is mailed to all interested users.
- A Web daemon that looks for changes to files in a particular depot subdirectory. If new file revisions are found there, they are synced to a client workspace that contains the live web pages.

Daemons can be used for almost any task that needs to occur when Perforce metadata has changed. Unlike triggers, which are used primarily for submission validation, daemons can also be used to write information (that is, submit files) to a depot.

Perforce’s change review daemon

The change review daemon described above can be downloaded from <http://www.perforce.com/perforce/loadsupp.html>. It runs under Python, which can be retrieved from <http://www.python.org/>. Before running the script, please be sure to read and follow the configuration instructions included in the script itself.

The change review daemon looks at the files included in each newly submitted changelist and emails those users who have “subscribed” to any of the files included in the changelist, letting those users know that the file(s) they’re interested in have changed.

Users subscribe to files by calling `p4 user` and entering filepatterns in the `Reviews:` field of the resulting form:

```
User:      sarahm
Email:     sarahm@elmco.com
Update:    1997/04/29 11:52:08
Access:    1997/04/29 11:52:08
FullName:  Sarah MacLonnogan
Reviews:
    //depot/doc/...
    //depot.../README
```

Users should enter their email addresses in the `Email:` field, and enter any number of filepatterns corresponding to the files in which they're interested into the `Reviews:` field. The daemon reports changes to these files to the users.

By including the special path `//depot/jobs` in the `Reviews:` field, users can also receive mail from the Perforce change review daemon whenever job data is updated.

The change review daemon implements the following scheme:

1. `p4 counter` is used to read and change a variable, called a *counter*, in the Perforce metadata. The counter used by this daemon, `review`, stores the number of the latest changelist that's been reviewed.
2. The Perforce depot is polled for submitted, unreviewed changelists with the `p4 review -t review` command.
3. `p4 reviews` generates a list of reviewers for each of these changelists.
4. The Python mail module mails the `p4 describe` changelist description to each reviewer.
5. The first three steps are repeated every three minutes, or at some other interval configured the time of installation.

The command used in the fourth step (`p4 describe`) is a straightforward reporting command. The other commands (`p4 review`, `p4 reviews`, and `p4 counter`) are used almost exclusively by review daemons.

Creating other daemons

You can use `p4review.py` as a starting point to create your own daemons, changing it as needed. As an example, another daemon might upload Perforce job information into an external bug tracking system after changelist submission. It would use the `p4 review` command with a new review counter to list new changelists, and use `p4 fixes` to get the list of jobs fixed by the newly submitted changelists. This information might then be fed to the external system, notifying it that certain jobs have been completed.

If you write a daemon of your own and would like to share it with other users, you can submit it into the Perforce Public Depot. For more information, go to <http://www.perforce.com> and follow the "Perforce Public Depot" link.

Commands used by daemons

Certain Perforce commands are used almost exclusively by review daemons.

These commands are:

Command	Usage
<code>p4 counter name [value]</code>	<p>When a <i>value</i> argument is not included, <code>p4 counter</code> returns the value of the variable <i>name</i>.</p> <p>When a <i>value</i> argument appears, <code>p4 counter</code> sets the value of the variable <i>name</i> to <i>value</i>.</p> <p>Requires at least <code>review</code> access to run.</p> <p>WARNING: The review counters named <code>journal</code>, <code>job</code>, and <code>change</code> are used internally by Perforce; <i>use of any of these three names as review numbers could corrupt the Perforce database.</i></p> <p>For Release 99.2 and above, Perforce will not let you change the values of <code>journal</code>, <code>job</code>, and <code>change</code>.</p> <p>Counters are represented internally as signed ints.</p> <p>For most platforms, the largest value that can be stored in a counter is $2^{31} - 1$, or 2147483647.</p> <p>A server running on a 64-bit platform can store counters up to $2^{63} - 1$, or 9223372036854775807</p>
<code>p4 counters</code>	List all counters and their values.
<code>p4 review -c change#</code>	<p>For all changelists between <i>change#</i> and the latest submitted changelist, this command lists the changelists' numbers, creators, and creators' email addresses.</p> <p>Requires at least <code>review</code> access to run.</p>

Command	Usage
<code>p4 reviews -c change# filespec</code>	Lists all users who have subscribed to review the named files or any files in the specified changelist. It is hard to imagine any use for this command outside of our own change review daemon.
<code>p4 changes -m 1 -s submitted</code>	Output a single line showing the changelist number of the last submitted changelist, as opposed to the highest changelist number known to the Perforce server.

Daemons and counters

If you're writing a change review daemon or other daemon that deals with submitted changelists, you may also wish to keep track of the changelist number of the last *submitted* changelist, which is the second field in the output of a `p4 changes -m 1 -s submitted` command.

This is *not* the same as the output of `p4 counter change`. The last changelist number known to the Perforce server (the output of `p4 counter change`) includes pending changelists created by users, but not yet submitted to the depot.

Scripting and buffering

Depending on your platform, the output of individual `p4` commands may be fully-buffered (output flushed only after a given number of bytes generated), line-buffered (as on a tty, one line sent per linefeed), or unbuffered.

In general, stdout to a file or pipe is fully-buffered, and stdout to a tty is line-buffered. If your trigger or daemon requires line-buffering (or no buffering), you can disable buffering by supplying the `-v0` debug flag to the `p4` command in question.

If you're using pipes to transfer standard output from a Perforce command (with or without the `-v0` flag), you may also experience buffering issues introduced by the kernel, as the `-v0` flag can only unbuffer the output of the command itself.

Tuning Perforce for Performance

Your Perforce server should normally be a light consumer of system resources. As your installation grows, however, you may wish to revisit your system configuration to ensure that it is configured for optimal performance.

The following chapter briefly outlines some of the factors that can affect the performance of a Perforce server, provides a few tips on diagnosing network-related difficulties, and offers some suggestions on decreasing server load for larger installations.

Tuning for Performance

The following variables can affect the performance of your Perforce server.

Memory

Server performance is highly dependent upon having sufficient memory. Two bottlenecks are relevant: the first can be avoided by ensuring that the server doesn't page when running large queries, and the second by ensuring that the `db.rev` table (or at least as much of it as practical) can be cached in main memory.

- Determining memory requirements for large queries is fairly straightforward: the server requires about 1KB/file of RAM to avoid paging; 10,000 files will require 10MB of RAM.
- To cache `db.rev`, the size of the `db.rev` file in an existing installation can be observed and used as an estimate. New installations of Perforce can expect `db.rev` to require about 150-200 bytes per revision, and roughly 3 revisions per file, or about 0.5KB of RAM per file.

Thus, if there is 1.5KB of RAM available per file, or 150MB for 100,000 files, the server will not page, even when performing an operation involving all files. It is still possible that multiple large operations will be performed simultaneously and thus require more memory to avoid paging. On the other hand, the vast majority of operations will only involve a small subset of files.

For most installations, a system with enough RAM for 1.5KB per file in the depot will suffice.

Filesystem performance

Perforce is judicious with regards to its use of disk I/O; its metadata is well-keyed and accesses are mostly sequential scans of limited subsets of the data.

The only disk-intensive activity is file check-in, where the Performer server must write and rename files in the archive. Server performance depends heavily upon the operating system's filesystem implementation, and in particular, whether directory updates are synchronous.

Although Performer does not recommend any specific filesystem, Linux servers are generally fastest (owing to Linux's asynchronous directory updating), but may have poor recovery if power is cut at the wrong time. The BSD filesystem (also used in Solaris) is relatively slow, but much more reliable. NTFS performance falls somewhere in between these two ranges. The filesystems used by IRIX and OSF have demonstrated an excellent combination of both speed and robustness.

Performance in systems where database and versioned files are stored on NFS-mounted volumes is typically dependent on the implementation of NFS in question and/or the underlying storage hardware. Performer has been tested and is supported under the Solaris implementation of NFS.

Under Linux and FreeBSD, database updates over NFS can be an issue as file locking is relatively slow; if the journal is NFS-mounted on these platforms, all operations will be slower. In general (but in particular on Linux and FreeBSD) we recommend that the Performer database, depot, and journal files be stored on disks local to the machine running the Performer server process.

These issues affect only the Performer Server process (`p4d`). Performer client programs, (such as `p4`, the Performer Command-Line Client) have always been able to work with client workspaces on NFS-mounted drives (for instance, workspaces in users' home directories).

Disk space allocation

Performer disk space usage is a function of three variables:

- Number and size of client workspaces
- Size of server database
- Size of server's archive of all versioned files

All three variables depend on the nature of your data and how heavily you use Performer.

The client file space required is the size of the files that your users will need in their client workspaces at any one time.

The server's database size can be calculated with a fair level of accuracy; as a rough estimate, it requires 0.5KB per user per file. (For instance, a system with 10,000 files and 50 users will require 250M of disk space for the database). The database can be expected to grow over time as histories of the individual files grow.

The size of the server's archive of versioned files depends on the sizes of the original files stored and grows as revisions are added. For most sites, allocate space equivalent to at least three times the aggregate size of the original files.

If you anticipate your database growing into the gigabyte range, you should ensure that your platform has adequate support for large filesystems. See "Allocate disk space for anticipated growth" on page 20.

The `db.have` file holds label contents and the list of files opened in client workspaces, and tends to grow the faster than other files in the database. If you are experiencing issues related to the size of your `db.have` file and are unable to quickly switch to a server with adequate support for large files, deleting unused clients and labels and reducing the scope of client views can help alleviate the problem.

Network

Perforce can run over any TCP/IP network. Although we have not yet seen network limitations, the more bandwidth the better. Presumably FDDI would be better than 10Mb/s Ethernet, but some users have reported that using a T1 (1.5 Mb/s) provides response times comparable to using Perforce locally. Perforce employees work successfully over ISDN (64 Kb/s) lines.

Perforce uses a TCP connection for each client interaction with the server. The server's port address is defined by `P4PORT`, but the TCP/IP implementation picks a client port number. After the command completes and the connection is closed, the port is left in a state called `TIME_WAIT` for two minutes. While the port number ranges from 1025 to 32767, generally only a few hundred or thousand can be in use simultaneously. It is therefore possible to occupy all available ports by invoking a Perforce client command many times in rapid succession, such as with a script.

Before release 99.2, both the server and client side of the connection remained in `TIME_WAIT`, which meant that a script running on one user's machine could deprive other users of service by tying up all available ports on the server side. As of Release 99.2, only the client side goes into `TIME_WAIT`, leaving the Perforce server free to handle other clients.

CPU

Perforce is based on a client/server architecture. Both the client and server are lightweight in terms of CPU resource consumption. By way of example, a server supporting 80 users on a low-end (140 MHz) SPARC Ultra server can use as little as 7 CPU-minutes per day, or about 0.5% of available processing power. Weighting this for peak use and headroom, such a server could support upwards of 800 users.

In general, CPU power is not a major consideration when determining the platform on which to install a Perforce server.

Diagnosing Slow Response Times

Perforce is normally a light user of network resources. While it is possible that an extremely large user operation could cause the Perforce server to respond slowly, consistently slow responses to `p4` commands are usually caused by network problems. Any of the following may cause slow response times:

1. misconfigured domain name system (DNS)
2. misconfigured Windows networking
3. difficulty accessing the `p4` executable on a networked file system

A good initial test is to run `p4 info`. If this does not respond immediately, then there is a network problem. Although solving network problems is beyond the scope of this manual, here are some suggestions for troubleshooting them:

Hostname vs. IP address

On a client machine, try setting `P4PORT` to the server's IP address instead of its hostname. For example, instead of using

```
P4PORT=host.domain:1666
```

try using:

```
P4PORT=1.2.3.4:1666
```

with your site-specific IP address and port number.

On most systems, you can determine the IP address of a host by invoking:

```
ping hostname
```

If `p4 info` responds immediately when you use the IP address, but not when you use the hostname, the problem is likely related to DNS.

Try `p4 info` vs. P4Win

If you are using P4Win, you can compare the response of P4Win's "Show Connection Info" (**Help -> Show Connection Info**) with the response from the command-line `p4 info`.

If the former is fast and the latter is slow, you have a DNS-related problem. (When the Perforce server receives a `p4 info` request, it does a reverse name lookup in order to send back the client and server hostnames along with other configuration information. When it

receives a P4Win “Show Connection Info” request, however, it simply returns the IP addresses.)

Note | This test is only valid for Release 99.1 and newer servers. In releases prior to 99.1, the server always did a reverse name lookup, whether the request was coming from `p4 info` or `P4win`

Windows wildcards

In some cases, `p4` commands using unquoted file patterns with a combination of depot syntax and wildcards, such as:

```
p4 files //depot/*
```

can result in a delayed response on Windows. You can prevent the delay by putting double quotes around the file pattern, like so:

```
p4 files "//depot/*"
```

The cause of the problem is the `p4` command’s use of a Windows function to expand wildcards. When quotes are not used, the function interprets `//depot` as a networked computer path and spends time in a futile search for a machine on the network named `depot`.

DNS lookups and the hosts file

On Windows, the `%SystemRoot%\system32\drivers\etc\hosts` file can be used to hardcode IP address-hostname pairs. You may be able to work around DNS problems by adding entries to this file.

The corresponding UNIX file is `/etc/hosts`.

Location of the “p4” executable

If none of the above diagnostic steps explains the sluggish response time, it’s possible that the `p4` executable itself is on a networked file system which is performing very poorly. To check this, try running:

```
p4 -V
```

This merely prints out the version information, without attempting any network access. If you get a slow response, network access to the `p4` executable itself may be the problem. Copying or downloading a copy of `p4` onto a local filesystem should improve response times.

Preventing Server Swamp

Generally, Perforce's performance depends on the number of files a user tries to manipulate in a single command invocation, not the size of the depot. That is, syncing a client view of 30 files from a 3,000,000-file depot should not be much slower than syncing a client view of 30 files from a 30-file depot.

The number of files affected by a single command is largely determined by:

- p4 command line arguments (or selected folders in the case of GUI operations).

Without arguments, most commands will operate on, or at least refer to, all files in the view.

- Client views, branch views, label views, and protections.

Because commands without arguments operate on all files in the view, it follows that the use of unrestricted views and unlimited protections can result in commands operating on all files in the depot.

When the server answers a request, it locks down the database for the duration of the computation phase. For normal operations, this is a successful strategy, as it can “get in and out” quickly enough to avoid a backlog of requests. Abnormally large requests, however, can take seconds, sometimes even minutes. If frustrated users hit CTRL-C and retry, the problem gets even worse; the server consumes more memory and responds even more slowly.

At sites with very large depots, unrestricted views and unqualified commands will make a Perforce server work much harder than it needs to. Users and administrators can ease load on their servers by:

- Using “tight” views
- Assigning protections
- Limiting `maxresults`
- Writing efficient scripts

Using tight views

The following “loose” view is trivial to set up but could invite trouble on a very large depot:

<code>//depot/...</code>	<code>//workspace/...</code>
--------------------------	------------------------------

In the loose view, the entire depot was mapped into the client workspace; for most users, this can be “tightened” considerably. The following view, for example, is restricted to specific areas of the depot:

//depot/main/srv/devA/...	//workspace/main/srv/devA/...
//depot/main/drv/lport/...	//workspace/main/drv/lport/...
//depot/rel2.0/srv/devA/bin/...	//workspace/rel2.0/srv/devA/bin/...
//depot/qa/s6test/dvr/...	//workspace/qa/s6test/dvr/...

Client views, in particular, but also branch views and label views, should also be set up to give users just enough scope to do the work they need to do.

Client, branch, and label views are set by a Perforce administrator or by individual users with the `p4 client`, `p4 branch`, and `p4 label` commands respectively.

Two of the techniques for script optimization (described in “Using branch views” on page 101 and “The temporary client trick” on page 102) rely on similar techniques. By limiting the size of the view available to a command, fewer commands need to be run, and when run, the commands require fewer resources.

Assigning protections

Protections (see “Administering Perforce: Protections” on page 61) are actually another type of Perforce view. Protections are set with the `p4 protect` command and control which depot files can be affected by commands run by users.

Unlike client, branch, and label views, however, the views used by protections can be set only by Perforce superusers. (Protections also control read and write permission to depot files, but the permission levels themselves have no impact on server performance.) By assigning protections in Perforce, a Perforce superuser can effectively limit the size of a user’s view, even if the user is using “loose” client specifications.

Protections can be assigned to either users or groups. For example:

write	user	sam	*	//depot/admin/...
write	group	rocketdev	*	//depot/rocket/main/...
write	group	rocketrel2	*	//depot/rocket/rel2.0/...

Perforce groups are created by superusers with the `p4 group` command. Not only do they make it easier to assign protections, but they provide useful fail-safe mechanisms in the form of `maxresults` and `maxscanrows`, described in the next section.

Limiting database queries

Each Perforce group has an associated *maxresults* and *maxscanrows* value. The default for each is “unlimited”, but a superuser can use `p4 group` to limit it for any given group.

Users in such groups are unable to run any commands which affect more database rows than the group’s *maxresults* limit. (For most commands, the number of database rows affected is roughly equal to the number of files affected.)

Like *maxresults*, *maxscanrows* prevents certain user commands from placing excessive demands on the server. (For most commands, the number of rows that could be scanned is roughly equal to the number of files affected, multiplied by the average number of revisions per file in the depot.)

To set these limits, fill in the `Maxresults:` or `Maxscanrows:` field in the `p4 group` form. If a user is listed in multiple groups, the *highest* of the *maxresults* (or *maxscanrows*) *limits* (but *not* including the default “unlimited” setting) for those groups is taken as the user’s *maxresults* (or *maxscanrows*) value.

Example: *Effect of setting `maxresults` and `maxscanrows`:*

As an administrator, you wish members of the group `rocketdev` to be limited to operations of 20,000 files or less, and to scan no more than 100,000 revisions:

Group:	rocketdev
Maxresults:	20000
Maxscanrows:	100000
Users:	
	bill
	ruth
	sandy

Suppose that Ruth has an unrestricted (“loose”) client view. When she types:

```
p4 sync
```

her `sync` command is rejected if the depot contains more than 20,000 files. She can work around this limitation either by restricting her client view, or, if she needs all of the files in the view, by syncing smaller sets of files at a time, like so:

```
p4 sync //depot/projA/...
p4 sync //depot/projB/...
```

Either way, she’ll get her files, but without tying up the server to process a single extremely large command.

If Ruth tries a command that scans every revision of every file, such as:

```
p4 filelog //depot/projA/...
```

and there are less than 20,000 files, but more than 100,000 revisions (perhaps the projA directory contains 8000 files, each of which has 20 revisions), the `maxresults` limit will not apply, but the `maxscanrows` limit will.

To remove any limits on the number of result lines processed (or database rows scanned) for a particular group, set the `Maxresults:` or `Maxscanrows:` value for that group to `unlimited`.

As these limitations can make life difficult for your users, do not use them unless you find that certain operations are slowing down your server. The `Maxresults:` value should never be less than 10,000, since certain operations performed by P4Win, the Perforce Windows Client, may require a `Maxresults:` value of between 5,000 and 8,000. Similarly, `Maxscanrows` should rarely need to be set below 50,000.

For more information, including a comparison of Perforce commands and the number of files they affect, type:

```
p4 help maxresults
p4 help maxscanrows
```

from the command line.

Maxresults and maxscanrows for users in multiple groups

As mentioned earlier, if a user is listed in multiple groups, the highest `maxresults` limit of all the groups a user belongs to is the limit that affects the user. The default value of “unlimited” is *not* a limit; if a user is in a group where `maxresults` is set to “unlimited”, he or she is still limited by the highest `maxresults` (or `maxscanrows`) limit of the other groups of which he or she is a member. A user’s commands are truly unlimited only when the user belongs to no groups, or when all of the groups of which the user is a member have their `maxresults` set to “unlimited”

A side effect of this is that you can’t create a group that assigns “unlimited” `maxresults` values to superusers, because if any of the users in such a group were to belong to another group, the “unlimited” limit from the superuser group would also apply to them. You can get around this by assigning a very high `maxresults` limit to your superusers group.

For example:

Group:	superusers
Maxresults:	10000000
Maxscanrows:	100000000

(The largest possible `maxresults` or `maxscanrows` limit is platform-dependent; on most platforms, this is a 32-bit integer.)

Scripting efficiently

The Perforce Command-Line Client, `p4`, allows you to do anything in a script that you can do interactively. The Perforce server can process commands far faster than users can issue them, so in an all-interactive environment, response time is excellent. However, `p4` commands issued by scripts -- triggers, review daemons, or command wrappers, for example -- can cause performance problems if you haven't paid attention to their efficiency. This is not because `p4` commands are inherently inefficient, but because the way one invokes `p4` as an interactive user isn't necessarily suitable for repeated iterations.

This section points out some common efficiency problems and solutions.

Iterating through files

Each Perforce command issued causes a connection thread to be created and a `p4d` subprocess to be started. Reducing the number of Perforce commands your script runs is the first step to making it more efficient.

To this end, scripts should never iterate through files running Perforce commands when they can accomplish the same thing by running one Perforce command on a list of files and iterating through the command results.

For example, try an approach like this:

```
for i in `p4 diff2 path1/... path2/...`
do
    [process diff output]
done
```

Instead of this:

```
for i in `p4 files path1/...`
do
    p4 diff2 path1/$i path2/$i
    [process diff output]
done
```

Using list input files

Any Perforce command that accepts a list of files as a command line argument can also read the same argument list from a file. Scripts can make use of the list input file feature by building up a list of files first, then passing the list file to `p4 -x`.

For example, if your script currently does something like:

```
for components in foo bar ola
do
    p4 edit ${component}.h
done
```

a more efficient alternative would be:

```
for components in foo bar ola
do
    echo ${component}.h >> LISTFILE
done
p4 -x LISTFILE edit
```

The `-x` flag instructs `p4` to read arguments, one per line, from the named file. If the file is specified as “-” (a dash), the standard input is read.

Using branch views

Branch views can be used with `p4 integrate` or `p4 diff2` to reduce the number of Perforce command invocations. For example, if you have a script that runs:

```
p4 diff2 pathA/src/...    pathB/src/...
p4 diff2 pathA/tests/... pathB/tests/...
p4 diff2 pathA/doc/...   pathB/doc/...
```

you can make it more efficient by creating a branch view that looks like this:

Branch:	pathA-pathB	
View:		
	pathA/src/...	pathB/src/...
	pathA/tests/...	pathB/tests/...
	pathA/doc/...	pathB/doc/...

and replacing the three commands with one:

```
p4 diff2 -b pathA-pathB
```

Limiting label references

Repeated references to large labels can be particularly costly. Commands that refer to files using labels as revisions will scan the whole label once for each file argument. To keep from hogging the Perforce server, your script should get the labeled files from the server, then scan the output for the files it needs.

For example, this:

```
p4 files path/...@label | egrep "path/foo.h|path/bar.h|path/ola.c"
```

will impose a lighter load on the Perforce server than either this:

```
p4 files path/foo.h@label path/bar.h@label path/ola.h@label
```

or this:

```
p4 files path/foo.h@label
p4 files path/bar.h@label
p4 files path/ola.h@label
```

The “temporary client” trick described below may also reduce the number of times you have to refer to files by label.

The temporary client trick

Most Perforce commands can process all the files in the current client view with a single command line argument. By making use of a temporary client view that contains the files on which you want to work, you may be able to reduce the number of commands you have to run, and/or to reduce the number of file arguments you need to give each command.

For instance, suppose your script runs these commands:

```
p4 sync pathA/src/...@label
p4 sync pathB/tests/...@label
p4 sync pathC/doc/...@label
```

You can combine the command invocations and reduce the three label scans to one by using a client spec that looks like:

Client:	XY-temp	
View:		
	pathA/src/...	//XY-temp/pathA/src/...
	pathB/tests/...	//XY-temp/pathB/tests/...
	pathC/doc/...	//XY-temp/pathC/doc/...

and running:

```
p4 -c XY-temp sync @label
```

Checkpoints for Database Tree Rebalancing

Perforce's internal database stores its data in structures called Bayer trees, more commonly referred to as B-trees. While B-trees are a very common way to structure data for rapid access, over time the process of adding and deleting elements to and from the trees can eventually lead to imbalances in the data structure.

Eventually, the tree may become sufficiently unbalanced that performance is negatively affected. The Perforce checkpoint and restore processes re-create the trees in a balanced manner, and consequently, you may see some increase in server performance following a restoration from a checkpoint.

Rebalancing the trees is normally only useful if the database files have become more than about 10 times the size of the checkpoint. Given the length of time required for the trees to become unbalanced during normal Perforce use, we expect that the majority of sites will never need to restore the database from a checkpoint (that is, rebalance the trees) for performance reasons.

This chapter describes certain information of specific interest to administrators who set up and maintain Perforce servers on Windows.

Note | Unless otherwise specified, the material presented here applies equally to Windows NT, Windows 2000, and Windows XP.

Using the Perforce installer

The Perforce installer program, `perforce.exe`, gives you the option to install either as a user (the Perforce client), a typical administrator (Perforce installed as a Windows service), a custom administrator (Perforce installed as a service with additional customization options), or to uninstall Perforce from your system.

If you have Administrator privileges, it is usually best to install Perforce as a service. If you don't, install it as a server.

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

Upgrade notes

The Perforce installer also automatically upgrades all types of Perforce servers (or services), even versions prior to 97.3. The upgrade process is extremely conservative; if anything fails at any step in the upgrade process, the installer will stop the upgrade, and you will still be able to use your old server (or service).

Installation options

When you invoke the installer, it presents an initial screen that lists the revisions of the Perforce software you're about to install. You are offered the choice between:

- a user install,
- a typical Administrator install,
- a customized Administrator install, or
- uninstalling Perforce.

User install

The “user install” installs only the Perforce Command-Line Client (`p4.exe`), Perforce Windows Client (P4Win), and (optionally) the third-party SCM plug-in. If you are running on Windows 95 or Windows 98, this will be the only installation option available to you. Under Windows 2000 or higher, this option requires Power User privileges.

You are prompted to specify the location of the client executables, the port (`P4PORT`) on which the client should attempt to contact the Perforce server, the default editor, and the default username.

When specifying the port for the client to use, don’t forget to include the hostname in the form `hostname:port`. See “Telling Perforce client programs which port to talk to” on page 13 for details on `P4PORT`.

If the installer detects older versions of Perforce client or server software on the machine, you are given the option to rename the old executables to prevent `PATH`-dependent conflicts.

Administrator typical

The “typical administrator install” allows the installation of both client and server software for Perforce. This option requires administrator privileges.

You are prompted to specify the directory for the client and server executables, the port on the local machine where the Perforce server or service will listen to client requests (`P4PORT`), the default editor, and the default username.

The installer selects default locations for the `P4LOG` error log file and the `journal` file. If an earlier version of Perforce was installed on the machine, these locations are based on those already in use.

If you have Administrator privileges, the installer installs Perforce and configures it to run as an auto-starting service. The service is set up and started after the installation is complete, and automatically restarts whenever the machine is rebooted. If you do not have Administrator privileges, a shortcut to run Perforce as a server is placed into your Start menu.

If the installer detects older versions of Perforce client or server software on the machine, you are given the option to rename the old executables to prevent `PATH`-dependent conflicts.

Administrator custom

The “custom administrator install” allows the installation of both client and server software for Perforce. This option requires administrator privileges.

As with the typical administrator install, you are prompted to specify the location of client and server executables, the port on the local machine where the Perforce server or service will listen to client requests, the default editor, and the default username.

Unlike the typical administrator install, you are allowed to optionally specify separate directories for the client and server executables, as well as server root, server port, and whether to set up Perforce as an auto-starting (or non-auto-starting) service or server process. The locations of any existing `P4LOG` file and `journal` file are displayed for reference, and may be changed later using `p4 set`.

If you try to install a Perforce service while another Perforce server is running, you will see the error message:

```
Setup has determined that a Perforce Server could be running. Please  
shut down all Perforce Servers before continuing the installation.
```

Failure to shut down the running Perforce server(s) will result in conflicts between the newly installed service and the existing server.

As with the other installation options, if the installer detects older versions of Perforce client or server software on the machine, you are given the option to rename the old executables to prevent `PATH`-dependent conflicts.

Uninstalling Perforce

Should you wish to remove Perforce from a Windows machine, run `perforce.exe` and select the **Uninstall** option. This option requires administrator privileges.

The uninstall procedure removes everything *except* your server data; the Perforce server, service, and client executables, registry keys, and service entries are all deleted. The database and depot files in your server root, however, are always preserved.

Windows services vs. Windows servers

To run any task as a Windows *server*, a user account must be logged in, as shortcuts in a user's `Startup` folder cannot be run until that user logs in. A Windows *service*, on the other hand, is invoked automatically at boot time, and runs regardless of whether or not a user is logged in to the machine.

Throughout most of the documentation set, the terms “Perforce server” or “`p4d`” are used to refer to “the process at the back end that manages the database and responds to requests from Perforce clients”. Under Windows, this can lead to ambiguity; the back-end process can run as either a service (`p4s.exe`, which runs as a thread) or as a server (`p4d.exe`, which runs as a regular process). From a Windows administrator's point of view, these are important distinctions. Consequently, the terminology used in this chapter uses the more precise definitions.

The Perforce service (`p4s.exe`) and the Perforce server (`p4d.exe`) executables are copies of each other; they are identical apart from their filenames. When run, they use the first three characters of the name with which they were invoked (that is, either `p4s` or `p4d`) to determine their behavior. For example, invoking copies named `p4smyserver.exe` or `p4dmyservice.exe` will invoke a service and a server, respectively.

Starting and stopping the Perforce service

If Perforce was installed as a service, a user with Administrator privileges can start and stop it using the **Services** applet in the **Control Panel**.

If you are running at Release 99.2 or above, you can also use the command:

```
p4 admin stop
```

to stop the Perforce service.

Starting and stopping the Perforce server

If Perforce was installed as a server, there should be a “Perforce Server” shortcut in your **Start** menu. To start the server, double-click on the shortcut. To stop the server, right-click on the “Perforce Server” button in the taskbar and select “**Close**”.

You can also start the Perforce server manually from an MS-DOS window. The server executable, `p4d.exe`, is normally found in your `P4ROOT` directory. To start the server, first make sure your current `P4ROOT`, `P4PORT`, `P4LOG`, and `P4JOURNAL` settings are correct, then run: `%P4ROOT%\p4d`

If you want to start a server using settings different than those set by `P4ROOT`, `P4PORT`, `P4LOG`, or `P4JOURNAL`, you can use `p4d` command line flags. For example:

```
c:\test\p4d -r c:\test -p 1999 -L c:\test\log -J c:\test\journal
```

will start a Perforce server process with a root directory of `c:\test`, listening to port 1999, logging errors to `c:\test\log`, and with a journal file of `c:\test\journal`.

Note that `p4d` command line flags are case sensitive.

If you are running at Release 99.2 or above, use the following command:

```
p4 admin stop
```

to stop the Perforce server.

Note	If you are running Release 99.1 or earlier, type Ctrl-C in the MS-DOS window, or simply Close the MS-DOS window. Although this works in all versions of Perforce, this method of shutting down the server is not necessarily “clean”, and with the availability of the <code>p4 admin stop</code> command in 99.2, is no longer recommended.
-------------	--

Installing the Perforce service on a network drive

By default, the Perforce service runs under the local `System` account. Because the `System` account has no network access, a real userid and password are required in order to make the Perforce service work if the metadata and depot files are stored on a network drive.

If you are installing your server root on a network drive, the Perforce installer (`perforce.exe`) requests a valid combination of userid and password at the time of installation. This user must have administrator privileges. (The service will also run as this user, rather than `System`)

While the Perforce NT service will run reliably using a network drive as the server root, there is still a marked performance penalty due to of increased network traffic and slower file access. For this reason, we recommend that the depot files and Perforce database be on a drive local to the machine on which the Perforce service is running.

Multiple Perforce services under Windows

By default, the Perforce installer for Windows installs a single Perforce server as a single service. If you wish to host more than one Perforce installation on the same machine (for instance, one for production and one for testing), then the additional Perforce servers must be started manually.

You can, however, install additional services to start up additional Perforce servers (that is, running as Windows services) automatically.

Note	You must download Perforce 99.1/10994 or a later release in order to use this procedure.
-------------	--

If your intent is to set up multiple services to increase the number of users you support without purchasing more user licenses, you are violating the terms of your Perforce License Agreement.

Before you begin, you should read and understand “Windows configuration parameter precedence” on page 108. Once you understand the precedence of environment variables in determining Perforce configuration, you’ll find setting up multiple Perforce services to be straightforward. To add a new service, you need to:

1. Create a new directory for the Perforce service.
2. Copy the server executable, service executable, and your license file into this directory.
3. Create the new Perforce service using the `svcinst.exe` utility, as described in the example below.
4. Set up the environment variables and start the new service.

We recommend that you install your first Perforce service using the Perforce installer. This first service will be called “Perforce” and its server root directory will contain a few files which are required for the other Perforce services.

Example: *Adding a second Perforce service.*

You wish to create a second Perforce service with a root in C:\p4root2 and a service name of “Perforce2”. You’re running Release 99.1/10994 or greater, and the svcinst executable is in the server root of his first Perforce NT service, which is in C:\perforce.

Verify that your p4d.exe executable is at Release 99.1/10994 or greater:

```
p4d -V
```

(We’ll assume the revision level is correct; if not, you’d have to download a newer version from <http://www.perforce.com> and upgrade the server before continuing.)

Create a P4ROOT directory for the new service:

```
mkdir c:\p4root2
```

Copy the server executables - both p4d.exe (the server) and p4s.exe (the service) - and your license file into the new directory:

```
copy c:\perforce\p4d.exe c:\p4root2
copy c:\perforce\p4d.exe c:\p4root2\p4s.exe
copy c:\perforce\license c:\p4root2\license
```

Use svcinst.exe, (the service installer) to create the “Perforce2” service:

```
svcinst create -n Perforce2 -e c:\p4root2\p4s.exe -a
```

After the “Perforce2” service has been created, set the service parameters for the “Perforce2” service:

```
p4 set -S Perforce2 P4ROOT=c:\p4root2
p4 set -S Perforce2 P4PORT=1667
p4 set -S Perforce2 P4LOG=log2
p4 set -S Perforce2 P4JOURNAL=journal2
```

Finally, use the service installer to start the “Perforce2” service:

```
svcinst start -n Perforce2.
```

The second service is now running, and both services will start automatically the next time you reboot.

Windows configuration parameter precedence

Under Windows, Perforce configuration parameters may be set in many different ways. When a Perforce client program (such as p4 or P4Win), or a Perforce server program (p4d) starts up, it reads its configuration parameters according to the following precedence:

1. The program’s command line flags have the highest precedence.

2. The `P4CONFIG` file, if `P4CONFIG` is set.
3. User environment variables.
4. System environment variables.
5. The Perforce user registry (set by `p4 set`).
6. The Perforce system registry (set by `p4 set -s`).

As of Release 99.1/10994, when a Perforce service (`p4s`) starts up, it reads its configuration parameters from the environment according to the following precedence:

1. Windows service parameters (set by `p4 set -S servicename`) have the highest precedence.
2. System environment variables.
3. The Perforce system registry (set by `p4 set -s`).

User environment variables can be set with any of the following:

- The MS-DOS `set` command.
- The `AUTOEXEC.BAT` file.
- **Control Panel>System>Environment>User Variables**

System environment variables can be set with:

- **Control Panel>System>Environment>System Variables**

Resolving Windows-related instabilities

There are many large sites running Perforce on Windows without incident. There are also sites in which Perforce service or server installation appears to be unstable; the server dies mysteriously, the service can't be started, and in extreme cases the system crashes. In most of these cases, this is an indication of recent changes to the machine or a corrupted operating system.

While not all Perforce failures are caused by OS-level problems, a number of symptoms may indicate the OS is at fault. Examples include: the system crashing, the Perforce server exiting without any error in its log and without Windows indicating that the server crashed, or the Perforce NT service not starting properly.

Many of these problems may be resolved by installing Service Packs. A machine running NT 4.0 with Service Pack 6 and a plain VGA video driver is very stable.

In particular, Perforce support personnel have also found that many sites have been able to reduce or eliminate OS-level instability problems by re-installing NT Service Pack 6;

consequently, we recommend that SP6 be installed on any NT system running the Perforce server or service.

In some cases, installing third-party software *after* installing a Service Pack can overwrite critical files installed by the service pack; reinstalling your most-recently installed service pack will often correct these problems. If you've installed another application after your last service pack, and server stability appears affected since the installation, consider reinstalling the service pack.

As a last resort, it may pay to install Perforce on another system to see if the same failures occur, or even to reinstall the OS and Perforce on the faulty system.

Users having trouble with P4EDITOR or P4DIFF

Your Windows users may experience difficulties using the Perforce Command-Line Client (`p4.exe`) if they use the `P4EDITOR` or `P4DIFF` environment variables.

The reason for this is that Perforce clients sometimes use the DOS shell (`cmd.exe`) to start programs such as user-specified editors or diff utilities. Unfortunately, the DOS shell knows that when a Windows command is run (such as a GUI-based editor like `notepad.exe`), it doesn't have to wait for the command to complete before terminating. When this happens, the Perforce client then mistakenly believes that the command has finished, and attempts to continue processing, often deleting the temporary files that the editor or diff utility had been using, leading to errors about temporary files not being found, or other strange behavior.

You can get around this problem in two ways:

- Unset the environment variable `SHELL`. Perforce clients under Windows only use `cmd.exe` when `SHELL` is set, otherwise they call `spawn()` and wait for the Windows programs to complete.
- Set the `P4EDITOR` or `P4DIFF` variable to the name of a DOS batch file, whose contents are the command:

```
start /wait program %1 %2
```

where *program* is the name of the editor or diff utility you wish to invoke. The `/wait` flag instructs the system to wait for the editor or diff utility to terminate, and the Perforce client will then behave properly.


Some Windows editors (most notably, Wordpad) do not behave properly, even when instructed to wait. There is presently no workaround for such programs.

Chapter 9 **Perforce Proxy**

Perforce is built to handle distributed development in a wide range of network topologies. Where bandwidth to remote sites is limited, P4P, the Perforce Proxy, improves performance by mediating between Perforce clients and servers to cache frequently transmitted file revisions. By intercepting requests for cached file revisions, P4P reduces demand on the Perforce server and network.

To improve performance obtained by multiple Perforce clients accessing a central Perforce server across a WAN, configure P4P on the side of the network close to the clients and configure the clients to access P4P, and then configure P4P to access the central Perforce server. (On a LAN, you can also obtain performance improvements by setting up proxies to divert workload from the central server's CPU and disks.)

The following diagram illustrates a typical P4P configuration:



In this configuration, file revisions requested by users at a remote development site are fetched first from a central Perforce server (`p4d` running on `central`) and transferred over a relatively slow WAN. Subsequent requests for that same revision, however, are delivered from the Perforce Proxy, (`p4p` running on `outpost`), over the remote development site's LAN, reducing both network traffic across the WAN and CPU load on the central server.

System Requirements

To use Perforce Proxy, you must have:

- A Perforce server at Release 2002.2 or higher
- Sufficient disk space on the proxy host to store a cache of file revisions

Installing P4P

UNIX

To install P4P on UNIX, do the following:

1. Download the `p4p` executable to the machine on which you wish to run the proxy.
2. Select a directory on this machine (`P4PCACHE`) in which to cache file revisions.
3. Select a port (`P4PORT`) on which `p4p` will listen for requests from Perforce client programs.
4. Select the target Perforce server (`P4TARGET`) for which this proxy will cache.

Windows

Install P4P from the Windows installer's custom/administrator installation dialog.

Running P4P

To run P4P, invoke the `p4p` executable, configuring it with environment variables or command-line flags. Flags you specify on the command line override environment variable settings.

For example, the following command line starts a proxy that communicates with a central Perforce server located on a host named `central`, listening on port 1666.

```
p4p -p 1999 -t central:1666 -r /var/proxyroot
```

To use the proxy, Perforce client programs connect to P4P on port 1999 on the machine where the proxy runs. P4P file revisions are stored under a directory named `/var/proxyroot`.

Running as a Windows service

To run P4P as a Windows service, specify the `-s` flag when you invoke `p4p.exe`, or rename the P4P executable to `p4ps.exe`.

P4P flags

The following command-line flags specific to the proxy are supported.

Flag	Meaning
-d	Run as daemon - fork first, then run (UNIX only)
-f	Do not fork - run as a single-threaded server (UNIX only)
-i	Run for <code>inetd</code> (socket on <code>stdin/stdout</code> - UNIX only)
-q	Run quietly; suppress startup messages
-s	Run as an NT service (Windows only)
	Running <code>p4p.exe -s</code> is equivalent to invoking <code>p4ps.exe</code>
-c	Do not compress files transmitted from the Perforce server to P4P (This option reduces CPU load on the central server at the expense of slightly higher bandwidth consumption)

The following general options are supported.

Flag	Meaning
-h or -?	Display a help message
-p <i>port</i>	Specify the port on which P4P will listen for requests from Perforce client programs. Default is <code>P4PORT</code> , or <code>1666</code> if <code>P4PORT</code> is not set.
-r <i>root</i>	Specify the directory where revisions are cached. Default is <code>P4PCACHE</code> , or the directory from which <code>p4p</code> is started if <code>P4PCACHE</code> is not set.
-t <i>port</i>	Specify the port of the target Perforce server (that is, the Perforce server for which P4P acts as a proxy). Default is <code>P4TARGET</code> or <code>perforce:1666</code> if <code>P4TARGET</code> is not set.
-v <i>level</i>	Specifies server trace level. Debug messages are stored in the log file. Default is <code>P4DEBUG</code> , or <code>none</code> if <code>P4DEBUG</code> is not set.
-L <i>logfile</i>	Specify the location of the log file. Default is <code>P4LOG</code> , or the directory from which <code>p4p</code> is started if <code>P4LOG</code> is not set.
-V	Display the version of the Perforce Proxy.

Administering P4P

No backups required

You never need to back up the P4P cache directory.

If necessary, P4P reconstructs the cache based on Perforce server metadata.

Managing disk space consumption

P4P caches file revisions in its cache directory. These revisions accumulate until you delete them. P4P does not delete its cached files or otherwise manage its consumption of disk space.

Warning! If you do not delete cached files, you will eventually run out of disk space. To recover disk space, remove files under the proxy's root. It is safe to delete the proxy's cached files while the proxy is running.

Determining if your Perforce client is using the proxy

If your Perforce client is using the proxy, the proxy's version information appears in the output of `p4 info`.

For example, if a Perforce server is running on `central:1666` and you direct your Perforce client to a Perforce Proxy running on `outpost:1999`, the output of `p4 info` resembles the following:

```
$ export P4PORT=outpost:1999
$ p4 info
User name: p4adm
Client name: admin-temp
Client host: remotesite22
Client root: /home/p4adm/tmp
Current directory: /home/p4adm/tmp
Client address: 192.168.0.123:55768
Server address: central:1666
Server root: /src/p4root
Server date: 2002/10/14 15:03:05 -0700 PDT
Server version: P4D/FREEBSD4/main/36609 (2002/09/30)
Proxy version: P4P/SOLARIS26/main/36884 (2002/10/14)
Server license: P4 Admin <p4adm> 20 users (expires 2003/02/01)
```

P4P and protections

To apply the IP address of a Perforce Proxy user's workstation against the protections table, prepend the string `proxy-` to the workstation's IP address.

For instance, consider an organization with a remote development site with workstations on a subnet of `192.168.10.0/24`. The organization also has a central office where local development takes place; the central office exists on the `10.0.0.0/8` subnet. A Perforce Server resides on the `10.0.0.0/8` subnet, and a Perforce Proxy resides on the `192.168.10.0/24` subnet. Users at the remote site belong to the group `remotedev`, and may occasionally visit the central office.

To ensure that members of the `remotedev` group use the proxy while working at the remote site, but do not use the proxy when visiting the local site, add the following lines to your protections table:

<code>list</code>	<code>group</code>	<code>remotedev</code>	<code>192.168.10.*</code>	<code>-//...</code>
<code>write</code>	<code>group</code>	<code>remotedev</code>	<code>proxy-192.168.10.*</code>	<code>//...</code>
<code>list</code>	<code>group</code>	<code>remotedev</code>	<code>proxy-10*</code>	<code>-//...</code>
<code>write</code>	<code>group</code>	<code>remotedev</code>	<code>10.*</code>	<code>//...</code>

The first line denies list access to all users in the `remotedev` group if they attempt to access Perforce without using the proxy from their workstations in the `192.168.10.*` subnet. The second line grants write access all users in `remotedev` if they are using a Perforce Proxy server and are working from the `192.168.10.*` subnet. Users of workstations at the remote site must use the proxy.

Similarly, the third and fourth lines deny list access to `remotedev` users when they attempt to use the proxy from workstations on the central office's subnet (`10.0.0.0/8`), but grant write access to `remotedev` users who access the Perforce server directly from workstations on the central office's subnet. When visiting the local site, users from the `remotedev` group must access the Perforce server directly.

Determining if specific files are being delivered from the proxy

Use the `-Zproxyverbose` flag with `p4` to display messages indicating whether file revisions are coming from the proxy (`p4p`) or the central server (`p4d`).

For instance:

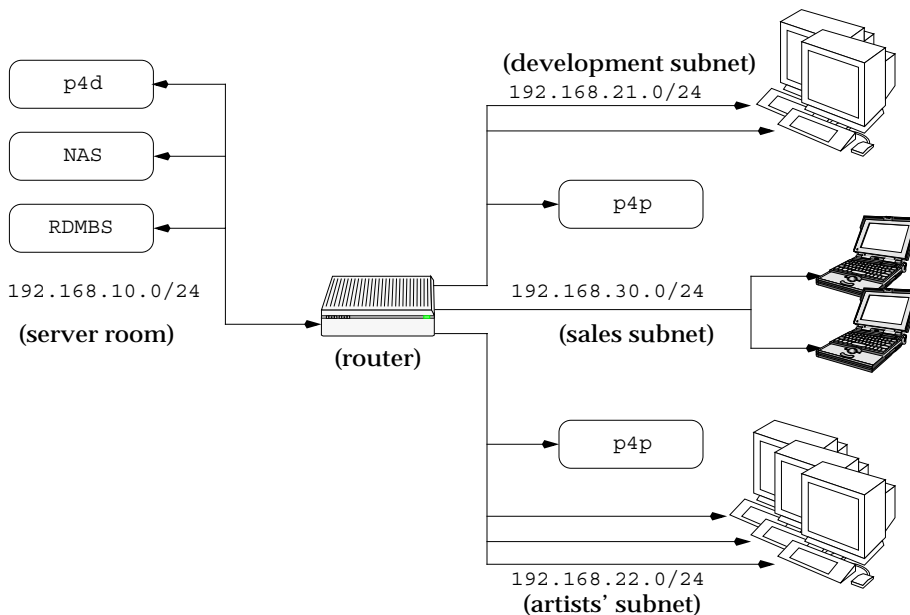
```
$ p4 -Zproxyverbose sync noncached.txt
//depot/main/noncached.txt - refreshing /home/p4adm/tmp/noncached.txt
$ p4 -Zproxyverbose sync cached.txt
//depot/main/cached.txt - refreshing /home/p4adm/tmp/cached.txt
File /home/p4adm/tmp/cached.txt delivered from proxy server
```

Maximizing performance improvement

Network topologies versus P4P

If network bandwidth on the same subnet as the central Perforce server is nearly saturated, deploying proxy servers on the same subnet will not likely result in a performance improvement. Instead, deploy the proxy servers on the other side of a router so that the traffic from the clients to the proxy server is isolated to a subnet separate from the subnet containing the central Perforce server.

For example:



Deploying an additional proxy server on a subnet when network bandwidth on the subnet is nearly saturated will not likely result in a performance improvement. Instead, split the subnet into multiple subnets and deploy a proxy server in each resulting subnet.

In the illustrated configuration, a server room houses a company's Perforce server (p4d), a network storage device (NAS), and a database server (RDBMS). The server room's network segment is saturated due heavy loads placed on it by a sales force constantly querying a database for live updates, and the developers and graphic artists frequently accessing large files through the Perforce server.

By deploying two instances of Perforce Proxy, one on the developers' subnet, and one on the graphic artists' subnet, all three groups benefit from improved performance due to decreased use on the server room's network segment.

Pre-loading the cache directory for optimal initial performance

P4P stores file revisions only when one of its clients requests them. That is, file revisions are not “prefetched”. The performance gains from P4P only occur after file revisions are cached.

After starting P4P, you can effectively prefetch the cache directory by creating a client workspace and syncing it to the head revision. All other clients that subsequently connect to the proxy immediately obtain the performance improvements provided by P4P.

For instance, a development site located in Asia with a P4P server targeting a Perforce server in North America can preload its cache directory by using an automated job that runs a `p4 sync` against the entire Perforce depot after most work at the North American site had been completed, but before its own developers arrived for work.

Distributing disk space consumption

P4P stores revisions as if there is only one depot tree. If this approach stores too much file data onto one filesystem, you can use symbolic links to spread the revisions across multiple filesystems.

For instance, if the P4P cache root is `/disk1/proxy` and the Perforce server it supports has two depots named `//depot` and `//released`, you can split data across disks, storing `//depot` on `disk1` and `//released` on `disk2` as follows:

```
mkdir /disk2/proxy/released
cd /disk1/proxy
ln -s /disk2/proxy/released released
```

The symbolic link means that when P4P attempts to cache files in the `//released` depot to `/disk1/proxy/released`, the files are stored on `/disk2/proxy/released`.

Reducing server CPU usage by disabling file compression

By default, P4P compresses communication with the central Perforce server, imposing additional overhead on the server.

To disable compression, specify the `-c` option when you invoke `p4p`. This option is particularly effective if you have excess network and disk capacity and are storing large numbers of binary file revisions in the depot, because the proxy (rather than the server) will decompress the binary files from its cache before sending them to Perforce clients.

Appendix A **Perforce Server (p4d)**

Reference

Synopsis

Invoke the Perforce server or perform checkpoint/journaling (system administration) tasks.

Syntax

```
p4d [ options ]
p4d.exe [ options ]
p4s.exe [ options ]
p4d -j [ -z ] [ args ... ]
```

Description

The first three forms of the command invoke the Perforce background process (“Perforce server”). The fourth form of the command is used for system administration tasks.

On UNIX and MacOS X, the executable is `p4d`.

On Windows, the executable is `p4d.exe` (running as a server) or `p4s.exe` (running as a service).

Exit Status

After successful startup, `p4d` does not normally exit. It merely outputs the startup message

```
Perforce server starting...
```

and runs in the background.

On failed startup, `p4d` returns a nonzero error code.

Also, if invoked with any of the `-j` checkpointing and/or journalling flags, `p4d` exits with a nonzero error code if any error occurs.

Options

Flag	Meaning
-d	Run as a daemon (in the background)
-f	Run as a single-threaded (non-forking) process
-i	Run from <code>inetd</code> on UNIX
-q	Run quietly (no startup messages)

Flag	Meaning
-s	Run <code>p4d.exe</code> as an NT service (equivalent to running <code>p4s.exe</code>)
-xu	Run database upgrades and exit.
-xi	Irreversibly reconfigure the Perforce server (and its metadata) to operate in internationalized mode. Do not use this flag unless you know you require internationalized mode. See the <i>Release Notes</i> for details.
-jc [prefix]	Journal-create; checkpoint and save/truncate journal.
-jd [file]	Journal-checkpoint; create checkpoint without saving/truncating journal.
-jj [prefix]	Journal-only; save and truncate journal without checkpointing.
-jr file	Journal-restore; restore metadata from a checkpoint and/or journal file.
-z	Compress (in <code>gzip</code> format) checkpoints and journals.
-h, -?	Print help message.
-V	Print server version.
-J journal	Specify a journal file. Overrides <code>P4JOURNAL</code> setting. Default is <code>journal</code> .
-L log	Specify a log file. Overrides <code>P4LOG</code> setting. Default is <code>stderr</code> .
-p port	Specify a port to listen to. Overrides <code>P4PORT</code> . Default <code>1666</code> .
-r root	Specify the server root directory. Overrides <code>P4ROOT</code> . Default is current working directory.
-v debuglevel	Set server trace flags. Overrides value <code>P4DEBUG</code> setting. Default is <code>null</code> .

Usage Notes

- On all systems, journaling is enabled by default. If `P4JOURNAL` is unset when a server starts, the default location for the journal is `$P4ROOT/journal`. If you wish to manually disable journaling, you must explicitly set `P4JOURNAL` to `off`.
- Take checkpoints and truncate the journal often, preferably as part of your nightly backup process.
- Checkpointing and journaling preserve only your Perforce metadata (data *about* your stored files). The stored files themselves (the files containing your source code) reside under `P4ROOT` and must be also be backed up as part of your regular backup procedure.

- If your users are using triggers, don't use the `-f` (non-forking mode) flag; the Perforce server needs to be able to spawn copies of itself ("fork") in order to run trigger scripts.
- After a hardware failure, the flags required for restoring your metadata from your checkpoint and journal files may vary, depending on whether or not data was corrupted.
- Because restorations from backups involving loss of files under `P4ROOT` often require the journal file, we strongly recommend that the journal file reside on a separate filesystem from `P4ROOT`. This way, in the event of corruption the filesystem containing `P4ROOT`, the journal is likely to remain accessible.
- The database upgrade flag (`-xu`) may require considerable disk space. See the *Release Notes* and the section "Important Notes for 2001.1 and later" on page 16 if upgrading to 2001.1 or later from a 2000.2 or earlier server.

Related Commands

To start the server, listening to port 1999, with journaling enabled and written to <code>journalfile</code> .	<code>p4d -d -p 1999 -J /opt/p4d/journalfile</code>
To checkpoint a server with a non-default journal file, the <code>-J</code> argument (or the environment variable <code>P4JOURNAL</code>) must match the journal file specified when the server was started.	Checkpoint with: <code>p4d -J /p4d/jfile -jc</code> or <code>P4JOURNAL=/p4d/jfile ; export P4JOURNAL</code> <code>p4d -jc</code>
To create a compressed checkpoint from a server with files in directory <code>P4ROOT</code>	<code>p4d -r \$P4ROOT -z -jc</code>
To create a compressed checkpoint with a user-specified prefix of " <code>ckp</code> " from a server with files in directory <code>P4ROOT</code>	<code>p4d -r \$P4ROOT -z -jc ckp</code>
To restore metadata from a checkpoint named <code>checkpoint.3</code> for a server with root directory <code>P4ROOT</code>	<code>p4d -r \$P4ROOT -jr checkpoint.3</code>
To restore metadata from a compressed checkpoint named <code>checkpoint.3.gz</code> for a server with root directory <code>P4ROOT</code>	<code>p4d -r \$P4ROOT -z -jr checkpoint.3.gz</code>

Index

A

- access level
 - and protections 62
- access levels 62
- admin access level 39, 63
- administrator
 - force flag 45
 - privilege required 107
- administrators
 - and job specifications 71
- allocating disk space 20
- AppleSingle 30
- arguments
 - passing to triggers 86
- automated checkpoints 27
- automating Perforce 40

B

- backing up 31
- backup
 - procedures 31
 - recovery procedures 33
- branches
 - namespace 59
- buffering
 - of input/output in scripts 90

C

- can 106
- case-sensitivity
 - and cross-platform development 23
 - UNIX and Windows 23, 50
- change review 87
- changelist numbers
 - highest possible 89
 - pending vs. submitted changelists 90
- changelists
 - deleting 43
 - editing 43
 - validating 83
- checkpoint

- as part of backup script 31
- creating 26
- creation of, automating 27
- defined 26
- ensuring completion of 32
- failed 27
- introduced 25
- managing disk space 20
- when to call support 27

checkpoints

- creating with `p4 admin` 27, 31

client

- and port 13

clients

- namespace 59

commands

- forcing 45

counter

- limits 89

CPU

- and performance 93

CR/LF conversion 55

creating checkpoints 26

creating users 39

creation of users

- preventing 40

cross-platform development

- and case sensitivity 23

D

daemon

- change review 87

daemons 83, 87–88

- changelist numbers 90

- creating 88

database files 53

- defined 25

- where stored 25

db.* files 25

debugging

- with server tracing 52
- defect tracking
 - integrating with Perforce 80
- deleting
 - changelists 43
 - depots 60
 - files, permanently 42
 - user groups 67
- deleting users 41
- depot
 - and Mac file formats 30
- depot files
 - see *versioned files* 30
- depots
 - defined 25
 - defining 58
 - deleting 60
 - listing 59
 - local 58
 - mapping field 59
 - multiple 56
 - namespace 59
 - remote 56, 60
 - remote, defining 58
- disabling journaling 29
- disk
 - performance 92
 - sizing 92
- disk space
 - allocating 20
 - and server trace flags 52
 - freeing up 42
 - required for upgrade 16
- distributed development 57
- DNS
 - and performance 94, 95
- drives
 - and `db.*` and journal file 19

E

- editing
 - changelists 43
- editor
 - Wordpad, limitation 110

- environment variables
 - P4PCACHE 112, 113
 - P4PORT 112
 - P4TARGET 112, 113
- error logging 23
- error messages
 - and `p4 verify` 43
- example
 - specifying journal files 29
- exclusionary mappings
 - and protections 65

F

- fields
 - of job template 72
- file formats
 - AppleSingle 30
- file names
 - mapping to file types 44
- file specification
 - and protections 62
- file types
 - mapping to file names 44
- files
 - access to, limiting 65
 - database 25
 - left open by users, reverting 41
 - matching Perforce file types to file names 44
 - permanent deletion of 42
 - subscribing to 67
 - verification of 43
 - versioned 25
- filesystems
 - and performance 92
 - large 21
 - NFS-mounted, caveats 21, 92
- firewall
 - defined 47
 - running Perforce through 46
- flags
 - and Perforce Proxy 113
 - f to force 45
 - server, listed 119

G

groups

- and protections 62, 66
- and subgroups 66
- deleting 67
- editing 66
- of users 66

H

hostname

- changing your server's 55

hosts

- and protections 62

hosts file

- on Windows and UNIX 95

I

-i

- and `inetd` 50
- automating job submissions 80
- automating user creation 40

`inetd` 50, 119

installing

- license file 19
- on network drives 22
- on NFS filesystems 21, 92
- on UNIX 11
- on Windows 15
- on Windows network drives 107
- Perforce Proxy 112

IP address

- changing your server's 55
- servers and `P4PORT` 49

IP forwarding

- and `ssh` 48

J

job fields

- data types 74

job specification 71–77

- and administrators 71
- and comments 76
- default format 71
- defining fields 73
- extended example 78
- warnings 77

job template

- default 71
- fields of 72
- viewing 72

jobs

- comments in 76
- other defect tracking systems 80

journal

- defined 28
- introduced 25
- managing size of 20
- where to store 20

journal file

- specifying 120
- store on separate drive 19

journaling

- disabling 29

L

label

- namespace 59

license 19

licensing information 18

limitations

- on information passed to triggers 86
- Wordpad 110

`list` access level 62

listing

- depot names 59

local depots 58

log file

- specifying 120

M

Mac

- and file formats 30

Macintosh

- OS X 11

mappings

- and depots 59

maxresults

- and multiple groups 99
- and P4Win 99
- and performance 98
- use of 98

- maxscanresults
 - and performance 98
 - use of 98
- maxscanrows
 - and multiple groups 99
 - and P4Win 99
- MD5 signatures 43
- memory
 - and performance 91
 - requirements 91
- metadata
 - see database files 25, 53
- moving servers 53
 - across architectures 54
 - between Windows and UNIX 55
 - new hostname 55
 - new IP address 55
 - same architecture 53
- multiple depots 56
 - and users 60
- multiple triggers
 - on one file 85
- N**
- naming
 - depots 59
- network
 - and performance 93, 94
 - Perforce Proxy configuration 111
 - problems, diagnosing 94
- network drives
 - and triggers 86
 - and Windows 22
- network interface
 - directing server to listen to specific 49
- NFS
 - and installation 21, 92
- non-forking 119
- O**
- obliterating files 42
- open access level 63
- operating systems
 - and large filesystem support 21
- OS X

- and UNIX 11
- P**
- p4 admin
 - and Windows 16, 106
 - creating checkpoints 27, 31
 - stopping server with 14, 34, 35
- p4 jobspec
 - warnings 77
- p4 set -s
 - setting variables for Windows services 109
- p4 triggers
 - form 84
- p4 typemap 44
- p4 verify 43
 - use of 31
- p4d
 - f option, and triggers 86
 - flags, listed 119
 - security 22, 86
 - specifying journal file 120
 - specifying log file 120
 - specifying port 120
 - specifying server root 120
 - specifying trace flags 120
- p4d.exe 15
- P4DEBUG 120
 - and proxy server 113
- P4JOURNAL 120
- P4LOG 120
 - and proxy server 113
- P4P 57
- P4P
 - see Perforce Proxy 112
- P4P
 - see Perforce Proxy 111
- P4PCACHE 112, 113
- P4PORT
 - and client 13
 - and proxy server 113
 - and server 13, 120
 - IP addresses and your server 49
 - Perforce Proxy 112

- P4ROOT 12, 120
- p4s.exe 15
- P4TARGET 112, 113
- passwords
 - setting 20, 39
- PDF files
 - and p4 typemap 44
- Perforce
 - uninstalling 105
- Perforce clients
 - and P4PORT 13
- Perforce Proxy 57, 111
 - options 113
- Perforce server
 - and P4PORT 13
 - and triggers 84
 - and Windows network drives 22
 - installing under NFS 21, 92
 - moving to another machine 53
 - running from inetd 50
 - UNIX UPGRADE 17
 - upgrading 16
 - upgrading under Windows 18
 - verifying 43
 - vs. service 15
- Perforce service
 - vs. server 15
- perforce.exe 15
- performance
 - and memory 91
 - and scripts 100
 - and wildcards under Windows 95
 - CPU 93
 - network 93, 111
 - preventing server swamp 96
 - slow, diagnosing 94
- permissions
 - see protections 64
- port
 - for client 13
 - for server 13
 - specifying 120
- ports
 - running out of TCP/IP 93
- pre-submit triggers
 - see *triggers* 83
- privileges
 - administrator 107
- protections 61–68
 - algorithm for applying 67
 - and commands 68
 - and groups 66
 - and performance 97
 - and superusers 61
 - commands affected by 68
 - default 64
 - exclusionary 65
 - multiple 64
 - schemes for defining 63
- protections table 61
- proxy 111
 - and remote development 57
- python 87
- R**
- RAM
 - and performance 91
- read access level 62
- recovery
 - procedures 33
- remote depots 56
 - and virtual users 60
 - defining 58
- resetting passwords 39
- review access level 63
- revision range
 - and obliterate 42
- rich text
 - and p4 typemap 44
- root
 - must not run p4d 22, 86
- S**
- scripting
 - buffering standard in/output 90
 - guidelines for efficient 100
 - with -i 40
- scripting Perforce 83–88

- secure shell 47
- security
 - and passwords 20
 - p4d must have minimal privileges 22, 86
 - preventing user impersonation 19
- server
 - and triggers 84
 - backing up 31
 - forking, and triggers 86
 - license file 19
 - licensing 18
 - migrating 53
 - port 13
 - proxy 111
 - recovery 33
 - root, specifying 120
 - running from `inetd` 50
 - running in background 119
 - running single-threaded 119
 - specifying journal file 120
 - specifying log file 120
 - specifying port 120
 - stopping on Windows 106
 - stopping with `p4 admin` 14, 34, 35
 - trace flags 52
 - upgrading 16
 - verifying 43
 - vs. service 15
 - Windows 15
- server flags
 - listed 119
- server root
 - and `P4ROOT` 12
 - creating 12
 - defined 12
 - specifying 120
- server upgrade
 - UNIX 17
 - Windows 18
- setting passwords 20, 39
- single-threaded 119
- ssh 47
- standard input/output
 - buffering 90
- stopping server
 - on Windows 106
 - with `p4 admin` 14, 34, 35
- subgroups
 - and groups 66
- submission, validating 83
- super access level 39, 63
- superuser
 - and triggers 84
 - force flag 45
 - Perforce, defining 19
- superusers
 - and protections 61
- symbolic links
 - and disk space 20
- T**
- TCP/IP
 - and port number 13
 - running out of ports 93
- technical support
 - when to call 27
- template
 - job, default 71
- trace flags
 - specifying 120
- triggers 83–86
 - and Windows 86
 - arguments, passing to 86
 - defined 83
 - form 84
 - multiple, on one file 85
 - naming 84
 - ordering 85
 - script, specifying arguments to 84
 - security and `p4d` 22, 86
 - server must be able to fork 86
 - single, on multiple filespecs 86
- troubleshooting
 - slow response times 94
- type mapping 44
- U**
- umask (1) 12

uninstalling Perforce 105

UNIX

 /etc/hosts file 95

 and case-sensitivity 51

 upgrading a server 17

upgrading

 server 16

users

 access control by groups 66

 and multiple depots 60

 and protections 62

 creating 39

 deleting 41

 files, limiting access to 65

 nonexistent 41

 preventing creation of 40

 preventing impersonation of 19

 resetting passwords 39

 virtual, and remote depots 60

V

validation, of changelists 83

variables

 setting for a Windows service 109

verifying server integrity 43

version information

 and Perforce Proxy 113

 clients and servers 19

versioned files 53

 defined 25

 format and location of 30

 introduced 25

 where stored 25

view

 scope of, and performance 96

W

warnings

 and job specifications 77

 database changes on upgrade 16, 17

 disk space and Perforce Proxy 114

 disk space and upgrade 16

 obliterating files 42

 security and p4d 22, 86

wildcards

 and protections 62

 and Windows performance 95

Windows

 and case-sensitivity 23, 52

 and p4 admin 16

 and server upgrade 18

 hosts file 95

 installer 15

 installing on network drive 22, 107

 server 15

 service, setting variables in 109

 stopping server 106

 triggers and network drives 86

Wordpad

 limitation 110

wrappers 83

write access level 63

