

PERFORCE

Server Deployment Package for Perforce Helix

User Guide (for Unix)

v18.2

Perforce Software, Inc.

Preface

This guide tells you how to set up a new production ready Perforce Helix Server installation using the Server Deployment Package (SDP). Recommendations for optimal system maintenance and performance are included as well. The SDP follows **best practices** for Perforce server configuration and administration. The SDP consists of standard configuration settings, scripts, and tools, which provide several key features:

- A volume layout designed for maximum data integrity and server performance.
- Automated offline checkpointing and backup procedures for server metadata (avoiding server downtime during the checkpointing process which can take hours for large installations)
- Reliable configuration for replication to other servers (particularly for HA/DR).

This guide assumes some familiarity with Perforce, and does not duplicate the basic information in the Perforce user documentation. For basic information on Perforce, consult *Introducing Perforce*. For system administrators, the *Perforce System Administrator's Guide* is essential reading. All documentation is available from the Perforce web site at <http://www.perforce.com>.

Please Give Us Feedback

Perforce welcomes feedback from our users. Please send any suggestions for improving this document or the SDP to consulting@perforce.com.

Table of Contents

1	Overview	6
2	Configuring the Perforce Server	7
2.1	Volume Layout and Hardware	7
2.2	Memory and CPU	8
2.3	General SDP Usage	9
2.3.1	Unix/Linux	9
2.3.2	Monitoring SDP activities	10
3	Installing the Perforce Server and the SDP	11
3.1	Installing on Unix/Linux Machines	11
3.1.1	Initial setup	11
3.1.2	Setting up init.d/systemd auto-start scripts	13
3.1.3	Upgrading an existing SDP installation	14
3.1.4	Configuration script	14
3.1.5	Use of SSL	15
3.1.6	Starting/Stopping Perforce Server Products	16
3.1.7	Archiving configuration files	18
3.2	Configuring protections, file types, monitoring and security	18
3.3	Other server configurables	18
4	Backup, Replication, and Recovery	20
4.1.1	Typical Backup Procedure	20
4.2	Full One-Way Replication	21
4.2.1	Replication Setup	21
4.3	Recovery Procedures	24
4.3.1	Recovering a master server from a checkpoint and journal(s)	24
4.3.2	Recovering a replica from a checkpoint	25
4.3.3	Recovering from a tape backup	25
4.3.4	Failover to a replicated standby machine	26
5	Server Maintenance	27
5.1	Server upgrades	27
5.2	Database Modifications	27
5.3	Maximizing Server Performance	27
5.4	Optimizing the database files	27
5.5	Proactive Performance Maintenance	28
5.5.1	Limiting large requests	28
5.5.2	Offloading remote syncs	28
6	Tools and Scripts	29
6.1	Core Scripts	29
6.1.1	p4_vars	29
6.1.2	p4_<instance>.vars	29
6.1.3	p4master_run	29
6.1.4	recreate_offline_db	29

6.1.5	live_checkpoint	29
6.1.6	daily_checkpoint	29
6.1.7	p4verify	30
6.1.8	p4review.py	30
6.1.9	p4login	30
6.1.10	p4d_instance_init	30
6.2	More Server Scripts	30
6.2.1	upgrade.sh	30
6.2.2	p4.crontab	31
6.3	Other Files	32
7	Appendix A – Directory Structure Configuration Script for Linux/Unix	33
8	Appendix B – Frequently Asked Questions/Troubleshooting	35
8.1.1	Journal out of sequence	35
8.1.2	Unexpected end of file in replica daily sync	35

1 Overview

The SDP has four main components:

- Hardware and storage layout recommendations for Perforce with a consistent standard directory structure which accommodates multiple instances on a single server machine where required
- Scripts to automate offline checkpoints and other critical maintenance activities, including running as a non-root user and ensuring case sensitivity settings are maintained
- Scripts to aid the setup and management of replication of Perforce instances to another server (for HA/DR – high availability/ disaster recovery).
- Scripts to assist with user account maintenance and other routine administration tasks.

Each of these components is covered in detail in this guide.

The SDP should be versioned in a depot (e.g. //perforce) as part of the installation process.

The directory structure of the SDP is shown below in Figure 1: SDP Package Directory Structure. This includes all SDP files, including documentation and maintenance scripts. A subset of these files are deployed to server machines during the installation process.

```
sdp
  doc
  Maintenance (Admin scripts)
  Server (Core SDP Files)
    setup (typemap, configure, etc)
    test (automated test scripts)
  Unix
    setup
    p4
    1
      bin
    common
      bin (Backup scripts, etc)
      triggers (Example triggers)
    config
    etc
      cron.d
      init.d
    lib
    test
```

Figure 1: SDP Package Directory Structure

2 Configuring the Perforce Server

This chapter tells you how to configure a Perforce server machine and an instance of the Perforce Server. These topics are covered more fully in the System Administrator's Guide and in the Knowledge Base; this chapter covers the details most relevant to the SDP.

The SDP can be installed on multiple server machines, and each server machine can host one or more Perforce server instances. (In this guide, the term *server* refers to a Perforce server instance unless otherwise specified.) Each server instance is assigned a number (or a name). This guide uses instance number 1 in the example commands and procedures. Other instance numbers/names can be substituted as required.

Optionally, instances can be given a short tag name, such as 'abc', rather than a number. Manual configuration is required in a couple of places to use tag names rather than the default numeric values, e.g. for the setting of appropriate P4PORT integer values (see `/p4/common/config/p4_1.vars`).

This chapter also describes the general usage of SDP scripts and tools.

2.1 Volume Layout and Hardware

To ensure maximum data integrity and performance, use three or four different physical volumes for each server instance. Three volumes can be used for all instances hosted on one server machine, but using three or four volumes per instance reduces the chance of hardware failure affecting more than one instance. The "hx" prefix is used to indicate Helix volumes (your own naming conventions/standards can be used instead).

- **Perforce metadata (database files), volumes 1 & 2:** Use the fastest volume possible, ideally SSD or RAID 1+0 on a dedicated controller with the maximum cache available on it. These volumes default to `/hxmetadata1` and `/hxmetadata2`. Having two means we can swap online and offline database just by switching a couple of links. It is fine to have these both pointing to the same physical volume.
- **Journals and logs:** Use a fast volume, ideally SSD or RAID 1+0 on its own controller with the standard amount of cache on it. This volume is normally called `/hxlogs` and should usually be backed up. If a separate logs volume is not available, put the logs on the `hxmetadata1` volume.
- **Depot data, archive files, scripts, and checkpoints:** Use a large volume, with RAID 5 on its own controller with a standard amount of cache or a SAN or NAS volume. This volume is the only volume that must be backed up. The backup scripts place the metadata snapshots on this volume. This volume can be backed up to tape or another long term backup device. This volume is normally called `/hxdepots`.

For optimal performance on UNIX machines, the XFS file system is recommended, but ext4 and others are fine.

If three controllers are not available, put the `hxlogs` and `hxdepots` volumes on the same controller.

Do not run anti-virus tools or back up tools against the `hxmetadata` volume(s) or `hxlogs`

volume(s), because they can interfere with the operation of the Perforce server.

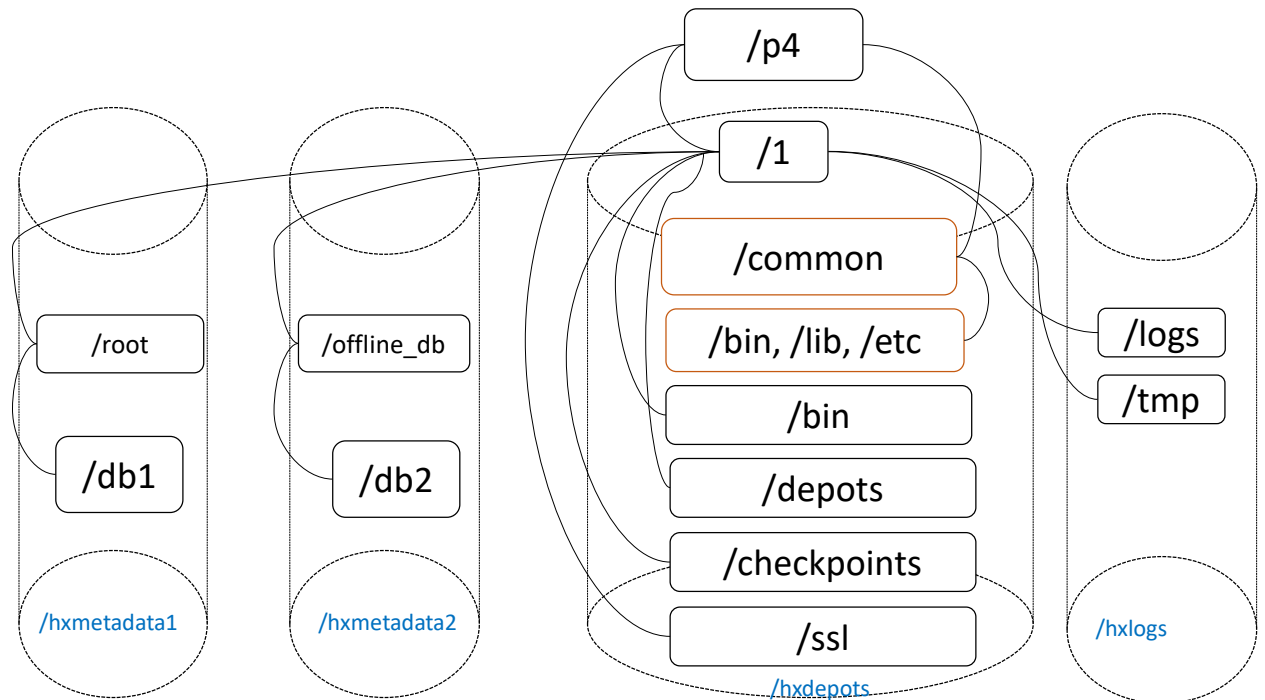


Figure 2: SDP Runtime Structure and Volume Layout



Back up everything on the `hxdepots` volume(s). Avoid backing up the `hxmetadata[1,2]` volumes directly, because doing so can interfere with the operation of a live Perforce server (due to locks on files), potentially corrupting data. The checkpoint and journal process archive the metadata onto the `hxdepots` volume. Backing up the `hxlogs` volume is usual

The SDP assumes (but does not require) the three volumes described above. On Unix/Linux platforms, the SDP `CREATES` a convenience directory containing links to the three volumes for each instance. This convenience directory is called `/p4`. The volume layout is shown in Figure 2: SDP Runtime Structure and Volume Layout, including the links that constitute the `/p4` directory for Unix/Linux platforms. (On Windows, a similar convenience directory `c:\p4` is created and `mklink` is used.) The convenience directory enables easy access to the different parts of the file system for each instance. For instance, the directory `/p4/1/root` has the database for instance 1, `/p4/1/logs` has the logs for instance 1, `/p4/1/bin` has the binaries for instance 1, and `/p4/common/bin` contains the scripts common to all instances.

Figure 2: SDP Runtime Structure and Volume Layout, viewed from the top down, displays a Perforce *application* administrator's view of the system, which shows how to navigate the directory structure to find databases, log files, and versioned files in the depots. Viewed from the bottom up, it displays a Perforce *system* administrator's view, emphasizing the physical volume where Perforce data is stored.

2.2 Memory and CPU

Make sure the server has enough memory to cache the `db.rev` database file and to prevent the server from paging during user queries. Maximum performance is obtained if the server has

enough memory to keep all of the database files in memory.

Below are some approximate guidelines for allocating memory.

- 1.5 kilobyte of RAM per file stored in the server.
- 32 MB of RAM per user.

Use the fastest processors available with the fastest available bus speed. Faster processors are typically more desirable than a greater number of cores and provide better performance since quick bursts of computational speed are more important to Perforce's performance than the number of processors. Have a minimum of two processors so that the offline checkpoint and back up processes do not interfere with your Perforce server. There are log analysis options to diagnose underperforming servers and improve things (contact support/consulting for details).

2.3 General SDP Usage

This section presents an overview of the SDP scripts and tools. Details about the specific scripts are provided in later sections.

2.3.1 Unix/Linux

Most scripts and tools reside in `/p4/common/bin`. The `/p4/instance/bin` directory contains scripts or links that are specific to that instance such as wrappers for the `p4d` executable.

Older versions of the SDP required you to always run important administrative commands using the `p4master_run` script, and specify fully qualified paths. This script loads environment information from `/p4/common/bin/p4_vars`, the central environment file of the SDP, ensuring a controlled environment. The `p4_vars` file includes instance specific environment data from `/p4/common/config/p4_instance.vars` e.g.

`/p4/common/config/p4_1.vars` The `p4master_run` script is still used when running `p4` commands against the server unless you set up your environment first by sourcing `p4_vars` with the instance as a parameter. Administrative scripts, such as `daily_backup.sh`, no longer need to be called with `p4master_run` however, they just need you to pass the instance number to them.

When invoking a Perforce command directly on the server machine, use the `p4_instance` wrapper that is located in `/p4/instance/bin`. This wrapper invokes the correct version of the `p4` client for the instance. The use of these wrappers enables easy upgrades, because the wrapper is a link to the correct version of the `p4` client. There is a similar wrapper for the `p4d` executable, called `p4d_instance`. This wrapper is important to handle case sensitivity in a consistent manner, e.g. when running a Unix server in case-insensitive mode.

Below are some usage examples for instance 1.

Example	Remarks
<code>/p4/common/bin/p4master_run 1</code> <code>/p4/1/bin/p4_1 admin stop</code>	Run <code>p4 admin stop</code> on instance 1
<code>/p4/common/bin/live_checkpoint.sh 1</code>	Take a checkpoint of the live database on instance 1
<code>/p4/common/bin/p4login 1</code>	Log in as the perforce user on instance 1.

Some maintenance scripts can be run from any client workspace, if the user has

administrative access to Perforce.

2.3.2 Monitoring SDP activities

The important SDP maintenance and backup scripts generate email notifications when they complete.

For further monitoring, you can consider options such as:

- Making the SDP log files available via a password protected HTTP server.
- Directing the SDP notification emails to an automated system that interprets the logs.

3 Installing the Perforce Server and the SDP

This chapter tells you how to install a Perforce server instance in the SDP framework. For more details about server installation, refer to the Perforce System Administrator's Guide.

Many companies use a single Perforce Server to manage their files, while others use multiple servers. The choice depends on network topology, the geographic distribution of work, and the relationships among the files being managed. If multiple servers are run, assign each instance a number and use that number as part of the name assigned to depots, to make the relationship of depots and servers obvious.

The default P4PORT setting used by the SDP is `instance666`. For example, instance **1** runs on port **1666**. Each Perforce instance uses its hostname as an identifying name; this identification is used for replicated servers. This can easily be changed in `/p4/common/bin/p4_vars`.

For any instances that are named rather than numbered, then the `/p4/common/bin/p4_vars` file must be customized to assign a numeric P4PORT value to each named instance.



To install the SDP, you must have root (super-user or administrator) access to the server machine (at least until you have created the appropriate directory structures with `mkdirs.sh` – after that you no longer need root access).

3.1 Installing on Unix/Linux Machines

To install Perforce Server and the SDP, perform the following basic steps that are discussed below:

- Set up a user account, file system, and configuration scripts.
- Run the configuration script.
- Start the server and configure the required file structure for the SDP.

3.1.1 Initial setup

Prior to installing the Perforce server, perform the following steps.

1. Create a group called `perforce` (can be a different name – see `OSGROUP` in `mkdirs.cfg`) and make it the `perforce` user's primary group. Typically this command:

```
sudo groupadd perforce
```
2. Create a user called `perforce` (can be a different name if you prefer, in which case modify the `OSUSER` entry in the `mkdirs.1.cfg` script - see item 12 below). Set the user's home directory to `/p4` on a local disk. Typically this command:

```
sudo useradd -d /p4 -s /bin/bash -m perforce -g perforce
```
3. Create or mount the server file system volumes (`/hxdepots`, `/hxmetadata1`, `/hxmetadata2`, `/hxlogs`). It is possible to have these be links to other directories/volumes if you wish. These directories should be owned by `$OSUSER:$OSGROUP` but this will be done automatically by `mkdirs.sh` (run as root).
4. Copy the SDP to the directory `/hxdepots/sdp`. We will refer to this directory as `$SDP`. Make the entire `$SDP` directory writable (`chmod -R $SDP`). These commands download directly (substitute latest release identifier from page:
<https://swarm.workshop.perforce.com/projects/perforce-software-sdp/files/downloads> - click on the version then right click on download button >

copy link address), e.g.

```
cd /hxdepots

wget https://swarm.workshop.perforce.com/projects/perforce-software-
sdp/download/downloads/sdp.Unix.2018.1.23583.p2.tgz

tar zxvf sdp.Unix.2018.1.23583.p2.tgz
```

The above creates the directory /hxdepots/sdp

- Download the appropriate p4 and p4d binaries for your release and platform from ftp.perforce.com (log in as anonymous) and place them in \$SDP/Server/Unix/p4/common/bin. **Do not** rename them to include the version number; this step is done automatically for you by the SDP installation.

```
cd $SDP/p4/common/bin

wget http://ftp.perforce.com/perforce/r18.2/bin.inux26x86_64/p4
wget http://ftp.perforce.com/perforce/r18.2/bin.inux26x86_64/p4d
wget http://ftp.perforce.com/perforce/r18.2/bin.inux26x86_64/p4broker

chmod +x p4*
```

- cd to \$SDP/Server/Unix/setup and copy `mkdirs.cfg` to an instance specific version such as `mkdirs.1.cfg` and edit it - set all of the variables in the configuration variables section for your company.

```
cp mkdirs.cfg mkdirs.instance.cfg
```

Examples:

```
cp mkdirs.cfg mkdirs.1.cfg

vi mkdirs.1.cfg

cp mkdirs.cfg mkdirs.Master.cfg

vi mkdirs.Master.cfg
```

- As the root user, cd to \$SDP/Server/Unix/setup, and run this:

```
mkdirs.sh instance
```

Examples:

```
mkdirs.sh 1

mkdirs.sh Master
```

This script configures the first Perforce Server instance. To configure additional instances, create a new instance `cfg` file and run `mkdirs.sh` again, specifying the instance number each time. For example, to configure a second and third instance, issue the following commands (having created and edited relevant `mkdirs.instance.cfg`):

```
mkdirs.sh 2

mkdirs.sh 3
```

- Put the Perforce license file for the server into /p4/1/root. Note, if you have multiple instances and have been provided with port-specific licenses by Perforce, the appropriate license file must be stored in the appropriate /p4/*instance*/root folder.
- Make the Perforce server a system service that starts and stops automatically when the machine reboots. Running `mkdirs.sh` creates a set of init scripts for various

Perforce server products in the instance-specific bin folder:

```
/p4/1/bin/p4d_1_init
/p4/1/bin/p4broker_1_init [only created if a p4broker executable
    found]
/p4/1/bin/p4p_1_init [only created if a p4p executable found]
/p4/1/bin/init/p4ftpd_1_init [move up 1 level if used]
/p4/1/bin/init/p4dtg_1_init [move up 1 level if used]
/p4/1/bin/p4web_1_init [only created if a p4web executable found]
```

Note that to create init scripts for p4p/p4broker/p4web you should download those executables to the /p4/common/bin directory before running makedirs.sh

The steps required to complete the configuration will vary depending on the Unix distribution being used.

3.1.2 Setting up init.d/systemd auto-start scripts

The following sample commands enable init scripts as system services on RedHat / CentOS (up to version 6) and SuSE (up to version 11). Run these commands as the root user:

```
cd /etc/init.d
ln -s /p4/1/bin/p4d_1_init
chkconfig --add p4d_1_init
chkconfig p4d_1_init on
```

Run the 'ln -s' and two 'chkconfig' commands for any other init scripts, besides p4d_1_init, that you wish to operate on for that instance and on the current machine, such as p4broker_1_init or p4web_1_init. Remove init scripts for any services not needed on that machine. Note that the broker and P4Web require additional configuration before those services will start.

RHEL 7, CentOS 7, SuSE 12, Ubuntu (v15.04) (and other) distributions utilize **systemd** / **systemctl** as the mechanism for controlling services, replacing the earlier init process. At present makedirs.sh does **not** generate the systemd configuration file(s) automatically, but a sample is included in the SDP distribution in (\$SDP/Server/Unix/setup/systemd), along with a README.md file that describes the configuration process, including for multiple instances. For simple installation run these commands as the root user:

```
cp $SDP/Server/Unix/setup/system/p4d_1.system /etc/systemd/system/
systemctl enable p4d_1
```

The above enables service for auto-start on boot. The following show management commands.

```
systemctl status -l p4d_1
systemctl start p4d_1
systemctl status -l p4d_1
systemctl stop p4d_1
systemctl status -l p4d_1
```

Note that if using systemctl we **do not** recommend you using `/p4/1/bin/p4d_1_init` script as then systemctl will not work correctly. Instead we recommend enabling the OSUSER account configured above to either have sudo access, or at least to directly be able to execute the systemctl for this service ([See here for pointers](#)).

Ubuntu (pre 15.04), MacOS and other Unix derivatives use different mechanisms to enable services. If your Linux distribution does not have the chkconfig or systemctl utilities, consult your distribution's documentation for information on enabling services.

3.1.3 Upgrading an existing SDP installation

If you have an earlier version of the Server Deployment Package (SDP) installed, you'll want to be aware of the new `-test` flag to the SDP setup script, `mkdirs.sh`. The following update instructions assume a simple, single-server topology.

See the instructions in the file `README.md` / `README.html` in the root of the SDP directory.

3.1.4 Configuration script

The `mkdirs.sh` script executed above resides in `$SDP/Server/Unix/setup`. It sets up the basic directory structure used by the SDP. Carefully review the config file `mkdirs.instance.cfg` for this script before running it, and adjust the values of the variables

as required. The important parameters are:

Parameter	Description
DB1	Name of the <code>hxmetadata1</code> volume (can be same as DB2)
DB2	Name of the <code>hxmetadata2</code> volume (can be same as DB1)
DD	Name of the <code>hxdepots</code> volume
LG	Name of the <code>hxlogs</code> volume
CN	Volume for <code>/p4/common</code>
SDP	Path to SDP distribution file tree
SHAREDATA	TRUE or FALSE – whether sharing the <code>/hxdepots</code> volume with a replica
ADMINUSER	P4USER value of a Perforce super user that operates SDP scripts, typically <code>perforce</code> or <code>p4admin</code> .
OSUSER	Operating system user that will run the Perforce instance, typically <code>perforce</code> .
OSGROUP	Operating system group that <code>OSUSER</code> belongs to, typically <code>perforce</code> .
CASE_SENSITIVE	Indicates if server has special case sensitivity settings
SSL_PREFIX	Set if SSL is required so either “ <code>ssl:</code> ” or blank for no SSL
P4ADMINPASS	Password to use for Perforce superuser account – can be edited later in <code>/p4/common/config/.p4password.p4_1.admin</code>
P4SERVICEPASS	Service User’s password for replication – can be edited later – same dir as above.
P4DNSNAME	Fully qualified DNS name of the Perforce master server machine

For a detailed description of this config file it is fully documented with in-file comments, or see Appendix A – Directory Structure Configuration Script for Linux/Unix.

3.1.5 Use of SSL

As documented in the comments in `mkdirs.cfg`, if you are planning to use SSL you need to set the value of `SSL_PREFIX=ssl:`

Then you need to put certificates in `/p4/ssl` after the SDP install or you can generate a self

signed certificate as follows:

Edit `/p4/ssl/config.txt` to put in the info for your company. Then run:

```
/p4/common/bin/p4master_run <instance> /p4/<instance>/p4d_<instance> -Gc
```

For example using instance 1:

```
/p4/common/bin/p4master_run 1 /p4/1/bin/p4d_1 -Gc
```

In order to validate that SSL is working correctly:

```
source /p4/common/bin/p4_vars 1
```

Check that P4TRUST is appropriately set in the output of:

```
p4 set
```

Update the P4TRUST values (answer yes when prompted):

```
p4 -p ssl:1666 trust
```

```
p4 -p ssl:localhost:1666 trust
```

Check the stored P4TRUST values:

```
p4 trust -l
```

Check not prompted for trust:

```
p4 login
```

```
p4 info
```

3.1.6 Starting/Stopping Perforce Server Products

The SDP includes templates for initialization (start/stop) scripts, “init scripts,” for a variety of Perforce server products, including:

- p4d
- p4broker
- p4p
- p4dtg
- p4ftpd
- p4web

The init scripts are named `/p4/instance/bin/svc_instance_init`.

For example, the init script for starting p4d for Instance 1 is `/p4/1/bin/p4d_1_init`. All init scripts accept at least `start`, `stop`, and `status` arguments. The perforce user can start p4d

by calling:

```
p4d_1_init start
```

And stop it by calling:

```
p4d_1_init stop
```



If you have configured systemd (CentOS 7+) then as noted above we do not recommend you using the `p4d_instance_init` scripts but instead **always use `systemctl` commands directly.**

Once logged into Perforce as a super user, the `p4 admin stop` command can also be used to stop p4d.

All init scripts can be started as the `perforce` user or the `root` user (except `p4web`, which must start initially as `root`). The application runs as the `perforce` user in any case. If the init scripts are configured as system services (non-systemd distributions), they can also be called by the `root` user using the `service` command, as in this example to start p4d:

```
service p4d_1_init start
```

Templates for the init scripts used by `mkdirs.sh` are stored in:

```
/p4/common/etc/init.d
```

There are also basic crontab templates for a Perforce master and replica server in:

```
/p4/common/etc/cron.d
```

These define schedules for routine checkpoint operations, replica status checks, and email reviews.

The Perforce should have a super user defined as named by the `P4USER` setting in `mkdir`.

To configure and start instance 1, follow these steps:

2. Start the Perforce server by calling `p4d_1_init start`.
3. Ensure that the admin user configured above has the correct password defined in `/p4/common/config/.p4passwd.p4_1.admin`, and then run the `p4login` script (which calls the `p4 login` command using the `.p4passwd.p4_1.admin` file).
4. For new servers, run this script, which sets several recommended configurables:

```
$SDP/Server/setup/configure_new_server.sh
```

For existing servers, examine this file, and manually apply the `p4 configure` command to set configurables on your Perforce server.

5. Initialize the perforce user's crontab with one of these commands:

```
crontab /p4/p4.crontab
```

and customise execution times for the commands within the crontab files to suite the specific installation.

The SDP uses wrapper scripts in the crontab: `run_if_master.sh`, `run_if_edge.sh`, `run_if_replica.sh`. We suggest you ensure these are working as desired, e.g.

```
/p4/common/bin/run_if_master.sh 1 echo yes
```

Any issues with the above indicate incorrect values for `$MASTER_ID`.

To verify that your server installation is working properly:

1. Issue the `p4 info` command, after setting appropriate environment variables by running:

```
source /p4/common/bin/p4_vars 1
```

2. If the server is running, it will display details about its settings.

Now that the server is running properly, copy the following configuration files to the `hxdepots` volume for backup purposes:

- Any init scripts used in `/etc/init.d`.
- A copy of the crontab file, obtained using `crontab -l`.
- Any other relevant configuration scripts, such as cluster configuration scripts, failover scripts, or disk failover configuration files.

3.1.7 Archiving configuration files

Now that the server is running properly, copy the following configuration files to the `hxdepots` volume for backup:

- The scheduler configuration.
- Cluster configuration scripts, failover scripts, and disk failover configuration files.

3.2 Configuring protections, file types, monitoring and security

After the server is installed and configured, most sites will want to modify server permissions (protections) and security settings. Other common configuration steps include modifying the file type map and enabling process monitoring. To configure permissions, perform the following steps:

1. To set up protections, issue the `p4 protect` command. The protections table is displayed.
2. Delete the following line:


```
write user * * //depot/...
```
3. Define protections for your server using groups. Perforce uses an inclusionary model. No access is given by default, you must specifically grant access to users/groups in the protections table. It is best for performance to grant users specific access to the areas of the depot that they need rather than granting everyone open access, and then trying to remove access via exclusionary mappings in the protect table even if that means you end up generating a larger protect table.
4. To set the server's default file types, run the `p4 typemap` command and define typemap entries to override Perforce's default behavior.

Add any file type entries that are specific to your site. Suggestions:

- For already-compressed file types (such as `.zip`, `.gz`, `.avi`, `.gif`), assign a file type of `binary+F1` to prevent the server from attempting to compress them again before storing them.
 - For regular binary files, add `binary+l` to make so that only one person at a time can check them out.
 - A sample file is provided in `$SDP/Server/config/typemap`
5. To make your changelists default to restricted (for high security environments):

```
p4 configure set defaultChangeType=restricted
```

3.3 Other server configurables

There are various configurables that you should consider setting for your server.

Some suggestions are in the file: `$SDP/Server/setup/configure_new_server.sh`

Review the contents and either apply individual settings manually, or edit the file and apply

the newly edited version. If you have any questions, please see the configurables section in Appendix of the Command Reference Guide (get the right version for your server!). You can also contact support regarding questions.

4 Backup, Replication, and Recovery

Perforce servers maintain *metadata* and *versioned files*. The metadata contains all the information about the files in the depots. Metadata resides in database (`db.*`) files in the server's root directory (`P4ROOT`). The versioned files contain the file changes that have been submitted to the server. Versioned files reside on the `hxdepots` volume.

This section assumes that you understand the basics of Perforce backup and recovery. For more information, consult the Perforce System Administrator's Guide and the Knowledge Base articles about replication.

4.1.1 Typical Backup Procedure

The SDP's maintenance scripts, run as *cron* tasks, periodically back up the metadata. The weekly sequence is described below.

Seven nights a week, perform the following tasks.

1. Truncate the active journal.
2. Replay the journal to the offline database. (Refer to Figure 2: SDP Runtime Structure and Volume Layout for more information on the location of the live and offline databases.)
3. Create a checkpoint from the offline database.
4. Recreate the offline database from the last checkpoint.

Once a week, perform the following tasks.

1. Verify all depot files.

Once every few months, perform the following tasks.

1. Stop the live server.
2. Truncate the active journal.
3. Replay the journal to the offline database. (Refer to Figure 2: SDP Runtime Structure and Volume Layout for more information on the location of the live and offline databases.)
4. Archive the live database.
5. Move the offline database to the live database directory.
6. Start the live server.
7. Create a new checkpoint from the archive of the live database.
8. Recreate the offline database from the last checkpoint.
9. Verify all depots.

This normal maintenance procedure puts the checkpoints (metadata snapshots) on the `hxdepots` volume, which contains the versioned files. Backing up the `hxdepots` volume with a normal backup utility like *robocopy* or *rsync* provides you with all the data necessary to recreate the server.

To ensure that the backup does not interfere with the metadata backups (checkpoints), coordinate backup of the `hxdepots` volume using the SDP maintenance scripts.

The preceding maintenance procedure minimizes server downtime, because checkpoints are

created from offline or saved databases while the server is running.



Be sure to back up the entire `hxdepots` volume using a normal backup utility.

With no additional configuration, the normal maintenance prevents loss of more than one day's metadata changes. To provide an optimal Recovery Point Objective (RPO), the SDP provides additional tools for replication.

4.2 Full One-Way Replication

Perforce supports a full one-way replication of data from a master server to a replica, including versioned files. The `p4 pull` command is the replication mechanism, and a replica server can be configured to know it is a replica and use the replication command. The `p4 pull` mechanism requires very little configuration and no additional scripting. As this replication mechanism is simple and effective, we recommend it as the preferred replication technique. Replica servers can also be configured to only contain metadata, which can be useful for reporting or offline checkpointing purposes. See the Distributing Perforce Guide for details on setting up replica servers.

If you wish to use the replica as a read-only server, you can use the P4Broker to direct read-only commands to the replica or you can use a forwarding replica. The broker can do load balancing to a pool of replicas if you need more than one replica to handle your load. Use of the broker may require use of a P4AUTH server for authentication.



Replication handles all server metadata and versioned file content, but not the SDP installation itself or other external scripts such as triggers. Use tools such as `robocopy` or `rsync` to replicate the rest of the `hxdepots` volume.

4.2.1 Replication Setup

To configure a replica server, first configure a machine identically to the master server (at least as regards the link structure such as `/p4`, `/p4/common/bin` and `/p4/instance/*`), then install the SDP on it to match the master server installation. Once the machine and SDP install is in place, you need to configure the master server for replication.

Perforce supports many types of replicas suited to a variety of purposes, such as:

- Real-time backup,
- Providing a disaster recovery solution,
- Load distribution to enhance performance,
- Distributed development,
- Dedicated resources for automated systems, such as build servers, and more.

We always recommend first setting up the replica as a read-only replica and ensuring that everything is working. Once that is the case you can easily modify server specs and

configurables to change it to a forwarding replica, or an edge server etc.

In the sample below, the replica name will be `replica1`, it is instance 1 on a particular host, the service user name is `svc_replica1`, and the master server's hostname is `svrmaster`.

The following sample commands illustrate how to setup a simple read-only replica.

First we ensure that `journalPrefix` is set appropriately for the master server (in this case we assume instance 1 rather than a named instance):

```
p4 configure set master#journalPrefix=/p4/1/checkpoints/p4_1
```

Then we set values for the replica itself:

```
p4 configure set replica1#P4TARGET=svrmaster:1667
p4 configure set "replica1#startup.1=pull -i 1"
p4 configure set "replica1#startup.2=pull -u -i 1"
p4 configure set "replica1#startup.3=pull -u -i 1"
p4 configure set "replica1#startup.4=pull -u -i 1"
p4 configure set "replica1#startup.5=pull -u -i 1"
p4 configure set "replica1#db.replication=readonly"
p4 configure set "replica1#lbr.replication=readonly"
p4 configure set replica1#serviceUser=svc_replica1
```

Then the following also need to be setup:

- Create a service user for the replica (Add the `Type: service` field to the user form before saving):

```
p4 user -f svc_replica1
```

- Set the service user's password:

```
p4 passwd svc_replica1
```

- Add the service user `svc_replica1` to a specific group `ServiceUsers` which has a timeout value of unlimited:

```
p4 group ServiceUsers
```

- Make sure the `ServiceUsers` group has super access in protections table:

```
p4 protect
```

Now that the settings are in the master server, you need to create a checkpoint to seed the replica. Run:

```
/p4/common/bin/daily_checkpoint.sh 1
```

When the checkpoint finishes, `rsync` the checkpoint plus the versioned files over to the replica:

```
rsync -avz /p4/1/checkpoints/p4_1.ckp.###.gz
    perforce@replica:/p4/1/checkpoints/
rsync -avz /p4/1/depots/ perforce@replica:/p4/1/depots/
```

(Assuming `perforce` is the OS user name and `replica` is the name of the replica server in the

commands above, and that ### is the checkpoint number created by the daily backup.)

Once the rsync finishes, go to the replica machine run the following:

```
/p4/1/bin/p4d_1 -r /p4/1/root -jr -z /p4/1/checkpoints/p4_1.ckp.###.gz
```

Login as the service user (specifying appropriate password when prompted), and making sure that the login ticket generated is stored in the same place as specified in the P4TICKETS configurable value set above for the replica (the following uses bash syntax):

```
source /p4/common/bin/p4_vars 1
/p4/1/bin/p4_1 -p svrmaster:1667 -u svc_replica1 login
```

Start the replica instance (either using `_init` script or `systemctl` if on system):

```
/p4/1/bin/p4d_1_init start
```

Now, you can log into the replica server itself and run `p4 pull -lj` to check to see if replication is working. If you see any numbers with a negative sign in front of them, replication is not working. The most likely cause of this is that the service user is not logged in. Rerun the steps above to login the service user and check again. If replication still is not working, check `/p4/1/logs/log` on the replica, and also look for authentication failures in the log for the master instance on `svrmaster`.

The final steps for setting up the replica server are to set up the crontab for the replica server, and set up the rsync trust certificates so that the replica scripts can run rsync without passwords.

The replica crontab is in `/p4/common/etc/init.d/crontab.replica`.

To configure the rsync trust:

On both the master and replica servers, go to the perforce user's home directory and run:

```
ssh-keygen -t rsa
```

Just use the defaults for the questions it asks.

Now from the master, run:

```
rsync -avz ~/.ssh/id_rsa.pub perforce@replica:~/.ssh/authorized_keys
```

and from the replica, run:

```
rsync -avz ~/.ssh/id_rsa.pub perforce@master:~/.ssh/authorized_keys
```

You can validate everything is working by manually running the replica crontab scripts as below. Please note that depending on the size of your checkpoints and the bandwidth between your replica and master server, these commands may take hours to rsync the checkpoints across.

```
/p4/common/bin/sync_replica.sh 1
/p4/common/bin/sync_replica.sh 1
```

Check the replica crontab template for default configuration.

- It runs `sync_replica.sh` 6 days a week at 3am
 - copy the latest checkpoint from the master server using rsync
 - restores it to the `offline_db` directory on the replica - this makes it quick to

get the replica up and running with defragmented db.* files

- It runs `sync_replica.sh` on the 7th day of the week at 3am
 - copies latest checkpoint from master using `rsync`
 - restores it to the `offline_db` directory as previously
 - stops the replica and replaces the main db.* files with the `offline_db` ones
 - The last step is the reason for the warning in the template crontab: "Don't run the line below on replicas that are being used" as the replica will down for a period

The log files will be in `/p4/1/logs`, so you can check for any errors from each script.

4.3 Recovery Procedures

There are three scenarios that require you to recover server data:

Metadata	Depotdata	Action required
lost or corrupt	Intact	Recover metadata as described below
Intact	lost or corrupt	Call Perforce Support
lost or corrupt	lost or corrupt	Recover metadata as described below. Recover the <code>hxdepots</code> volume using your normal backup utilities.

Restoring the metadata from a backup also optimizes the database files.

4.3.1 Recovering a master server from a checkpoint and journal(s)

The checkpoint files are stored in the `/p4/instance/checkpoints` directory, and the most recent checkpoint is named `p4_instance.ckp.number.gz`. Recreating up-to-date database files requires the most recent checkpoint, from `/p4/instance/checkpoints` and the journal file from `/p4/instance/logs`.

To recover the server database manually, perform the following steps from the root directory of the server (`/p4/instance/root`).

1. Stop the Perforce Server by issuing the following command:

```
/p4/common/bin/p4master_run instance \
/p4/instance/bin/p4_instance admin stop
```

2. Delete the old database files in the `/p4/instance/root/save` directory
3. Move the live database files (db.*) to the `save` directory.
4. Use the following command to restore from the most recent checkpoint.

```
/p4/instance/bin/p4d_instance -r /p4/instance/root -jr -z \
/p4/instance/checkpoints/p4_instance.ckp.most_recent#.gz
```

5. To replay the transactions that occurred after the checkpoint was created, issue the following command:

```
/p4/instance/bin/p4d_instance -r /p4/instance/root -jr \
/p4/instance/logs/journal
```

6. Restart your Perforce server.

If the Perforce service starts without errors, delete the old database files from `/p4/instance/root/save`.

If problems are reported when you attempt to recover from the most recent checkpoint, try

recovering from the preceding checkpoint and journal. If you are successful, replay the subsequent journal. If the journals are corrupted, contact Perforce Technical Support. For full details about backup and recovery, refer to the Perforce System Administrator's Guide.

4.3.2 Recovering a replica from a checkpoint

This is very similar to creating a replica in the first place as described above.

If you have been running the replica crontab commands as suggested, then you will have the latest checkpoints from the master already copied across to the replica.

See the steps in the script `weekly_sync_replica.sh` for details (note that it deletes the `state` and `rdb.lbr` files from the replica root directory so that the replica starts replicating from the start of a journal).

Remember to ensure you have logged the service user in to the master server (and that the ticket is stored in the correct location as described when setting up the replica).

4.3.3 Recovering from a tape backup

This section describes how to recover from a tape or other offline backup to a new server machine if the server machine fails. The tape backup for the server is made from the `hxdepots` volume. The new server machine must have the same volume layout and user/group settings as the original server. In other words, the new server must be as identical as possible to the server that failed.

To recover from a tape backup, perform the following steps.

1. Recover the `hxdepots` volume from your backup tape.
2. Create the `/p4` convenience directory on the OS volume.
3. Create the directories `/metadata/p4/instance/root/save` and `/metadata/p4/instance/offline_db`.
4. Change ownership of these directories to the OS account that runs the Perforce processes.
5. Switch to the Perforce OS account, and create a link in the `/p4` directory to `/depotadata/p4/instance`.
6. Create a link in the `/p4` directory to `/hxdepots/p4/common`.
7. As a super-user, reinstall and enable the `init.d` scripts
8. Find the last available checkpoint, under `/p4/instance/checkpoints`
9. Recover the latest checkpoint by running:

```
/p4/instance/bin/p4d_instance -r /p4/instance/root -jr -z last_ckp_file
```

10. Recover the checkpoint to the `offline_db` directory:

```
/p4/instance/bin/p4d_instance -r /p4/instance/offline_db -jr -z last_ckp_file
```

11. Reinstall the Perforce server license to the server root directory.
12. Start the perforce service by running `/p4/1/bin/p4d_1_init start`
13. Verify that the server instance is running.
14. Reinstall the server crontab or scheduled tasks.

15. Perform any other initial server machine configuration.
16. Verify the database and versioned files by running the `p4verify . sh` script. Note that files using the `+k` file type modifier might be reported as `BAD!` after being moved. Contact Perforce Technical Support for assistance in determining if these files are actually corrupt.

4.3.4 Failover to a replicated standby machine

See `DR-Failover-Steps-Linux.docx`

5 Server Maintenance

This section describes typical maintenance tasks and best practices for administering server machines.

5.1 Server upgrades

Upgrading a server instance in the SDP framework is a simple process involving a few steps.

- Download the new p4 and p4d executables for your OS from ftp.perforce.com and place them in `/p4/common/bin`
- Run:

```
/p4/common/bin/upgrade.sh instance
```
- If you are running replicas, upgrade the replicas first, and then the master (outside -> in)

5.2 Database Modifications

Occasionally modifications are made to the Perforce database from one release to another. For example, server upgrades and some recovery procedures modify the database.

When upgrading the server, replaying a journal patch, or performing any activity that modifies the `db.*` files, you must restart the offline checkpoint process so that the files in the `offline_db` directory match the ones in the live server directory. The easiest way to restart the offline checkpoint process is to run the `live_checkpoint` script after modifying the `db.*` files, as follows:

```
/p4/common/bin/live_checkpoint.sh instance
```



Be aware that this script will lock your production server and prevent any user commands for the duration of the checkpoint. Look at the latest `/p4/instance/logs/checkpoint_log.log` for an indication of how long this will take.

This script makes a new checkpoint of the modified database files in the live root directory, then recovers that checkpoint to the `offline_db` directory so that both directories are in sync. This script can also be used anytime to create a checkpoint of the live database.

This command should be run when an error occurs during offline checkpointing. It restarts the offline checkpoint process from the live database files to bring the offline copy back in sync. If the live checkpoint script fails, contact Perforce Consulting at consulting@perforce.com.

5.3 Maximizing Server Performance

The following sections provide some guidelines for maximizing the performance of the Perforce Server, using tools provided by the SDP. More information on this topic can be found in the System Administrator's Guide and in the Knowledge Base.

5.4 Optimizing the database files

The Perforce Server's database is composed of b-tree files. The server does not fully rebalance and compress them during normal operation. To optimize the files, you must checkpoint and restore the server. This normally only needs to be done very few months.

To minimize the size of back up files and maximize server performance, minimize the size of the `db.have` and `db.label` files. The scripts for **Error! Reference source not found.**, and

Error! Reference source not found., help achieve this goal. For best server performance, run these scripts weekly via `/p4/sdp/Maintenance/maintenance`

5.5 Proactive Performance Maintenance

This section describes some things that can be done to proactively to enhance scalability and maintain performance.

5.5.1 Limiting large requests

To prevent large requests from overwhelming the server, you can limit the amount of data and time allowed per query by setting the `maxresults`, `maxscanrows` and `maxlocktime` parameters to the lowest setting that does not interfere with normal daily activities. As a good starting point, set `maxscanrows` to `maxresults * 3`; set `maxresults` to slightly larger than the maximum number of files the users need to be able to sync to do their work; and set `maxlocktime` to 30000 milliseconds. These values must be adjusted up as the size of your server and the number of revisions of the files grow. To simplify administration, assign limits to groups rather than individual users.

To prevent users from inadvertently accessing large numbers of files, define their client view to be as narrow as possible, considering the requirements of their work. Similarly, limit users' access in the `protections` table to the smallest number of directories that are required for them to do their job.

Finally, keep triggers simple. Complex triggers increase load on the server.

5.5.2 Offloading remote syncs

For remote users who need to sync large numbers of files, Perforce offers a proxy server. P4P, the Perforce Proxy, is run on a machine that is on the remote users' local network. The Perforce Proxy caches file revisions, serving them to the remote users and diverting that load from the main server.

P4P is included in the Windows installer. To launch P4P on Unix machines, copy the `/p4/common/etc/init.d/p4p_1_init` script to `/p4/1/bin/p4p_1_init`. Then review and customize the script to specify your server volume names and directories.

P4P does not require special hardware but it can be quite CPU intensive if it is working with binary files, which are CPU-intensive to attempt to compress. It doesn't need to be backed up. If the P4P instance isn't working, users can switch their port back to the main server and continue working until the instance of P4P is fixed.

6 Tools and Scripts

This section describes the various scripts and files provided as part of the SDP package. To run main scripts, the machine must have Python 2.7, and a few scripts require Perl 5. The Maintenance scripts can be run from the server machine or from client machines.

The following various scripts.

6.1 Core Scripts

The core SDP scripts are those related to checkpoints and other scheduled operations, and all run from `/p4/common/bin`.

6.1.1 `p4_vars`

Defines the environment variables required by the Perforce server. This script uses a specified instance number as a basis for setting environment variables. It will look for and open the respective `p4_<instance>.vars` file (see next section).

This script also sets server logging options and configurables.

Location: `/p4/common/bin`

6.1.2 `p4_<instance>.vars`

Defines the environment variables for a specific instance, including P4PORT etc.

Location: `/p4/common/config`

6.1.3 `p4master_run`

This is the wrapper script to other SDP scripts. This ensures that the shell environment is loaded from `p4_vars`. It provides a `'-c'` flag for silent operation, used in many `crontab` so that email is sent from the scripts themselves.

Location: `/p4/common/bin`

6.1.4 `recreate_offline_db`

Recovers the `offline_db` database from the latest checkpoint and replays any journals since then. If you have a problem with the offline database then it is worth running this script first before running `live_checkpoint`, as the latter will stop the server while it is running which can take hours.

Run this script if an error occurs while replaying a journal during weekly or daily checkpoint process.

Location: `/p4/common/bin`

6.1.5 `live_checkpoint`

Stops the server, creates a checkpoint from the live database files, recovers the `offline_db` database from that checkpoint to rebalance and compress the files, then recovers the checkpoint in the `offline_db` directory to ensure that the database files are optimized.

Run this script when creating the server and if an error occurs while replaying a journal during the off-line checkpoint process. Be aware it locks live database for the duration of the checkpoint which can take hours.

Location: `/p4/common/bin`

6.1.6 `daily_checkpoint`

This script is configured to run six days a week using `crontab` or the Windows scheduler. The script truncates the journal, replays it into the `offline_db` directory, creates a new checkpoint from the resulting database files, then recreates the `offline_db` directory from the new

checkpoint.

This procedure rebalances and compresses the database files in the `offline_db` directory, which are rotated into the live database directory once a week by the `weekly_checkpoint` script.

Location: `/p4/common/bin`

6.1.7 p4verify

Verifies the integrity of the depot files. This script is run by `crontab` on a regular basis.

Location: `/p4/common/bin`

6.1.8 p4review.py

Sends out email containing the change descriptions to users who are configured as reviewers for affected files (done by setting the `Reviews:` field in the user specification). This script is a version of the `p4review.py` script that is available on the Perforce Web site, but has been modified to use the server instance number. It relies on a configuration file in `/p4/common/config`, called `p4_<instance>.p4review.cfg`. On Windows, a driver called `run_p4review.cmd`, located in the same directory, allows you to run the review daemon through the Windows scheduler.

This is not required if you have installed Swarm which also performs notification functions and is easier for users to configure.

Location: `/p4/common/bin`

6.1.9 p4login

Executes a `p4 login` command, using the administration password configured in `mkdirs.cfg` and stored in a text file: `/p4/common/config/.p4passwd`

Location: `/p4/common/bin`

6.1.10 p4d_instance_init

Starts the Perforce server. Do not use if you have configured `systemctl` for `systemd` Linux distributions such as CentOS 7.x

This script sources `/p4/common/bin/p4_vars`, then `/p4/common/bin/p4d_base`.

Note: In clustered environments, put this script in the `/p4/instance/bin` directory and configure your cluster software to launch it from this location.

Location: `/p4/instance/bin` with a symlink to it from `/etc/init.d` (or a copy in `/etc/init.d` in a clustered environments). Templates for init scripts for other Perforce server products exist in `p4/common/etc/init.d`

6.2 More Server Scripts

These scripts are helpful components of the SDP that run on the server, but are not included in the default `crontab` schedules.

6.2.1 upgrade.sh

Runs a typical upgrade process, once new `p4` and `p4d` binaries are available in

/p4/common/bin.

This script will:

- Rotate the journal (for clean recovery point)
- Apply all necessary journals to offline_db
- Stop the server
- Create an appropriately versioned link for new p4/p4d/p4broker etc
- Link those into /p4/1/bin (per instance)
- Run p4d -xu on live and offline_db to perform database upgrades
- Restart server

Location: /p4/common/bin

6.2.2 p4.crontab

Contains crontab entries to run the server maintenance scripts. The p4.crontab.solaris script is for Solaris.

Location: /p4/sdp/Server/Unix/p4/common/etc/cron.d

6.3 Other Files

The following table describes other files in the SDP distribution. These files are usually not invoked directly by you; rather, they are invoked by higher-level scripts.

File	Location	Remarks
dummy_ip.txt	\$SDP/Server/config	Instructions for using a license on more than one machine. Typically used to enable a standby server. Contact Perforce Licensing before using.
backup_functions.sh	/p4/common/bin	Unix/Linux only. Utilities for maintenance scripts.
p4admin_verify_client.bat	/p4/common/bin	Unix/Linux only. Used by p4verify.sh.
p4d_base	/p4/common/bin	Unix/Linux only. Template for Unix/Linux init.d scripts.
template.(pl sh)	/p4/common/bin	Sample script templates for Bash and Perl scripts.
mirror_ldap_groups.pl	/p4/common/bin	Script to mirror selected groups from an LDAP server (e.g. Active Directory).
Perl Modules (*.pm files)	/p4/common/lib	Modules used by some Perl scripts.
change.txt	\$SDP/Maintenance	Template for new pending changelist.

7 Appendix A – Directory Structure Configuration Script for Linux/Unix

This script describes the steps performed by the `mkdirs.sh` script on Linux/Unix platforms. Please review this appendix carefully before running these steps manually. Assuming the three-volume configuration described in the Volume Layout and Hardware section are used, the following directories are created. The following examples are illustrated with “1” as the server instance number.

Directory	Remarks
<code>/p4</code>	Must be under / on the OS volume
<code>/hxdepots/p4/1/bin</code>	Files in here are generated by the <code>mkdirs.sh</code> script.
<code>/hxdepots/p4/1/depots</code>	
<code>/hxdepots/p4/1/tmp</code>	
<code>/hxdepots/p4/common/config</code>	Contains <code>p4_<instance>.vars</code> file.
<code>/hxdepots/p4/common/bin</code>	Files from <code>\$SDP/Server/Unix/p4/common/bin</code> .
<code>/hxdepots/p4/common/etc</code>	Contains <code>init.d</code> and <code>cron.d</code> .
<code>/hxlogs/p4/1/logs/old</code>	
<code>/hxmetadata2/p4/1/db2</code>	Contains offline copy of main server databases (linked by <code>/p4/1/offline_db</code>).
<code>/hxmetadata1/p4/1/db1/save</code>	Used only during <code>recreate_db_checkpoint.sh</code> for extra redundancy.

Next, `mkdirs.sh` creates the following symlinks in the `/hxdepots/p4/1` directory:

Link source	Link target	Command
<code>/hxmetadata1/p4/1/db1</code>	<code>/p4/1/root</code>	<code>ln -s /hxmetadata1/p4/1/root</code>
<code>/hxmetadata2/p4/1/db2</code>	<code>/p4/1/offline_db</code>	<code>ln -s /hxmetadata1/p4/1/offline_db</code>
<code>/hxlogs/p4/1/logs</code>	<code>/p4/1/logs</code>	<code>ln -s /hxlogs/p4/1/logs</code>

Then these symlinks are created in the `/p4` directory:

Link source	Link target	Command
<code>/hxdepots/p4/1</code>	<code>/p4/1</code>	<code>ln -s /hxdepots/p4/1</code>
<code>/hxdepots/p4/common</code>	<code>/p4/common</code>	<code>ln -s /hxdepots/p4/common</code>

Next, `mkdirs.sh` renames the Perforce binaries to include version and build number, and then creates appropriate symlinks.

The structure is shown in this example, illustrating values for two instances, with instance #1 using Perforce 2018.1 and instance #2 using 2018.2. Files are shown in red and symlinks in

green.

In /p4/common/bin:

```
p4_2018.1_bin → p4_2018.1.685046
p4d_2018.1_bin → p4d_2018.1.685046
p4_2018.2_bin → p4_2018.2.700949
p4d_2018.2_bin → p4d_2018.2.700949
p4_1_bin → p4_2018.1_bin
p4d_1_bin → p4d_2018.1_bin
p4_2_bin → p4_2018.2_bin
p4d_2_bin → p4d_2018.2_bin
```

In /p4/1/bin:

```
p4_1 → /p4/common/bin/p4_1_bin
p4d_1 → /p4/common/bin/p4d_1_bin
```

In /p4/2/bin:

```
p4_2 → /p4/common/bin/p4_2
p4d_2 → /p4/common/bin/p4d_2
```

8 Appendix B – Frequently Asked Questions/Troubleshooting

This appendix lists common questions and problems encountered by SDP users. Do not hesitate to contact consulting@perforce.com if additional assistance is required.

8.1.1 Journal out of sequence

This error is encountered when the offline and live databases are no longer in sync, and will cause the offline checkpoint process to fail. Because the scripts will replay all outstanding journals, this error is much less likely to occur. This error can be fixed by running the `live_checkpoint.sh` script. Alternatively, if you know that the checkpoints created from previous runs of `daily_checkpoint.sh` are correct, then restore the `offline_db` from the last known good checkpoint.

8.1.2 Unexpected end of file in replica daily sync

Check the start time and duration of the `daily_checkpoint.sh` cron job on the master. If this overlaps with the start time of the `sync_replica.sh` cron job on a replica, a truncated checkpoint may be rsync'd to the replica and replaying this will result in an error.

Adjust the replica's cronjob to start later to resolve this.

Default cron job times, as installed by the SDP are initial estimates, and should be adjusted to suit your production environment.