



Perforce Software, Inc.

510.864.7400

consulting@perforce.com

www.perforce.com

Disaster Recovery Failover Plan Windows

Doc Version 1.3

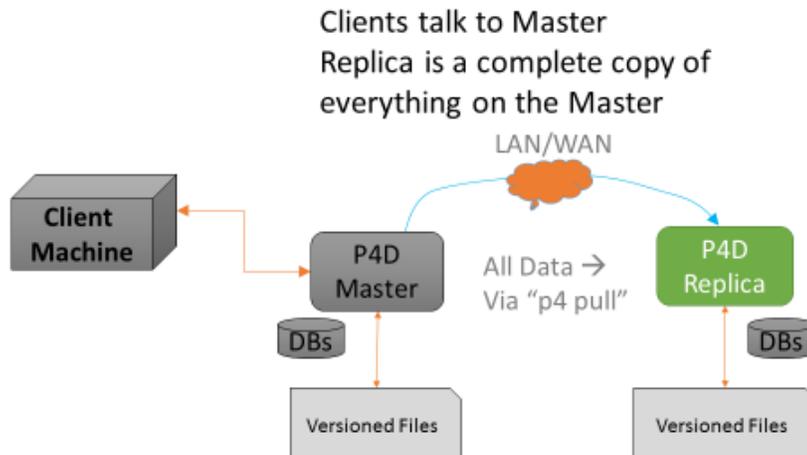
17 November, 2016

Table of Contents

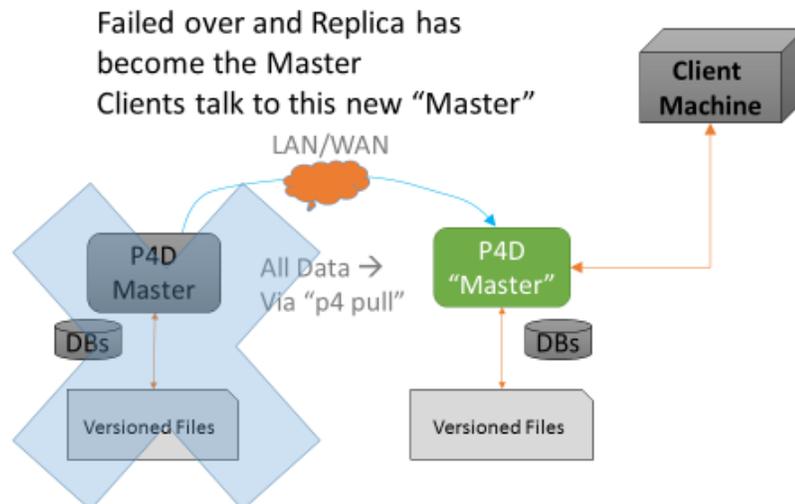
1	Overview.....	2
2	Failover Procedure (for Windows).....	3
2.1	Introduction.....	3
2.2	Check replica status	3
2.3	Change the replica to become the master	3
2.4	Update the scheduled jobs	4
2.5	Change the original master to become the replica	4

1 Overview

The DR server is being replicated using the Perforce built in server to server replication via the “p4 pull” command and the replication of checkpoints. This document provides the steps to failover from the main server to the DR server.



We wish to go from the above situation to:



2 Failover Procedure (for Windows)

2.1 Introduction

This document is specific to the Windows SDP and commands and paths assume that. If you are running the SDP on Unix then please consult the equivalent guide for Unix.

This guide also assumes the generic SDP installation process. If you have customized the SDP then you will need to customize this guide – the general principles still apply.

The following values are defined in `sdp_config.ini` within the `<metadata>` drive.

- `<metadata>` - default value is E:
- `<depotdata>` - default value is F:
- `<logs>` - default value is G:

2.2 Check replica status

This is just a verification that replication is up to date before we start the whole process. Do these steps if the master is still functioning. If the master is not running and not able to run, you would normally start by assuming that the replica is up-to-date.

On the replica machine, run “`p4 pull -l`” to check the status of the replication. You want to wait for the replica to fully catch up before shutting down the master and replica instances. You can also verify that the replica is up to date by checking the values on the servers. For instance, run `p4 counters` on both the master and the replica to make sure they match. Once you have verified the replication is up to date:

stop the replica by running “`net stop p4_1`” on the replica server

stop the master by running “`net stop p4_1`” on the master server

2.3 Change the replica to become the master

Change the DNS entry on your network or move the VIP from the master to the new master (previously the replica) if you are using a VIP address.

On the replica server:

- edit “`<metadata>:\p4\1\root\server.id`” and change the name to “Master” (without quotes) and save the file. Note that the file is read-only and needs to be saved as read-only.

Now you can restart the server on the replica machine as the master by running “`net start p4_1`”. Next run:

```
del <metadata>:\p4\1\root\state
```

```
del <metadata>:\p4\1\root\rdb.lbr
```

2.4 Update the scheduled jobs

On the new master server (originally the replica), run:

```
at /delete
at 03:00 /every:M,T,W,Th,F,Su
    <depotdata>:\p4\1\bin\daily_backup.bat
at 03:00 /every:S <depotdata>:\p4\1\bin\weekly_backup.bat
at 06:00 /every:S <depotdata>:\p4\1\bin\p4verify.bat
```

On the original master server, now the replica, use task scheduler, and schedule:

```
Every M,T,W,Th,F,Su <depotdata>:\p4\1\bin\sync_replica.bat
to run at 06:00am
Every S <depotdata>:\p4\1\bin\weekly_sync_replica.bat to
run at 06:00am
```

Both tasks must be schedule to run as a local account on the machine or network rather than the SYSTEM account because they need to be able to map a drive to the (new) master server.

2.5 Change the original master to become the replica

On the original master (the new replica), edit “<metadata>\p4\1\root\server.id and change it to “replica1” (without quotes) or whatever the value in your original installation was.

Delete the following files:

- <logs>:\p4\1\logs\journal
- <metadata>:\p4\1\root\state (if it exists)
- <metadata>:\p4\1\root\rdb.lbr” (if it exists)

All of the commands below are run on the new replica (old master) machine.

You also need to login into the (new) master server as the service user on the (new) replica so that replication can run:

```
p4 -p master_server_address:1667 -u service login
```

Enter the service user’s password. (Same as p4admin password.)

Now start the server as a replica by running

```
net start p4_1
```

Check replication by running:

```
p4 pull -1
```

Finally, run these final two steps:

```
<depotdata>:\p4\1\bin\sync_replica.bat
```

The command that task scheduler normally runs in order to make sure it is working properly.

NOTE: If the master server has been down for more than 7 days, you will have to reset the replica as follows before starting it. The reason for the 7 days value is that the number of old checkpoints and journals to keep is set to 7 in the sdp_config.ini file. The replication uses the old journals in order to catch up from where it last stopped, so if the old ones have rotated off, you have to reset the replica with a new checkpoint and updated versioned files.

To reset the replica, from the replica run:

```
net use z: \\master_server\<depotdata_drive>$  
attrib -r <depotdata>:\p4\1\checkpoints\*.*  
del <depotdata>:\p4\1\checkpoints\*.*  
xcopy /S z:\p4\1\checkpoints <depotdata>:\p4\1\checkpoints  
xcopy /S z:\p4\1\depots <depotdata>:\p4\1\depots  
del <metadata>:\p4\1\root\db.*  
del <metadata>:\p4\1\offline_db\db.*  
del <logs>:\p4\1\logs\journal  
dir <depotdata>:\p4\1\checkpoints
```

We need to use the highest numbered p4_1.ckp.*.gz file in that directory. Assume it is 10 for this example:

```
<depotdata>:\p4\1\bin\p4d -r <metadata>:\p4\1\root -jr -z  
    <depotdata>:\p4\1\checkpoints\p4_1.ckp.10.gz  
<depotdata>:\p4\1\bin\p4 -p master_server:1666 -u service  
    login  
net start p4_1  
<depotdata>:\p4\1\bin\p4d -r <metadata>:\p4\1\offline_db -  
    jr -z <metadata>:\p4\1\checkpoints\p4_1.ckp.10.gz
```