Perforce Replication

The Definitive Guide

Sven Erik Knop
Senior Consultant

PERFORCE SOFTWARE

SOME QUOTES

Deutsch's Eight Fallacies of Distributed Computing:

- The network is reliable.
- Latency is zero.
- Bandwidth is infinite.
- The network is secure.
- Topology doesn't change.
- There is one administrator.
- Transport cost is zero.
- The network is homogeneous.

(Peter Deutsch)

CURRENT PROBLEMS

- Perforce server benefits from replication
  - High availability
  - Disaster recovery
  - Load sharing (for example for reports and build servers)
- rsync has issues
  - Cannot handle running Perforce server gracefully
  - Building file list takes too long
  - No transactional safety
  - External dependency and processes

SOLUTION

- Server-to-Server replication
  - Directly supported by Perforce “out-of-the-box”
  - Asynchronous based on journal file
  - Supports both Metadata-only and full replication
  - No need for external scripts, complete solution

- Replica server is running in read-only mode
  - Requires separate license file (free of charge)

ARCHITECTURE

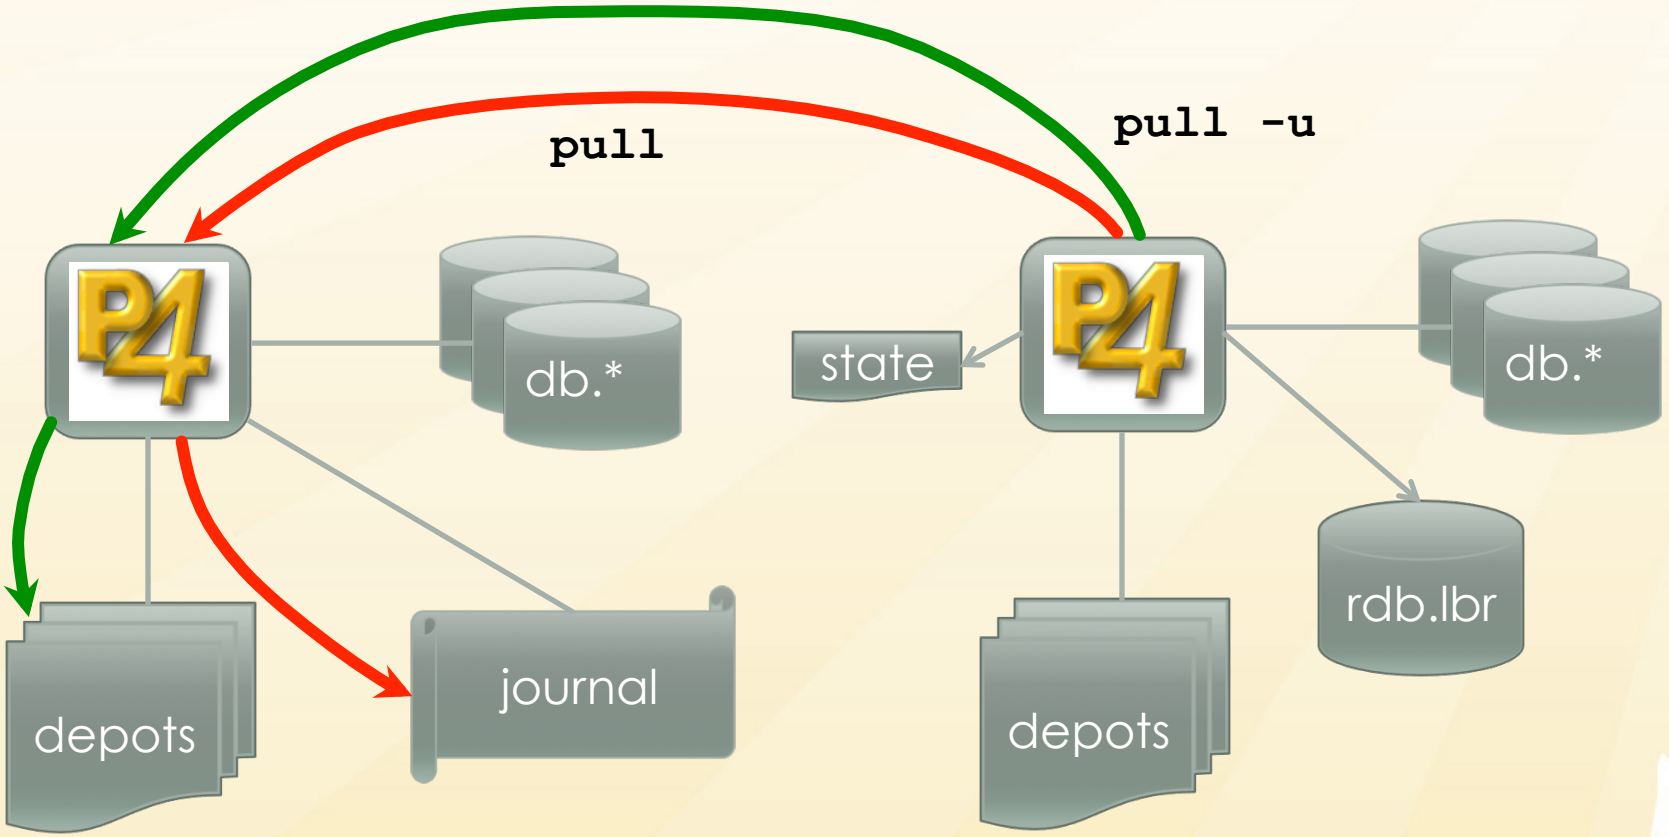

P4 PULL

- Run against the replica server

Command	Effect
`p4 pull`	Retrieve missing journal entries, then terminate
`p4 pull -i <N>`	Continuously pull every <N> seconds
`p4 pull -u`	Retrieve missing file revisions, then terminate
`p4 pull -u -i <N>`	Continuously pull file revisions
`p4 pull -l`	List missing file revisions or errors

- Can be run as background task inside the replica

HOW DOES 'PULL' KEEP TRACK?

- **state** file
  - Text file normally located in the replica P4ROOT directory
  - `journal#/offset`
  - Allows replication to be interrupted
  - Master server can rotate journal file
    - Specify '-J prefix' if master uses journal prefix for checkpoints

- **rdb.lbr** database
  - Binary file located in the replica P4ROOT directory
  - Contains information on missing archive revisions

JOURNAL ROTATION AND PREFIX

- **Master**
  - `p4 admin checkpoint/journal [-z] prefix`
  - Do not use –z (compression)!
  - If you use a prefix, use the same prefix for ‘p4 pull’
- **Replica**
  - `p4 pull –J prefix [-i <N>]`
  - Journal will rotated in sync with the master (in P4ROOT)

CONFIGURATION

- ‘p4 pull’ is designed to be a background process
  - Started from the replica server
  - One process for retrieving metadata
  - Several additional processes to retrieve archive data

- Use the new ‘p4 configure set’
  - `p4 configure set monitor=3`
  - `p4 configure set repl1#statefile=repl1_state`

PREPARE IN THE MASTER

P4NAME=Master

checkpoint

restore

P4NAME=Repl_1

```
monitor=1
Repl_1#monitor=3
Repl_1#P4TARGET=master:1666
Repl_1#P4PORT=repl1:1666
Repl_1#P4LOG=repl1_log
Repl_2#monitor=3
Repl_2#P4TARGET=master:1666
Repl_2#P4PORT=repl2:1999
Repl_2#P4LOG=repl2_log
```

```
monitor=1
Repl_1#monitor=3
Repl_1#P4TARGET=master:1666
Repl_1#P4PORT=repl1:1666
Repl_1#P4LOG=repl1_log
Repl_2#monitor=3
Repl_2#P4TARGET=master:1666
Repl_2#P4PORT=repl2:1999
Repl_2#P4LOG=repl2_log
```

P4NAME determines which configuration is active

CONFIGURATION PARAMETERS

Parameter	Values (examples)
P4PORT	1666
P4TARGET	master:1666
db.replication	readonly
lbr.replication	readonly
serviceUser	service_replica
monitor	1
startup.1	pull –i 1 [-J prefix]
startup.2	pull –u –i 1
startup.3	pull –u –i 1

SERVICEUSER

- Special user for background processes
  - Type: Service
- Ignores AUTH_CHECK trigger, local password instead
- Needs entry in the protection table, typically ‘super’
- Does not consume a license
- Can only run a few limited commands

p4 login	p4 logout	p4 passwd
p4 info	p4 user	

- Needs to be logged in before replication can start
  - P4TICKETS

MONITORING

- p4 monitor show –a (on replica)
  - `695 R service    72:22:23 pull -i 1`
  - `696 R service    72:22:23 pull -u -i 1`
  - `697 R service    72:22:23 pull -u -i 1`
- p4 logtail (on master, with server=1..3)
  - rmt-Journal
  - rmt-FileFetch
- p4 pull –l
  - Reports pending archive file transfers
- p4 verify

CONNECT TO THE REPLICA

- Replica is read-only
- ‘p4 login’ requires database change for ticket

- Solution:
  - Replica forwards request to Master
  - ‘p4 pull’ retrieves ticket from Master

- ‘p4 login’ can experience delay

- Alternative: P4AUTH pointing to master server
  - Faster in LAN, but slower in WAN

COMMANDS ON THE REPLICA

- Only read-only commands are allowed
  - ‘p4 sync –p’, ‘p4 print’, but not ‘p4 sync’
- Clients used against the replica must be created on the master server
  - Will be replicated across
- Timestamps do not get updated

USE CASES: HIGH AVAILABILITY

- Recommended: identical hardware to master
- Asynchronous solution:
  - Replica can be a few seconds behind master
  - HA server usually within the same LAN
- Failover procedure:
  - Stop replica
  - Restart with P4NAME set to master name
  - ➔ Replica becomes the master

- Currently no fallback to former master
  - Need to build a new replica
  - Do not automate failover!

USE CASES

USE CASES: DISASTER RECOVERY

- Hardware requirements typically less than HA
  - DR Server is not expected to handle the same load
- DR replica can be several minutes behind master
  - DR server needs to be physically separated from master
  - RPO (recovery point objective) depends on bandwidth
- Failover scenarios similar to HA

USE CASES: BUILD SERVER

- Usable for full builds
  - No incremental file update of the workspace
  - Need to use ‘p4 sync –p’ or ‘p4 print’ to retrieve files
  - Use long-lasting ticket or P4AUTH to avoid delay waiting for ‘p4 login’
- Can use P4Broker to redirect read-only commands to replica (http://kb.perforce.com/article/1354)
- Alternative to P4Proxy for builds

USE CASES: REPORTING

- Reporting does not require access to archive files
- Reporting replica:
  - lbr.replication=none
  - No ‘p4 pull –u’ background tasks
- Can also be used with P4Broker
- Use ‘p4 replicate’ instead of ‘p4 pull’ to filter
  - No background task, needs to be run separately

THE FUTURE

- Backup using replica server
  - Missing link between replica checkpoint and master journal
- Federated servers
  - Advanced proxy with some Metadata
- Better reporting and recovery options

QUESTIONS?

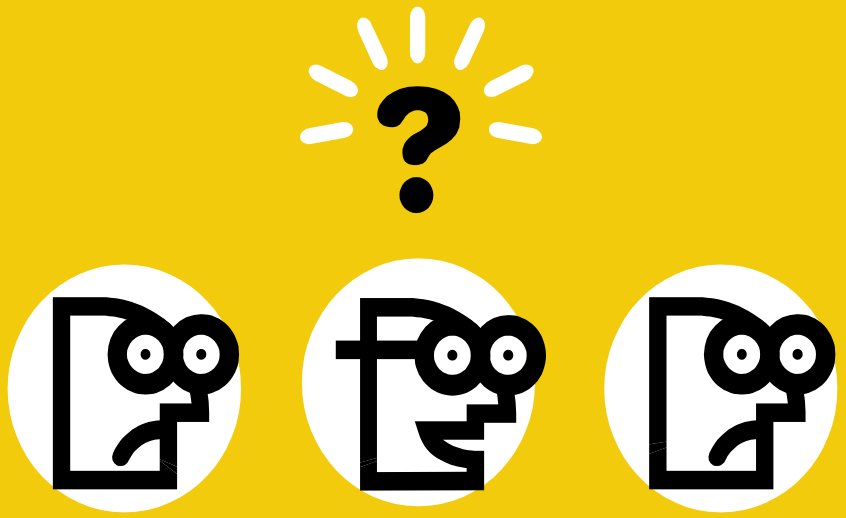

It's QUESTION TIME!!