

Mini-review + cheat sheet: ZFS on FUSE

Dan Şerban

Instalarea și pornirea serviciului

- **Compilare din surse (branșa "critical"):**
`git clone -b critical`
`http://git.zfs-fuse.net/official`
`zfs-fuse.git`
- **Dependencies:**
`glibc, fuse, zlib, libaio`
- **FUSE e instalat? Modulul e încărcat?**
`lsmod | grep fuse`
- **OK, se pornește serviciul (ca root):**
`zfs-fuse &`

The ZFS Way

- **The three primary goals of ZFS:**
 - **Highly scalable data repository**
 - **Ease of administration**
 - **Guaranteed on-disk data integrity**
- **Motto:**
**Throwing away 20 years worth
of obsolete assumptions
about storage**

ZFS features

- **Complete storage subsystem**
integrated view of the logical and physical structure of the data
- **Copy-On-Write transactional semantics**
- **Self-healing**
 - ZFS auto-corrects silent data corruption
 - there is never a need to fsck a ZFS dataset
- **RAID-Z uses dynamic stripe width**
it never has to do read-modify-write, only write
- **On-the-fly compression**
- **Deduplication**
- **NFS, SMB, iSCSI sharing**
- **Hybrid storage pools with L2ARC**

Terminologie

- **VDEV**
- **Pool**
- **DataSet**
- **ZVOL**
- **ZIL**
- **ARC**
- **L2ARC**
- **RTWWWMP**

**Read The Wonderfully Well-Written
Man Pages:**

man zpool

man zfs

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#  
# zpool status -v  
no pools available  
#  
# zpool list  
no pools available  
#  
# zfs list  
no datasets available  
#  
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#  
# zpool create rlugdemopool mirror /dev/sda10 /dev/sda11 /dev/sda12  
#
```

```
# zpool status -v  
pool: rlugdemopool  
state: ONLINE  
scrub: none requested  
config:
```

NAME	STATE	READ	WRITE	CKSUM
rlugdemopool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
sda10	ONLINE	0	0	0
sda11	ONLINE	0	0	0
sda12	ONLINE	0	0	0

```
errors: No known data errors
```

```
#  
# zpool list  
NAME          SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT  
rlugdemopool  5.09G  76.5K  5.09G  0%  ONLINE  -
```

```
#  
# zfs list  
NAME          USED  AVAIL  REFER  MOUNTPOINT  
rlugdemopool  72K  5.01G  21K  /rlugdemopool
```

```
#  
#
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#  
# zpool detach rlugdemopool /dev/sda12  
#
```

```
# zpool status -v  
pool: rlugdemopool  
state: ONLINE  
scrub: none requested  
config:
```

NAME	STATE	READ	WRITE	CKSUM
rlugdemopool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
sda10	ONLINE	0	0	0
sda11	ONLINE	0	0	0

```
errors: No known data errors
```

```
#  
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#  
# zpool add rlugdemopool spare /dev/sda12  
#
```

```
# zpool status -v  
pool: rlugdemopool  
state: ONLINE  
scrub: none requested  
config:
```

NAME	STATE	READ	WRITE	CKSUM
rlugdemopool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
sda10	ONLINE	0	0	0
sda11	ONLINE	0	0	0
spares				
sda12	AVAIL			

```
errors: No known data errors
```

```
#  
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#
# zpool offline rlugdemopool /dev/sda11
#
# zpool status -v
pool: rlugdemopool
state: DEGRADED
status: One or more devices has been taken offline by the administrator.
        Sufficient replicas exist for the pool to continue functioning in a
        degraded state.
action: Online the device using 'zpool online' or replace the device with
        'zpool replace'.
scrub: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
rlugdemopool	DEGRADED	0	0	0
mirror	DEGRADED	0	0	0
sda10	ONLINE	0	0	0
sda11	OFFLINE	0	0	0
spares				
sda12	AVAIL			

```
errors: No known data errors
```

```
#
```

```
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#
# zpool replace rlugdemopool /dev/sda11 /dev/sda13
#
# zpool status -v
pool: rlugdemopool
state: ONLINE
scrub: resilver completed after 0h0m with 0 errors on Wed May 12 11:51:56 2010
config:
```

NAME	STATE	READ	WRITE	CKSUM	
rlugdemopool	ONLINE	0	0	0	
mirror	ONLINE	0	0	0	
sda10	ONLINE	0	0	0	
sda13	ONLINE	0	0	0	78K resilvered
spares					
sda12	AVAIL				

```
errors: No known data errors
```

```
#
```

```
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#
# zpool scrub rlugdemopool
#
# zpool status -v
pool: rlugdemopool
state: ONLINE
scrub: scrub completed after 0h0m with 0 errors on Wed May 12 12:24:54 2010
config:
```

NAME	STATE	READ	WRITE	CKSUM
rlugdemopool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
sda10	ONLINE	0	0	0
sda13	ONLINE	0	0	0
spares				
sda12	AVAIL			

```
errors: No known data errors
```

```
#
```

```
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#
# zpool iostat 5
          capacity      operations      bandwidth
pool      used  avail  read  write  read  write
-----
r lugdemopool  91.5K  5.09G    0    0    58    325
r lugdemopool  91.5K  5.09G    0    0     0     0
r lugdemopool  91.5K  5.09G    0    0     0     0
^C
#
# █
```

Hands-on with ZFS

Terminal - root@skynet:~

File Edit View Terminal Go Help

```
#  
# zfs create rlugdemopool/dserban  
#  
# zfs list  
NAME                USED  AVAIL  REFER  MOUNTPOINT  
rlugdemopool        116K  5.01G   21K    /rlugdemopool  
rlugdemopool/dserban  21K   5.01G   21K    /rlugdemopool/dserban  
#  
# 
```

Hands-on with ZFS

Terminal - root@skynet:/ruggedemopool

File Edit View Terminal Go Help

```
#  
# zfs set compression=on ruggedemopool/dserban  
#  
# zfs get compression ruggedemopool/dserban  
NAME                PROPERTY    VALUE      SOURCE  
ruggedemopool/dserban  compression on         local  
#  
#
```

Hands-on with ZFS

Terminal - root@skynet:/rplugdemopool

File Edit View Terminal Go Help

```
#
# zfs snapshot rplugdemopool/dserban@rplugdemosnap
#
# ls -l /rplugdemopool/dserban/*
ls: cannot access /rplugdemopool/dserban/*: No such file or directory
#
# touch /rplugdemopool/dserban/fisier.txt
#
# ls -l /rplugdemopool/dserban/*
-rw-r--r-- 1 root root 0 May 12 13:25 /rplugdemopool/dserban/fisier.txt
#
# zfs rollback rplugdemopool/dserban@rplugdemosnap
#
# ls -l /rplugdemopool/dserban/*
ls: cannot access /rplugdemopool/dserban/*: No such file or directory
#
#
```

Hands-on with ZFS

Terminal - root@skynet:/r lugdemopool

File Edit View Terminal Go Help

```
#
# zfs clone r lugdemopool/dserban@r lugdemosnap r lugdemopool/clonadupasnapshot
#
# zfs promote r lugdemopool/clonadupasnapshot
#
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
r lugdemopool        168K  5.01G   24K    /r lugdemopool
r lugdemopool/clonadupasnapshot  21K  5.01G   21K    /r lugdemopool/clonadupasnapshot
r lugdemopool/dserban   15K  5.01G   21K    /r lugdemopool/dserban
#
# zfs unmount r lugdemopool/dserban
#
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
r lugdemopool        168K  5.01G   24K    /r lugdemopool
r lugdemopool/clonadupasnapshot  21K  5.01G   21K    /r lugdemopool/clonadupasnapshot
r lugdemopool/dserban   15K  5.01G   21K    /r lugdemopool/dserban
#
# zfs destroy r lugdemopool/dserban
#
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
r lugdemopool        126K  5.01G   24K    /r lugdemopool
r lugdemopool/clonadupasnapshot  21K  5.01G   21K    /r lugdemopool/clonadupasnapshot
#
#
```

Hands-on with ZFS

Terminal - root@skynet:/r lugdemopool

File Edit View Terminal Go Help

```
#  
# zfs list -H -t snapshot  
r lugdemopool/clonadupasnapshot@r lugdemosnap 0 - 21K -  
#  
#
```

Hands-on with ZFS

Terminal - root@skynet:/r lugdemopool

File Edit View Terminal Go Help

```
#
# zfs get all | grep share
r lugdemopool                sharenfs                off
default
r lugdemopool                shareiscsi              off
default
r lugdemopool                sharesmb                 off
default
r lugdemopool/clonadupasnapshot sharenfs                off
default
r lugdemopool/clonadupasnapshot shareiscsi              off
default
r lugdemopool/clonadupasnapshot sharesmb                 off
default
r lugdemopool/clonadupasnapshot@r lugdemosnap shareiscsi              off
default
#
#
```

Hands-on with ZFS

Terminal - root@skynet:/ruggedemopool

File Edit View Terminal Go Help

```
#
# zpool history
History for 'ruggedemopool':
2010-05-12.11:25:08 zpool create ruggedemopool mirror /dev/sda10 /dev/sda11 /dev/sda12
2010-05-12.11:28:04 zpool detach ruggedemopool /dev/sda12
2010-05-12.11:41:41 zpool add ruggedemopool spare /dev/sda12
2010-05-12.11:44:18 zpool offline ruggedemopool /dev/sda11
2010-05-12.11:51:56 zpool replace ruggedemopool /dev/sda11 /dev/sda13
2010-05-12.12:24:24 zpool scrub ruggedemopool
2010-05-12.12:24:55 zpool scrub ruggedemopool
2010-05-12.12:33:42 zfs create ruggedemopool/dserban
2010-05-12.12:45:03 zfs set compression=on ruggedemopool/dserban
2010-05-12.12:45:57 zfs set sharenfs=on ruggedemopool/dserban
2010-05-12.12:46:14 zfs set sharesmb=on ruggedemopool/dserban
2010-05-12.12:53:59 zfs set compression=on ruggedemopool/dserban
2010-05-12.13:05:41 zfs snapshot ruggedemopool/dserban@ruggedemosnap
2010-05-12.13:15:17 zfs rollback ruggedemopool/dserban@ruggedemosnap
2010-05-12.13:17:50 zfs destroy ruggedemopool/dserban@ruggedemosnap
2010-05-12.13:24:17 zfs snapshot ruggedemopool/dserban@ruggedemosnap
2010-05-12.13:25:32 zfs rollback ruggedemopool/dserban@ruggedemosnap
2010-05-12.13:29:31 zfs clone ruggedemopool/dserban@ruggedemosnap ruggedemopool/clonadupasnapshot
2010-05-12.13:30:13 zfs promote ruggedemopool/clonadupasnapshot
2010-05-12.13:32:03 zfs destroy ruggedemopool/dserban
```

```
#
```

```
# █
```

Hands-on with ZFS

- Remote backup of ZFS snapshot:

```
zfs send rlugdemopool/dserban@rlugdemosnap |  
ssh remotehost zfs receive remotezfs/backup
```

Întrebări / Feedback