

# Multi-machine mirroring choices

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2008-07/msg00207.html>

---

- *From:* Sven Willenberger <[sven@xxxxxxx](mailto:sven@xxxxxxx)>
  - *Date:* Tue, 15 Jul 2008 10:07:14 -0400
- 

With the introduction of zfs to FreeBSD 7.0, a door has opened for more mirroring options so I would like to get some opinions on what direction I should take for the following scenario.

Basically I have 2 machines that are "clones" of each other (master and slave) wherein one will be serving up samba shares. Each server has one disk to hold the OS (not mirrored) and then 3 disks, each of which will be its own mountpoint and samba share. The idea is to create a mirror of each of these disks on the slave machine so that in the event the master goes down, the slave can pick up serving the samba shares (I am using CARP as the samba server IP address).

My initial thought was to have the slave set up as an iscsi target and then have the master connect to each drive, then create a gmirror or zpool mirror using local\_data1:iscsi\_data1, local\_data2:iscsi\_data2, and local\_data3:iscsi\_data3. After some feedback (P.French for example) it would appear as though iscsi may not be the way to go for this as it locks up when the target goes down and even though I may be able to remove the target from the mirror, that process may fail as the "disk" remains in "D" state.

So that leaves me with the following options:

- 1) ggated/ggatec + gmirror
- 2) ggated/ggatec + zfs (zpool mirror)
- 3) zfs send/recv incremental snapshots (ssh)

1) I have been using ggated/ggatec on a set of 6.2-REL boxes and find that ggated tends to fail after some time leaving me rebuilding the mirror periodically (and gmirror resilvering takes quite some time). Has ggated/ggatec performance and stability improved in 7.0? This combination does work, but it is high maintenance and automating it is a bit painful (in terms of re-establishing the gmirror and rebuilding and making sure the master machine is the one being read from).

2) Noting the issues with ggated/ggatec in (1), would a zpool be better at rebuilding the mirror? I understand that it can better determine which drive of the mirror is out of sync than can gmirror so a lot of the "insert" "rebuild" manipulations used with gmirror would not be needed here.

## Multi-machine mirroring choices

3) The send/recv feature of zfs was something I had not even considered until very recently. My understanding is that this would work by a) taking a snapshot of master\_data1 b) zfs sending that snapshot to slave\_data1 c) via ssh on pipe, receiving that snapshot on slave\_data1 and then d) doing incremental snapshots, sending, receiving as in (a)(b)(c). How time/cpu intensive is the snapshot generation and just how granular could this be done? I would imagine for systems with little traffic/changes this could be practical but what about systems that may see a lot of files added, modified, deleted to the filesystem(s)?

I would be interested to hear anyone's experience with any (or all) of these methods and caveats of each. I am leaning towards ggate(dc) + zpool at the moment assuming that zfs can "smartly" rebuild the mirror after the slave's ggate processes bug out.

Sven

***Attachment: signature.asc***

*Description:* This is a digitally signed message part