

## Re: Creating RAID 1+0 with LVM

**Source:** <http://unix.derkeiler.com/Newsgroups/comp.unix.solaris/2005-08/0106.html>

---

**From:** Frank Cusack ([fcusack\\_at\\_fcusack.com](mailto:fcusack_at_fcusack.com))

**Date:** 08/02/05

Date: Mon, 01 Aug 2005 19:35:32 -0700

On Tue, 02 Aug 2005 00:12:42 GMT Darren Dunham <[ddunham@redwood.taos.com](mailto:ddunham@redwood.taos.com)> wrote:

> Neil Truby <[neil.truby@ardenta.com](mailto:neil.truby@ardenta.com)> wrote:

>>>> Obviously I'm missing something ...

>>>

>>> While the configuration syntax is 0+1, LVM actually creates 1+0 RAIDs.

>>> But don't believe me; you can easily test this.

>>>

>>> -frank

>

>> I do believe you ... but how can I test it? :-)

>

> Pull a disk. Watch that writes still hit both sides of the mirror where

> it remains redundant. Reconnect the disk and see that the resync is

> only for the slice with that disk, not the others.

I would have suggested, pull two drives, one on each side of the mirror. Of course the two drives must not be the two submirrors of the same part of the stripe. Observe that the md does not fail. This proves 1+0 behavior.

eg

stripe 1 = drives 1a, 1b, 1c

stripe 2 = drives 2a, 2b, 2c

md = mirror of stripe 1 + 2

pull drive 1a

pull drive 2b

If this were really 0+1, pulling drive 1a would fail half of the mirror (drives 1b and 1c would no longer be used at all). Pulling drive 2b would then fail the other half and the md would fail. But in 1+0 it would continue chugging along.

-frank