
Subject: Re: Re: openvz ploop images on moosefs mount
Posted by [Corin Langosch](#) on Mon, 02 Apr 2012 10:49:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Am 02.04.2012 12:19, schrieb Aleksandar Ivanisevic:

> Out of interest, have you tried running VEs off of moosefs, directly,
> without ploop? How does that work, if at all?

Yes, but only for a quick test and not with any production data:

- setup moosefs as usual and mount it on host
- create (sparse) file for container on moosefs mount
- use this file to setup a loop device using losetup
- make fs on loop device and mount it to the container's private folder
- start the container as usual

I don't really like this approach because it involves a lot of (slow) layers and double buffering issues and so I'd expect quite bad performance. If you try it yourself and do some benchmarks I'd be very happy to know the results.

I didn't try to use moosefs directly (without loop device) as a backing storage for containers because I made very bad experience (mostly performance and stability) with all kinds of network storages in the past.

From my past experience having a distributed block device (and not a distributed file system) is the way to go, as the system then has almost "native" performance due to the page cache (if not too many layers are involved as above).

I'm currently evaluating distributed network block devices, but unluckily there doesn't one to exist yet:

<http://serverfault.com/questions/375570/distributed-fault-tolerant-network-block-device/375821#375821>

Corin
