
Subject: Re: Using a layered filesystem as private dir?
Posted by [jjs - mainphrame](#) on Thu, 05 Jan 2012 19:07:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have postfix servers running on openvz and in general give it high marks, but there's little point in trying to make it something it's not. If you have the budget, the extra features of virtuoizzo are well worth the money. There are good reasons why virtuoizzo is more expensive, but if you can only afford openvz, my advice would be to let it do what it does best, and don't obsess over disk space which is a rather affordable commodity these days.

Joe

On Thu, Jan 5, 2012 at 10:52 AM, Kirill Korotaev <dev@parallels.com> wrote:

> As Scott mentioned we have VZFS in commercial version of Parallels
> Containers.
> It helps to save a lot of IOPS by sharing files between containers and is
> fully POSIX compliant.
>
> Thanks,
> Kirill
>
>
> On Jan 5, 2012, at 15:32 , Rick van Rein wrote:
>
> > Hello,
> >
> > I've just started using OpenVZ, and it feels more natural than the
> > alternatives I've seen -- my compliments!
> >
> > I can get a host running from a ZFS volume like /tank/vzdemo, which then
> > also gets shown at /var/lib/vz/vz-\$VEID. But what I really want to
> > do is use a layered FS (like aufs) as the private directory for the
> > container. But trying to do that leads to an error:
> >
> > bash# mount -t aufs -o br:/tank/vzdemo=rw:/tank/squeeze=ro none /mnt
> > bash# grep VE_ /etc/vz/conf/777.conf
> > VE_PRIVATE=/mnt
> > bash# vzctl create 777
> > Private area already exists in /mnt
> > Creation of container private area failed
> >
> > What is this trying to say? Is there a way to do what I am trying
> > to do? Did I understand well that the private area is a directory,
> > not a device?
> >
> >

> > Thanks,
> > -Rick
> >
> >
> > P.S. To capture any "why" questions :- I am trying to share as many
> > resources as possible. Containers beat Xen/KVM/VMware in that
> > respect, and when I can share the base OS and only have a thin
> > layer on top, it should mean that even the buffer cache is
> > shared between containers. It also means that upgrades can be
> > reduced to a minimum of repetition.
> >
