
Subject: Re: Using a layered filesystem as private dir?
Posted by pavel@pronkiy.ru **on** Thu, 05 Jan 2012 19:32:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

hi!

my configuration: openvz + squashfs4.0 + aufs
2.1-standalone.tree-32-20110228

/vz/squashfs/gentoo/ - mounted squashfs gentoo image or other distr
/vz/private/123/ - mount squashed gentoo dir and aufs layer from storage
dir /vz/storage/123/
/vz/private/124/ - mount squashed gentoo dir and aufs layer from storage
dir /vz/storage/124/
etc..

basic example:

```
hac_ostpl_sq_file = /vz/template/cache/gentoo.sq
hac_storage_private_dir = /vz/storage/vpsid/
hac_ostpl_squashfs_dir = /vz/squashfs/gentoo/
hac_private_dir = /vz/private/vpsid/
```

```
# start ve
mount -o loop $hac_ostpl_sq_file $hac_ostpl_squashfs_dir/
mount -t aufs -o
br=${hac_storage_private_dir}=rw:${hac_ostpl_squashfs_dir}=r o none
${hac_private_dir}
vzctl start vpsid
```

```
# stop ve
vzctl stop vpsid
umount /vz/private/vpsid
```

Making new ve:

1. generate /etc/vz/conf/vpsid.conf
 2. mount aufs
 3. vzctl start vpsid
- profit.

```
-=>> mount
/vz/template/cache/gentoo.sq on /vz/squashfs/gentoo type squashfs
(ro,relatime)
none on /vz/private/203 type aufs (rw,relatime,si=4b64c719)
/vz/private/203 on /vz/root/203 type simfs (rw,relatime)
proc on /vz/root/203/proc type proc (rw,relatime)
sysfs on /vz/root/203/sys type sysfs (rw,relatime)
rc-svcdir on /vz/root/203/lib/rc/init.d type tmpfs
```

```
(rw,nosuid,nodev,noexec,relatime,size=1024k,nr_inodes=131072 ,mode=755)
devpts on /vz/root/203/dev/pts type devpts
(rw,nosuid,noexec,relatime,gid=5,mode=620)
shm on /vz/root/203/dev/shm type tmpfs
(rw,nosuid,nodev,noexec,relatime,size=524288k,nr_inodes=1310 72)
...
...
```

work perfect on 2.6.32-rhel6 (patched by me and J.R. Okajima author aufs)
If you need support: pavel@pronkiy.ru

> 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 runnig 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.
>>
