
Subject: Re: nfs mount vs. mount --bind
Posted by [freedev](#) on Sun, 28 Jun 2009 00:08:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi all,

I have a hundred of containers that must mount the same identical NFS share list, so I have for example: 5 NFS mount on hardware node and 5 hundreds of bind inside the containers.
But I'm experiencing a strange behavior when I stop a container: during the execution of unmount script (in /vz/private/\$VEID/scripts/umount), not only the bind mount is unmounted but also the related NFS Volumes is unmounted.

So for example: if we have a NFS volume:

IPADDRESS:/VOLUME/share01 --> /usr/local/storage/share01

and a bind mount:

/usr/local/storage/share01 --> /vz/root/\$VEID/usr/local/mnt/share01

when we stop the container and umount /vz/root/\$VEID/usr/local/mnt/share01
also the NFS volume /usr/local/mnt/share01 is unmounted.

This problem is critical because if only one Container is stopped all remaining bind inside the containers become stale.

I found a workaround using the umount command with -l (lazy) parameter, so the bind is unmounted correctly without affect the NFS Volumes.

I apologize for a so long introduction, but the question is: could someone tell me if this behavior is related to OpenVZ?

The workaround we found it seems work well, someone think any counterindication or any suggestion?

Thanks in advance,
Vincenzo