

>
> Hi Keith,
> I'm quite sure XFS is the "problem".
> Kirill response is technically perfect and his knowledge is far ahead of mine, I would have not even imagined that disk quota might be implied. But I suggest you to have a try with ext4 or ext3, the both work fine, even with any tuning option you may need.

Ext4 (since you will be able to use ploop)

>
>
> On 31/05/2012 23:41, Keith Keller wrote:
>>
>> Hello all, thanks for the quick responses!
>>
>> On 2012-05-31, Kirill Kolyshkin<kolyshkin@gmail.com> wrote:
>>
>>>
>>> Assuming you don't have DISK_QUOTA=no in global config (ie /etc/vz/vz.conf),
>>>
>>
>> That's correct.
>>
>>
>>>
>>> the figures shown come from vzquota, and in order for vzquota to work correctly
>>> when you want to copy something to container you have to have it mounted (ie
>>> vzctl status 21 should show the word 'mounted' among others) and copy the
>>> data to VE_ROOT (ie /vz/root/21) but not to VE_PRIVATE (/vz/private/21).
>>>
>>> If you have already done it wrong (I assume you did *), you have to recalculate
>>> vzquota, the easiest way is to stop container, do vzquota drop 21 and
>>> start container
>>> again. This should fix your issue.
>>>
>>

```

>> You are indeed correct that I originally copied data to /vz/private/21.
>> But when I attempted to drop the quota, it still reports 0 blocks used.
>> I wonder if Massimiliano's comment is relevant?
>>
>> On 2012-05-31, Massimiliano
>> <massimiliano.sciabica@kiiama.com> wrote:
>>
>>>
>>> Hi,
>>> I had a similar issue when I first tried to improve performance of VPS
>>> with high number of small files.
>>> When a VPS reports 0% usage it is usually due to a not ext2 family
>>> filesystem. What fs where you using?
>>>
>>
>> I am using XFS. The FAQ mentions that disk quotas do not work with XFS,
>> so perhaps that's why it isn't displaying quite right (see below).
>>
>>
>>>
>>> I believe this is because of file system crash and the fsck (or journal
>>> replay)
>>> which truncated your files. In other words, this is not directly
>>> related to what you have described above.
>>>
>>
>> Perhaps--on boot, I didn't notice any unusual messages from fsck, though
>> I admit I wasn't paying an enormous amount of attention, and the logs
>> don't have anything interesting to report either. Does OpenVZ do a lot
>> of caching of disk writes from within a container? (It's obviously too
>> late now to see what xfs_repair thinks of the filesystem, but FWIW it
>> didn't find anything unusual.)
>>
>>
>>>
>>> Speaking of kernel crashes, it's nice to have some console logger
>>> installed,
>>> such as netconsole so whenever you have an oops you can report the bug.
>>> See http://wiki.openvz.org/Remote\_console\_setup
>>>
>>
>> Yes, I just set this up after the first crash--silly oversight not to
>> have done it right away. :)
>>
>>
>>>>
>>>> # Disk quota parameters (in form of softlimit:hardlimit)
>>>> DISKSPACE="1000G:2000G"

```

```
>>>>
>>>
>>> It looks like you have set disk quota values to more than your really
>>> have. Since this doesn't make sense
>>> my question is -- did you meant to disable disk space limit entirely?
>>> If yes, you can just have
>>> DISK_QUOTA=no in this config.
>>>
>>
>> Well, I am not entirely sure what I want, to be honest. If it's true
>> that having VE_ROOT and VE_PRIVATE on an XFS filesystem means disk
>> quotas don't work right, then perhaps I should either use ext3 (or
>> ext4?) on that filesystem, or disable disk quotas for all containers.
>>
>> As an experiment on the latter, I set DISK_QUOTA=no in vz.conf, and now
>> I get:
>>
>> # vzctl exec 21 df -h
>> Filesystem      Size  Used Avail Use% Mounted on
>> /dev/simfs       1.0T  332G  693G   33% /
>> none             16G   4.0K   16G    1% /dev
>>
>> But it would be convenient to have disk quotas. Is there a preference
>> for ext3 or ext4 for the host filesystem?
>>
>> --keith
>>
>>
>>
>
```
