
Subject: ploop for snapshotting openvz containers on amazon ec2/xen

Posted by [fredish](#) on Mon, 13 May 2013 18:49:12 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

I have started using ploop with openvz successfully on centos 6 for snapshotting containers, and it works marvelously. However in moving to Amazon and running on centos 6 via pv-grub, the basic openvz functionalities work per normal, but not with ploop. I believe Xen is doing something unexpected with the locking. I haven't been able to google up any specific info for ec2/xen + openvz + ploop. Is anyone aware of any gotchas or have pointers?

regards, Fred Patton

When running the following script ovz-container-setup.sh--
[ec2-user@blah ~]\$ sudo ./ovz-container-setup.sh 201

which consists of:

```
sudo vzctl create $1 --ostemplate ubuntu-12.04-x86_64 --config unlimited --layout ploop
--diskspace 3G
sudo vzctl set $1 --hostname box$1 --ipadd 10.0.0.$1 --nameserver 10.0.0.1 --userpasswd
root:jl4j3298 --onboot yes --save
sudo vzctl start $1
sudo vzctl snapshot $1 --id some-guid
```

OUTPUT:

```
Creating image: /vz/private/201.tmp/root.hdd/root.hdd size=3145728K
Creating delta /vz/private/201.tmp/root.hdd/root.hdd bs=2048 size=6291456 sectors
Storing /vz/private/201.tmp/root.hdd/DiskDescriptor.xml
Adding delta dev=/dev/ploop36619 img=/vz/private/201.tmp/root.hdd/root.hdd (rw)
mke2fs 1.42.3 (14-May-2012)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
196608 inodes, 785915 blocks
39295 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=805306368
24 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912
```

Allocating group tables: done
Writing inode tables: done
Creating journal (16384 blocks): done
Writing superblocks and filesystem accounting information: done

tune2fs 1.42.3 (14-May-2012)

Creating balloon file .balloon-c3a5ae3d-ce7f-43c4-a1ea-c61e2b4504e8

Mounting /dev/ploop36619p1 at /vz/private/201.tmp/root.hdd/root.hdd.mnt fstype=ext4 data="

Unmounting device /dev/ploop36619

I/O error : No such file or directory

/vz/private/201.tmp/root.hdd/DiskDescriptor.xml:1: parser error : Document is empty

^

/vz/private/201.tmp/root.hdd/DiskDescriptor.xml:1: parser error : Start tag expected, '<' not found

^

Removing stale lock file /vz/lock/201.lck

Warning: distribution not specified in CT config, using defaults from /etc/vz/dists/default

Starting container...

Warning: distribution not specified in CT config, using defaults from /etc/vz/dists/default

stat(/vz/private/201): No such file or directory

stat(/vz/private/201): No such file or directory

Can't umount /vz/root/201: Invalid argument

stat(/vz/private/201): No such file or directory

Unable to start init, probably incorrect template

Container start failed

Killing container ...

Container was stopped

stat(/vz/private/201): No such file or directory

stat(/vz/private/201): No such file or directory

Can't umount /vz/root/201: Invalid argument

Error: failed to apply some parameters, not saving configuration file!

Container private area /vz/private/201 does not exist

Snapshot feature is only available for ploop-based CTs

ADDITIONAL CONTEXT:

[ec2-user@blah ~]\$ sudo dmesg | tail -n 50

[0.211801] device-mapper: ioctl: 4.22.6-ioctl (2011-10-19) initialised: dm-devel@redhat.com

[0.237828] <30>udev[87]: starting version 173

[0.292946] xlblk_init: register_blkdev major: 202

[0.296020] alloc irq_desc for 275 on node 0

[0.296026] alloc kstat_irqs on node 0

[0.305178] blkfront: xvde1: barriers disabled

[0.606316] EXT4-fs (xvde1): INFO: recovery required on readonly filesystem

[0.606345] EXT4-fs (xvde1): write access will be enabled during recovery

[6.305800] EXT4-fs (xvde1): orphan cleanup on readonly fs

```

[ 6.308531] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7113
[ 6.308619] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7109
[ 6.309350] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7076
[ 6.309370] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7074
[ 6.309404] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7072
[ 6.309426] EXT4-fs (xvde1): ext4_orphan_cleanup: deleting unreferenced inode 7068
[ 6.309446] EXT4-fs (xvde1): 6 orphan inodes deleted
[ 6.309460] EXT4-fs (xvde1): recovery complete
[ 6.392109] EXT4-fs (xvde1): mounted filesystem with ordered data mode. Opts:
[ 6.664406] dracut: Remounting /dev/disk/by-label/x2f with -o noatime,ro
[ 6.685461] EXT4-fs (xvde1): mounted filesystem with ordered data mode. Opts:
[ 6.691476] dracut: Mounted root filesystem /dev/xvde1
[ 6.763928] dracut: Switching root
[ 12.227947] <30>udev[224]: starting version 173
[ 14.594784] Initialising Xen virtual ethernet driver.
[ 14.595974] alloc irq_desc for 274 on node 0
[ 14.595980] alloc kstat_irqs on node 0
[ 17.844949] NET: Registered protocol family 10
[ 29.386043] eth0: no IPv6 routers present
[ 36.884041] venet0: no IPv6 routers present
[ 49.671537] ip_tables: (C) 2000-2006 Netfilter Core Team
[ 50.059239] ip6_tables: (C) 2000-2006 Netfilter Core Team
[ 50.153451] Enabling conntracks and NAT for ve0
[ 50.153466] nf_conntrack version 0.5.0 (16384 buckets, 65536 max)
[ 50.458267] RPC: Registered named UNIX socket transport module.
[ 50.458279] RPC: Registered udp transport module.
[ 50.458285] RPC: Registered tcp transport module.
[ 50.458291] RPC: Registered tcp NFSv4.1 backchannel transport module.
[ 51.065970] Slow work thread pool: Starting up
[ 51.066120] Slow work thread pool: Ready
[ 51.066160] FS-Cache: Loaded
[ 51.121922] Registering the id_resolver key type
[ 51.122080] FS-Cache: Netfs 'nfs' registered for caching
[ 51.221254] ploop_dev: module loaded
[ 75.822126] ploop36619: unknown partition table
[ 76.882564] ploop36619:
[ 76.900774] ploop36619: p1
[ 80.539312] ploop36619: p1
[ 80.567750] EXT4-fs (ploop36619p1): mounted filesystem with ordered data mode. Opts:
[ 80.643295] CT: 201: started
[ 80.737492] CT: 201: stopped

```

```
[ec2-user@blah ~]$ sudo df -iT
```

```

Filesystem Type Inodes IUsed IFree IUse% Mounted on
/dev/xvde1 ext4 524288 57886 466402 12% /
tmpfs tmpfs 936809 1 936808 1% /dev/shm

```

```
[ec2-user@ip-10-254-73-168 ~]$ sudo ploop stat -d /dev/ploop34720
```

bio_in	16786
bio_fast	120
bio_full	0
bio_out	28
bio_alloc	75
bio_alloc_whole	64
bio_splits	0
coal_back	0
coal_forw	16547
coal_back2	0
coal_forw2	0
coal_oback	0
coal_oforw	0
coal_mback	0
coal_mforw	0
coal_overlap	0
coal_flush	0
bio_barriers	0
bio_rzero	13
bio_wzero	0
bio_syncwait	75
bio_fsync	19
bio_cows	0
bio_whole_cows	0
merge_neg_cluster	0
merge_neg_disable	0
fast_neg_nomap	116
fast_neg_noem	0
fast_neg_shortem	0
fast_neg_backing	0
bio_lockouts	28
map_lockouts	1
merge_lockouts	0
map_reads	4
map_merges	0
map_single_writes	14
map_multi_writes	10
map_multi_updates	61
bio_trans_whole	0
bio_trans_copy	0
bio_trans_alloc	0
bio_trans_index	0
bio_flush_in	9
bio_fua_in	1
bio_flush_out	9
bio_fua_out	1
bio_flush_skip	0

