
Subject: Is ploop stable/Any progress on quotas for simfs

Posted by [seanfulton](#) on Tue, 15 Oct 2019 23:12:31 GMT

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

When ploop first came out we used it from lost many, many containers to corrupt ploop filesystems. We went back to simfs. Now looking at upgrading some nodes to OpenVZ 7 and it looks like ploop is the only way to get quotas for containers.

So for those who have upgraded, how stable is it?

Is there any progress on first-level quotas for SIMFS???

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [ccto](#) on Wed, 23 Oct 2019 12:19:19 GMT

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

We have servers running with both OpenVZ 6, OpenVZ 7. Gradually migrating to OpenVZ 7 platform

We use ploop, snapshot (for backup), compact. ploop is OK.

We do not receive report about loss of files.

simfs inside OpenVZ 7 does not support 2nd level quota.

If guest is CentOS 5, ploop does not support 2nd level quota too in OpenVZ 7 host.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [seanfulton](#) on Wed, 23 Oct 2019 19:41:26 GMT

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

This is very helpful, thank you. I have not been able to get SIMFS containers to have a container quota, I thought that was a first-level quota. It says it's supported, but when I create a VE in SIMFS and enter it, df shows the whole filesystem, not just the little slice (like it used to in OpenVZ 6). Is that normal?

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [websavers](#) on Sun, 17 Nov 2019 19:07:48 GMT

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

We've also been using OpenVZ 7 with ploop (transferred most containers from OpenVZ6 SIMFS) for about 2 years now. We were a bit hesitant to use ploop because of those past reports on these forums about recovery and such, however there hasn't yet been a single data consistency problem that the virtuozzo subsystem hasn't auto-repaired for us upon boot of the container. Granted we use RAID 5 and have replaced multiple disks with full parity recovery by the RAID

controller and not a single problem with booting containers yet. Fingers' crossed that remains true.

Upsides: because ploop containers are singular image files on the node, migrating containers has never been faster (full network bandwidth), which is a huge upside. Quota works great within the containers as well.

The biggest downsides of ploop over SIMFS are:

1. BACKUPS (A) If you want quick restore from backup capability, your backup script needs to snapshot, (to create a disk not currently locked) and you need to sacrifice incremental backups and compression, thus backups take up a huge amount of space. (B) If you want file-by-file compressable, dedupable, incremental backups (like using borg backup) that can be used to restore an entire container / node then you need to ensure that your restore script accomodates for container (and ploop) creation, mount it, and then restore the data (This is preferred by far if you ask me, for backup storage costs alone, but takes much more work to implement).
2. There's a fair amount of storage overhead used by ploop, even with its nightly compact system running. On a node with 1.7TB total usable storage, where 90% of it is used by around 20 ploop containers, about 10% of that 90% is overhead -- that makes for around 150GB- 200GB of wasted space. I've analyzed a number of different nodes and they all have similar overhead. I don't expect zero overhead from such a system, but I would definitely like to see that down to less than 5%.

The backup thing isn't *that* big of a deal, but it's definitely a bit more involved than SIMFS where you simply create the container and pop your files in place.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [ccto](#) on Mon, 18 Nov 2019 02:30:35 GMT

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

1. You may mount the ploop snapshot to some mountpoint (e.g.) /mnt, and make the incremental file-based backup of that mountpoint. Then, you can restore it file by file.

2. Yes, from our experience, ploop contains overhead (yes, around 10%).

There is a pcompact tools, which somehow compact the ploop file automatically upon certain threshold.

Comparing to other, like KVM qcow2, there are also some storage overhead too.

From my personal experience, if you have a number of containers, each container contains a number of files, ploop performs faster than simfs over time.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [websavers](#) on Mon, 18 Nov 2019 14:05:18 GMT

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

ccto wrote on Sun, 17 November 2019 22:30:1. You may mount the ploop snapshot to some mountpoint (e.g.) /mnt, and make the incremental file-based backup of that mountpoint. Then, you can restore it file by file.

Indeed! That's what we do. Just pointing out that it's more involved than SIMFS

ccto wrote on Sun, 17 November 2019 22:30: Comparing to other, like KVM qcow2, there are also some storage overhead too.

For sure, but this is a comparison with SIMFS -- just making it clear what kinds of differences can be expected. I think the benefits outweigh the downsides, but that might not be the case for everyone.

ccto wrote on Sun, 17 November 2019 22:30: From my personal experience, if you have a number of containers, each container contains a number of files, ploop performs faster than simfs over time.

That may be the case; I haven't done a direct comparison on the same hardware to know for sure. One thing we can be sure of is that it's definitely more performant migrating ploop containers as it doesn't need to xfer file by file.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [seanfulton](#) on Fri, 10 Jan 2020 20:31:23 GMT

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

I want to reply to my own thread here. I built several OpenVZ servers and migrated my old simfs containers to ploop. Not really trusting it, but not having much choice.

Last week I needed to expand a ploop container. I did so and it filled up the physical drive that it was on. Now, that's not a great situation, given that I allocated 4500G to the container and the drive was 5.4T, it should not have gotten full.

But, I spent three days trying to fix it. I resized, compacted, etc. I beat the hell out of it.

Each time the activity would hang forcing a reboot. I am happy to say that the ploop drive was still completely mountable to the end. It never got corrupted, never had an error with that.

So at this point I would say ploop is pretty reliable and stable.

sean

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [websavers](#) on Sat, 11 Jan 2020 15:57:45 GMT

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

Hey Sean, that's some great testing! Thankfully we haven't experienced the resizing issues -- all of our containers are 500GB or less, so that may explain it. We *do* experience the overhead (everyone using ploop must) and I wish that could be improved a bit, even if it could be reduced to 5% rather than ~10%.

For our <500GB containers, resizing and compacting using the following command has *not* resulted in the need to restart the container.

```
prl_disk_tool compact --hdd /vz/private/$CTID/root.hdd/
```

That said our results agree on the lack of corruption in ploop drives so far... it's been a little over 2 years since we switched to ploop and so far not a single corruption has occurred. We'll keep our fingers crossed!

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [seanfulton](#) on Sat, 11 Jan 2020 17:40:06 GMT

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

What sort of overhead are you seeing with PLOOP containers? Right now I am migrating a CENTOS6 container and the container's disk usage shows 500M but df on the drive it is on is showing 1.2G of usage.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [websavers](#) on Sat, 11 Jan 2020 21:28:20 GMT

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

seanfulton wrote on Sat, 11 January 2020 13:40 What sort of overhead are you seeing with PLOOP containers? Right now I am migrating a CENTOS6 container and the container's disk usage shows 500M but df on the drive it is on is showing 1.2G of usage.

Here's a few containers we currently have for comparison. The first number is the total storage of all files in the container.

- WS2-163: 64.68 GB of 150 GB used -- du reports 78GB = 17% overhead
- WS2-230: 65.38 GB of 100 GB used -- du reports 77GB = 15% overhead
- WS2-253: 2.55 GB of 15 GB used -- du reports 4.3GB = 41% overhead
- WS2-301: 45.06 GB of 75 GB used -- du reports 54GB = 17% overhead

Assuming 41% is an outlier, and/or more likely to be a problem on smaller containers (which I believe is accurate from what I've seen), we're talking an average of 16% overhead. If you've got a 2TB drive, that means you've lost over 327GB to ploop overhead which in the above examples could easily be used for 4-6 more containers. Whereas if it were closer to 5% you'd only be losing 100GB, which would be an easier loss to stomach, given the advantages of ploop.

Subject: Re: Is ploop stable/Any progress on quotas for simfs

Posted by [HHawk](#) on Thu, 16 Jan 2020 10:55:55 GMT

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

websavers wrote on Sat, 11 January 2020 21:28seanfulton wrote on Sat, 11 January 2020 13:40What sort of overhead are you seeing with PLOOP containers? Right now I am migrating a CENTOS6 container and the container's disk usage shows 500M but df on the drive it is on is showing 1.2G of usage.

Here's a few containers we currently have for comparison. The first number is the total storage of all files in the container.

- WS2-163: 64.68 GB of 150 GB used -- du reports 78GB = 17% overhead
- WS2-230: 65.38 GB of 100 GB used -- du reports 77GB = 15% overhead
- WS2-253: 2.55 GB of 15 GB used -- du reports 4.3GB = 41% overhead
- WS2-301: 45.06 GB of 75 GB used -- du reports 54GB = 17% overhead

Assuming 41% is an outlier, and/or more likely to be a problem on smaller containers (which I believe is accurate from what I've seen), we're talking an average of 16% overhead. If you've got a 2TB drive, that means you've lost over 327GB to ploop overhead which in the above examples could easily be used for 4-6 more containers. Whereas if it were closer to 5% you'd only be losing 100GB, which would be an easier loss to stomach, given the advantages of ploop.

Similar results here. One time we even have an overhead of almost 1.5 TB, because of resizing the ploop (backup) container.

I still prefer SimFS, but alas we have no real choice. Never had these issues with SimFS. Always had the space which was in use.