

---

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

Posted by [wsap](#) 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 virtio 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 accommodates 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.

---