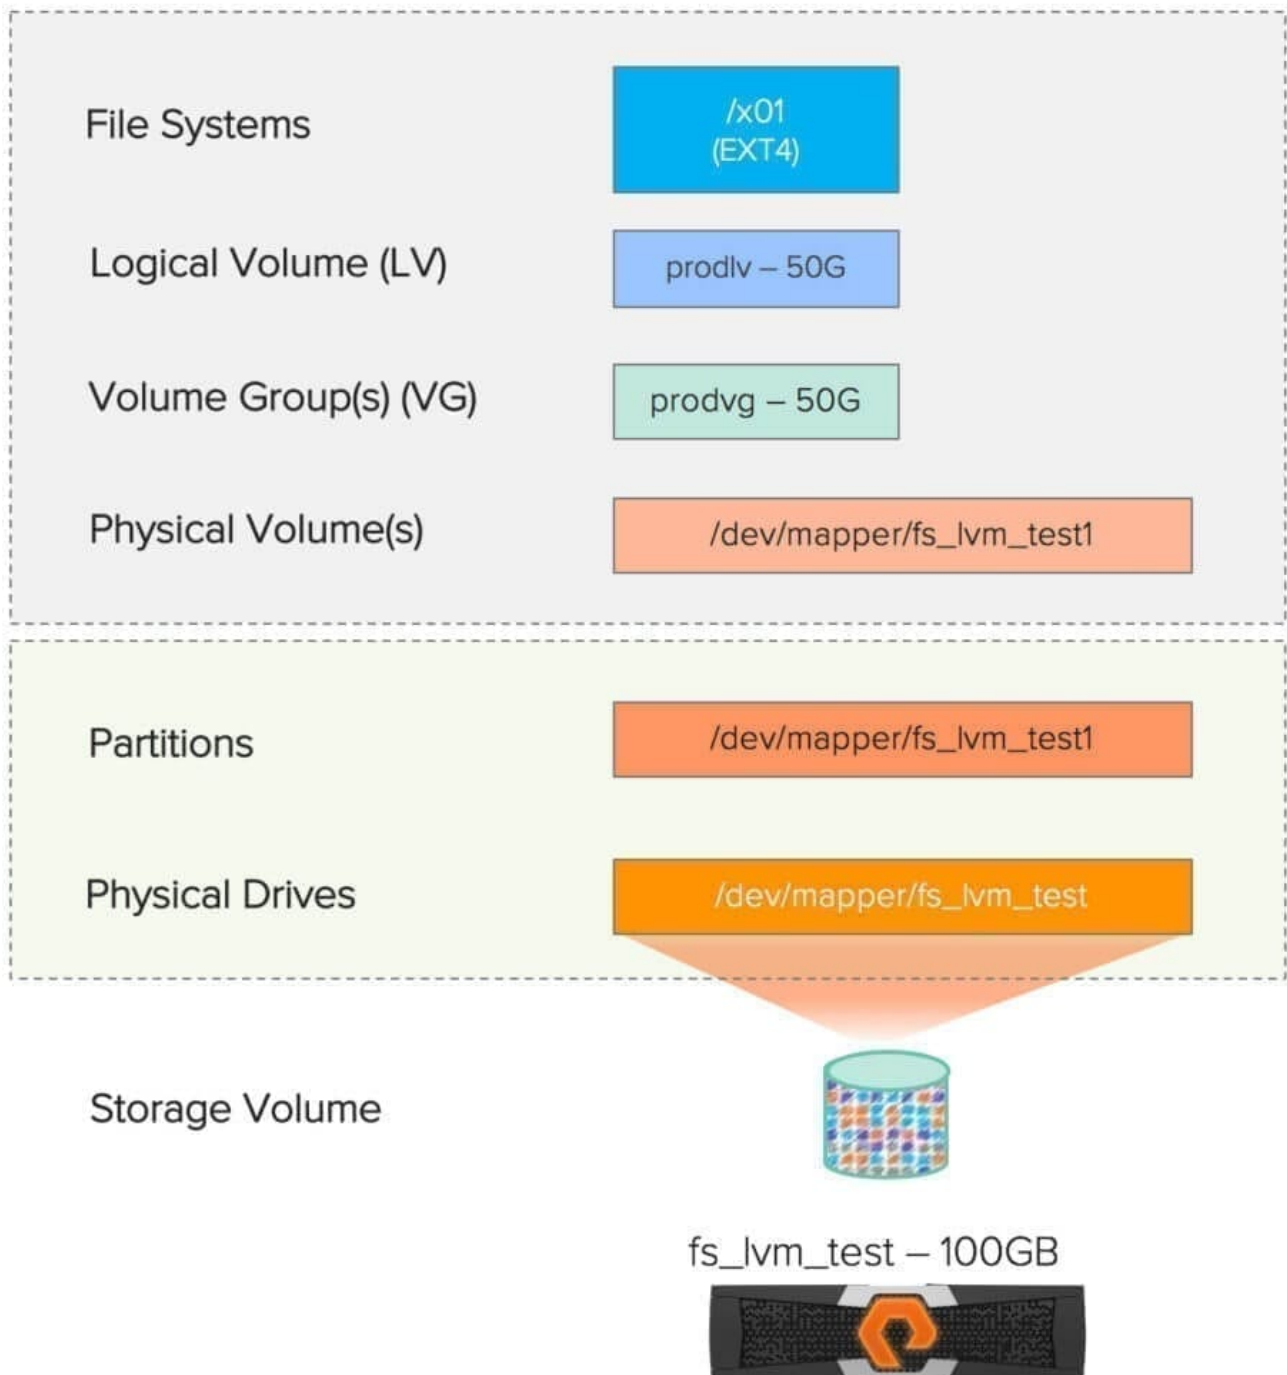


# Cloning LVM Volumes Using Pure FlashRecover Snapshots



As explained in my earlier [post](#) Pure FlashRecover [Snapshots](#) is one of the most used feature within Purity by our customers who use these snapshots to clone their databases. In case of filesystems, the volumes that are cloned may even be part of LVM (Logical Volume Manager).

In Linux world, when volume(s) part of LVM is cloned using storage level snapshots, it doesn't pose an issue to mount them on another host than the source. The challenge is when the cloned volumes that are part of LVM has to be mounted on the same host. Cloned volume that is part of LVM cannot coexist on the same host as system will throw duplicate errors on physical volume (PV) and volume group (VG) as they will have same UUID.

Is it possible to attach the cloned volume that is part of LVM to the same host? Yes, this is very much possible by using the **vgimportclone** command. The command `vgimportclone` takes care of renaming the volume group and changing the associated UUID with PV and VG.

Here is the high level steps to clone a volume that is part of LVM and attaching it to the same host.

1) Clone the volume and attach to the same host and rescan.

2) Make sure the new lun shows up when you run `multipath -ll` (or `multipath -d -l`)

(Note: If you use the alias in `/etc/multipath.conf`, do not include the cloned volume as device multipathing ignores the mapping. I did not spend time figuring out why this is the case and if anyone knows the reason, please update me)

3) `vgimportclone -n <clonedvgname> /dev/mapper/<clonedlun>`

4) `pvscan -cache`

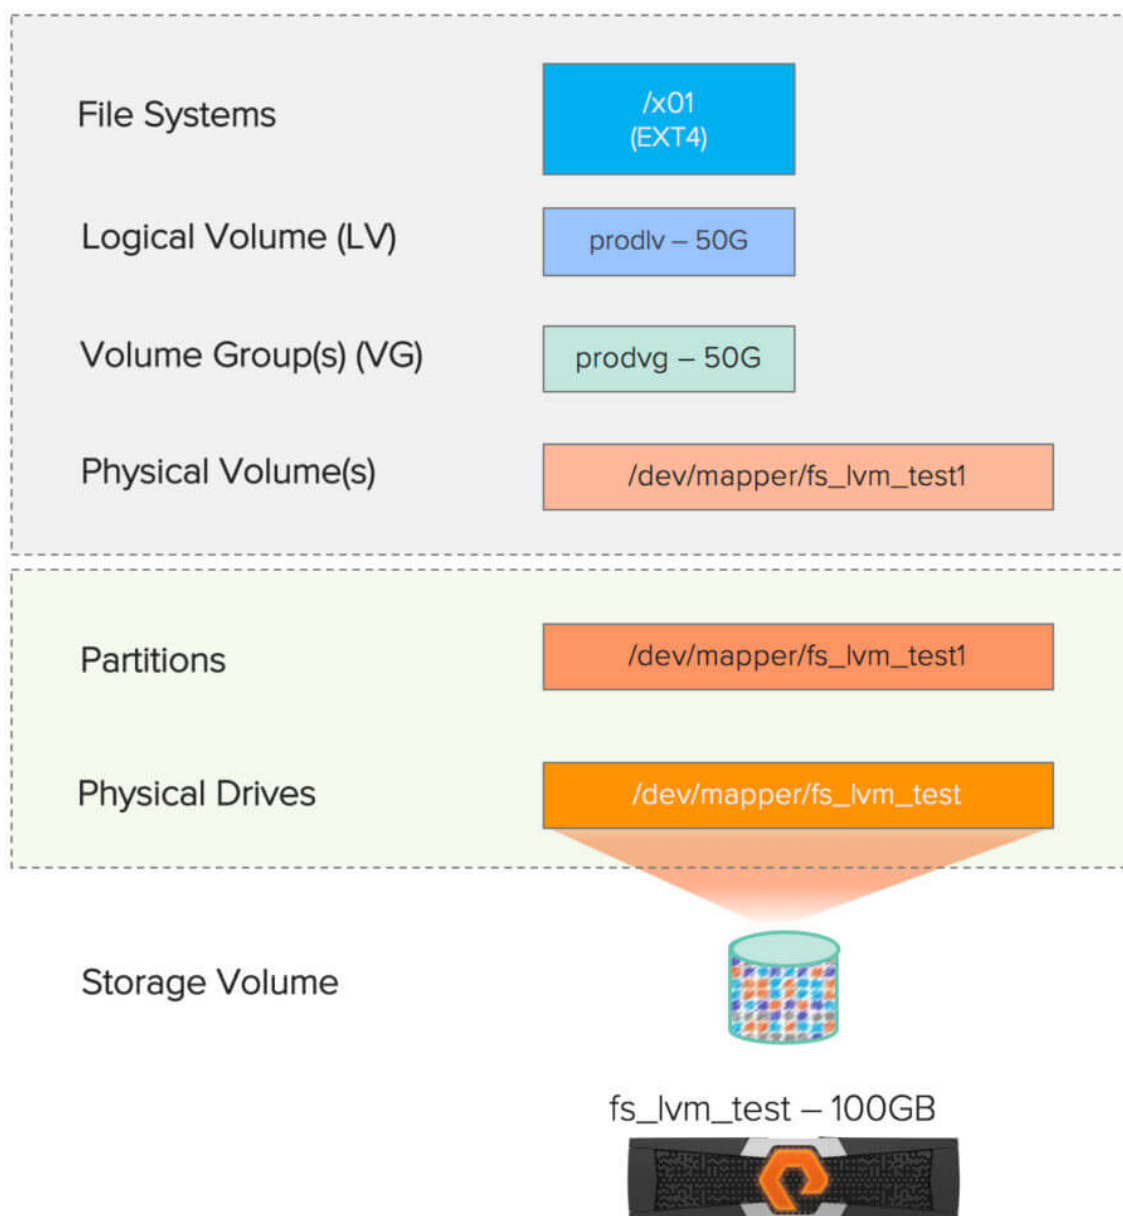
5) `vgchange -ay <clonedvgname>`

6) `pvdisplay` (should show both the volumes now)

7) If needed, use `lvrename` to change the logical volume which might have the same name as the source.

Similar to the filesystem [clone](#) procedure, make sure to unmount the filesystem when overwriting this volume for periodic refresh.

Here is an example showcasing the sequence of steps to clone a volume that is part of LVM and attaching it to the same host. The following picture shows the existing source LVM details on the Linux host. The volume `fs_lvm_test` from Pure FlashArray is connected to the host and the partition `/dev/mapper/fs_lvm_test1` is configured as the PV on which the `prodv` was created with the logical volume `prodlv` of size 50G which is mounted as EXT4 filesystem on `/x01`.



Here are the contents of the filesystem before taking the snapshot.

[crayon-6515bbef72337150897221/]

## FlashRecover Snapshot 101

Using Pure FlashArray GUI we took the snapshot of the volume fs\_lvm\_test and copied into fs\_lvm\_clone and attached it to the host. You can certainly automate this process by taking snapshot through CLI or REST based APIs.

Checking the volume detail of fs\_lvm\_clone showed 3624a93708c44af151613837200013c35 as its serial number.

Rescanned the host for new SCSI devices and checked the multipath output to see the cloned volume.

[crayon-6515bbef72345080406617/]

At this point, ran vgimportclone with the new volume group (qavg) passing the physical volume /dev/mapper/mpathe1

[crayon-6515bbef72349253370861/]

## More on Snapshots

Ran pvscan with -cache option to refresh the LVM cache information.

[crayon-6515bbef7234b974780276/]

Activated the volume group using vgchange with -ay option.

[crayon-6515bbef72356953228384/]

Running pvdisplay showed the cloned volume (/dev/mapper/mpathe1) as a PV.

[crayon-6515bbef72358053811797/]

The lvdisplay showed the logical volume from the cloned volume which had the same name as source (prodlv).

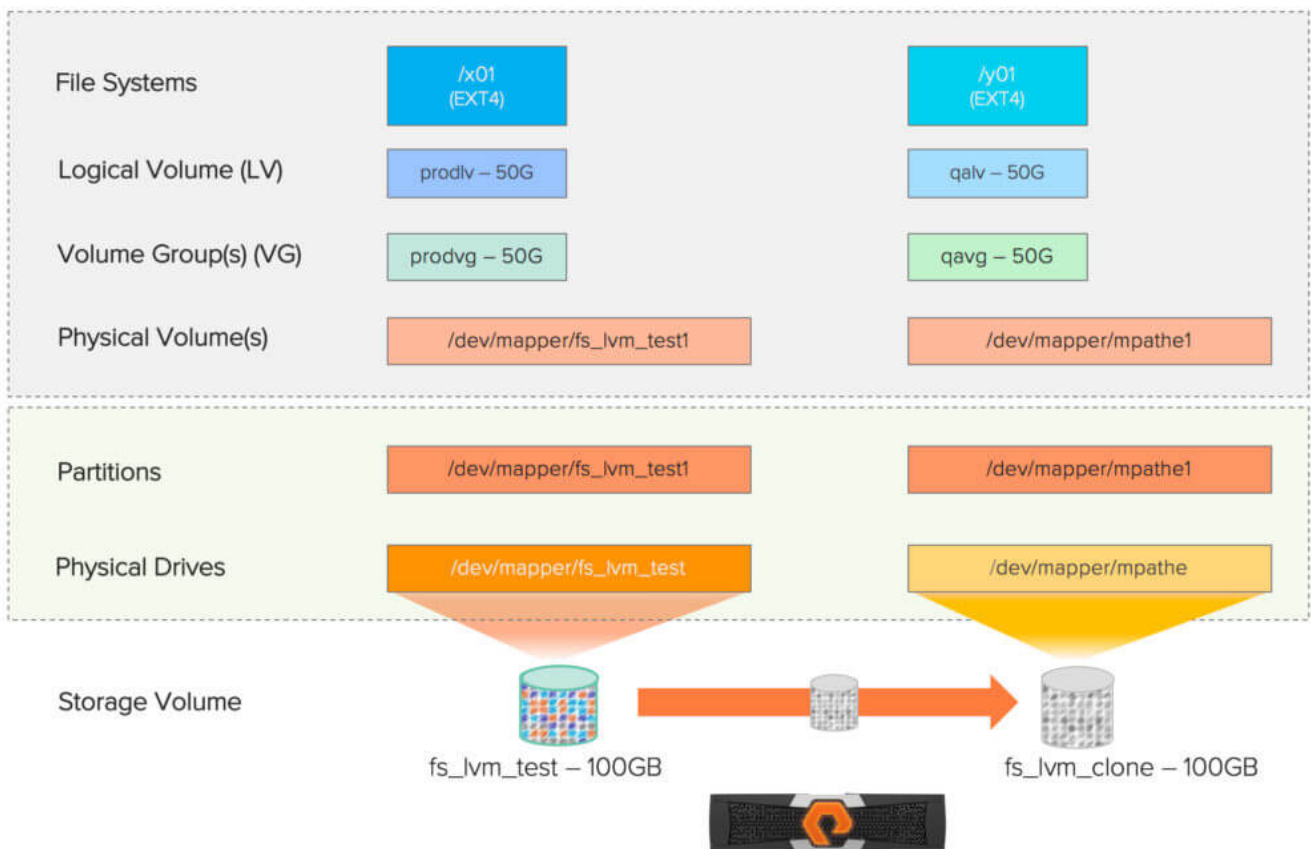
[crayon-6515bbef7235a931988341/]

To avoid confusion, we went ahead and changed the logical volume name in the cloned LVM to qalv.

Also mounted the logical volume as EXT4 filesystem on /y01 to validate the content.

[crayon-6515bbef7235c128498861/]

Here is what the system sees after cloning and mounting onto the same host.



As you can see, it is not only easy to clone the volumes that are part of LVM but it is also easier to mount them on to the same host without any issues.

More about [FlashRecover](#)