

Move OSDs disks between hosts

reinstall PVE

After reinstallation of PVE, CEPH automatically detects OSD but keeps it as down. To run OSD it is need to: `ceph-volume lvm activate --all`

1st info

If you used `ceph-deploy` and/or `ceph-disk` to set up these OSDs (that is, if they are stored on labeled GPT partitions such that `upstart` is automagically starting up the `ceph-osd` daemons for you without you putting anything in `/etc/fstab` to manually mount the volumes) then all of this should be plug and play for you--including step #3. By default, the startup process will 'fix' the CRUSH hierarchy position based on the hostname and (if present) other positional data configured for 'crush location' in `ceph.conf`. The only real requirement is that both the `osd` data and journal volumes get moved so that the daemon has everything it needs to start up.

2nd info

[move osd to another node](#)

if you don't have a separate (shared) journal device, you can just down and out the OSD and physically move it from one host to another. (hot-)plugging it in should automatically start the OSD service on the new host, and you can mark the `osd` as "in" on the GUI.

if you have a separate journal device only used for the particular `osd`, you can move it together with the OSD and it should work as well.

if you have a shared journal device, you might be able to move all the OSDs using it and the journal, all at the same time - I have never tried this though.

if you only want to move one of several OSDs sharing a journal, this is AFAIK not possible, and you need to actually remove it and create a new OSD from scratch on the new host, using the moved disk.

From:

<https://niziak.spox.org/wiki/> - **niziak.spox.org**

Permanent link:

https://niziak.spox.org/wiki/vm:proxmox:ceph:move_disks

Last update: **2022/07/25 09:41**

