LVM + dm-cache (unencrypted)

LVM dm-cache is a feature of the Linux device mapper, which uses a fast storage device to boost data read/write speeds of a slower one. It achieves this by transparently copying blocks of frequently accessed data to the faster storage device in the background. On subsequent reads/writes the faster storage device is queried first. If the requested data blocks are not on there, it automatically falls back on the slower source storage device.

This makes it possible to combine the benefits of SSD speeds with the low cost and high storage capacity of HDDs, when comparable pure SSD-based storage with the same capacity is too expensive or otherwise unavailable.

> **NOTE:** This partition scheme is tailored towards a desktop computer setup with enough RAM and no SWAP (and therefore no hibernate/suspend-to-disk support).

> **CAUTION:** This setup does **NOT** utilize LUKS disk encryption.

This guide assumes the following:

- `/dev/nvme0n1` is the primary disk (cache device)
- `/dev/sda` is the secondary disk (origin device)

Nomenclature

Term	Description
Physical Volume (PV)	On-disk partitioning format to be combined in a VG to a common storage pool
Volume Group (VG)	Grouping of one or more PVs to provide a combined storage pool from which storage can be requested in the form of LVs.
Logical Volume (LV)	Logical partition format which can be accessed like a block device to hold file systems and data.

Preparing the cache device

First the available disks need to be determined. This can easily be achieved with `fdisk`:

```
fdisk -l
```

To start the actual partitioning process start `cfdisk` and point it to the **disk** you wish to partition:

> **WARNING:** Make sure to select your actually desired device!

```
cfdisk /dev/nvme0n1
```

Partition the disk in the following way:

FS Type	Size	Mount Point	Comment
vfat	1G	/boot	EFI System
LVM	(remaining)		Linux LVM

Preparing the origin device

Partition the disk by starting `cfdisk` and pointing it to the **disk** for the origin device:

> **WARNING:** Make sure to select your actually desired device!

```
cfdisk /dev/sda
```

Partition the disk in the following way:

FS Type	Size	Mount Point	Comment
LVM	(all)		Linux LVM

Creating physical volumes, volume group and logical volumes

To create physical volumes as the basis for the LVM setup, use `pvcreate` and point it to the **partitions** you created in the two previous steps:

```
pvcreate /dev/nvme0n1p2    # SSD
pvcreate /dev/sda1         # HDD
```

Continue by creating a volume group with `vgcreate` that spans both physical volumes you just created:

> **NOTE:** `vg0` is used as an example here. Use whatever you like.

```
vgcreate vg0 /dev/nvme0n1p2 /dev/sda1
```

Next, create logical volumes inside the volume group with `lvcreate`, using 100% of the available space on the HDD and specifying the cache pool on the SSD:

```
lvcreate -l 100%FREE -n lv_root vg0 /dev/sda1
lvcreate --type cache-pool -n lv_cache -l 100%FREE vg0 /dev/nvme0n1p2
```

Finally, link the cache pool to the origin device with `lvconvert`:

```
lvconvert --type cache --cachepool vg0/lv_cache vg0/lv_root
```

Formatting devices

Format the partitions with the appropriate `mkfs` subcommand:

```
mkfs.fat -F 32 /dev/nvme0n1p1       # EFI System Partition
mkfs.btrfs /dev/mapper/vg0-lv_root  # Btrfs root file system
```

Mount the root Btrfs file system:

```
mount /dev/mapper/vg0-lv_root /mnt
```

Next, create the subvolumes with the `btrfs` user space tools:

```
btrfs subvolume create /mnt/@
btrfs subvolume create /mnt/@home
```

Unmount the root file system again:

```
umount -R /mnt
```

Mount the `@` subvolume at `/mnt`:

```
mount /dev/mapper/vg0-lv_root -o noatime,compress-force=zstd,space_cache=v2,subvol=@ /mnt
```

Create directories for subsequent mount points:

```
mkdir -p /mnt/{boot,home}
```

Mount the remaining file systems:

```
mount /dev/nvme0n1p1 /mnt/boot
mount /dev/mapper/vg0-lv_root -o noatime,compress-force=zstd,space_cache=v2,subvol=@home /mnt/home
```

Revision #12
Created 1 April 2021 14:44:16 by Sebin
Updated 6 March 2024 07:18:50 by Sebin