All You Need to Know About SUSE Linux Enterprise Server 12 Memory Optimization

Sander van Vugt

Author, training and consultant

sandervanvugt.nl

Virtual versus Resident Memory

Understanding Virtual Memory

What would you do
with 35 TB of RAM?

Understanding Swap

Consider this:

- from **top**

```
KiB Mem: 8000M total,7900M used, 100M free, 40M buffers
KiB Swap: 1000M total, 1000M used, 0 free, 300M cached
```

- from /proc/meminfo

```
Active: 5000M
Inactive: 3000M
Active (anon): 1000M
Inactive (anon): 4000M
Active (file): 300M
Inactive (file): 0M
```

Consider this:

- from **top**

```
KiB Mem: 8000M total,7900M used, 100M free, 40M buffers
KiB Swap: 1000M total, 1000M used, 0 free, 300M cached
```

- from /proc/meminfo

```
Active: 5000M
Inactive: 3000M
Active (anon): 1000M
Inactive (anon): 4000M
Active (file): 300M
Inactive (file): 0M
```

DEMO

Memory Over-Allocation and OOMM

DEMO

Huge Pages

Using Bcache

Speeding up Cache with bcache

- Use SSD disk to speed up cache and in fact create a hybrid disk environment
  - New feature, available since 3.10 kernels
  - Install bcache-tools to use it

Speeding up Cache with bcache

- Only new devices can be configured with bcache, the procedure wipes all on the device
  - create a partition for the file system on the HDD, and a partition for the bcache device on the SDD (fdisk etc.)
  - **make-bcache –C /dev/sdc1 –B /dev/sdb1**
    - Creates a /dev/bcache0 device
    - **mkfs.ext4 /dev/bcache0**
    - **mount /dev/bcache0 /mnt/data**
    - **bcache-status -a**

Enabling Writeback

- Multiple modes for caching writes are supported
  - **cat /sys/block/bcache0/bcache/cache_mode**
  - **[writetrough] writeback writearound none**
    - writethrough: write is done synchronously to cache and backing store
    - writeback: writes occur to cache first and write to disk is postponed
    - writearound: writes go directly to disk and bypass cache; cache is read exclusively
- **bcache-status** shows the current mode

Questions?

Thank you.

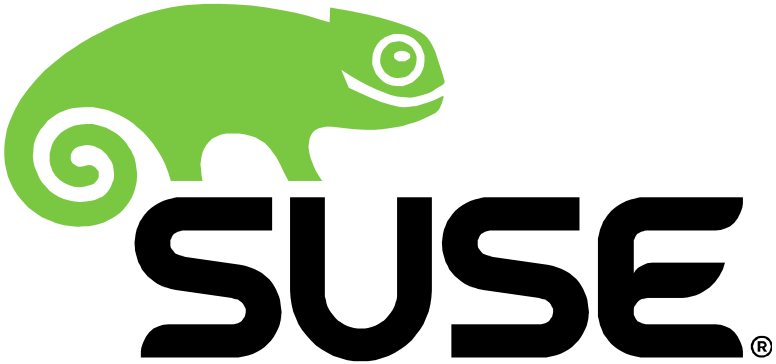
SUSE®
We adapt. You succeed.