

# Analyzing and Optimizing Linux Memory

Sander van Vugt

[mail@sandervanvugt.nl](mailto:mail@sandervanvugt.nl)

[www.sandervanvugt.com](http://www.sandervanvugt.com)

# Understanding Memory

- free -m?

```
[root@centos init]# free -m
```

	total	used	free	shared
buffers	cached			
Mem:	988	768	220	0
109	324			
-/+ buffers/cache:		333	654	
Swap:	1983	0	1983	

- Test it: `echo 3 > /proc/sys/vm/drop_caches`
- Why do you always keep some memory in buffers/cache?

# This talk

- For system administrators
- It picks out two common situations encountered on Linux memory
- It does NOT give a complete overview of everything (we would need a Linux Memory miniconf for the complete overview)

# About Linux Memory

- /proc/meminfo gives much details
  - Look at active/inactive memory
  - anon versus file

```
Active:                191028 kB
Inactive:              35944 kB
Active(anon):          169752 kB
Inactive(anon):        2704 kB
Active(file):          21276 kB
Inactive(file):        33240 kB
```

- Also consider virtual memory

```
VmallocTotal: 34359738367 kB
```

# What is virtual memory?

- It's an address space that programs can reserve memory in
- It's virtual, so it doesn't exist!
- It is NOT swap space
- Only at the moment that a process really needs the memory, it becomes resident memory

# Overcommit

- Because of virtual memory, Linux can overcommit a lot of memory
- overcommit is good, because no process uses all memory it claims anyway
- overcommit can bring lots of trouble!
- Look at the following slide

top - 14:23:46 up 2:28, 3 users, load average: 0.08, 0.17, 0.28  
Tasks: 164 total, 1 running, 163 sleeping, 0 stopped, 0 zombie  
Cpu(s): 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st  
Mem: 1012352k total, 550200k used, 462152k free, 17284k buffers  
Swap: 2031608k total, 0k used, 2031608k free, 99176k cached

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1922	root	20	0	234m	38m	10m	S	0.0	3.9	0:05.06	Xorg
4926	root	20	0	527m	38m	13m	S	0.0	3.9	0:01.16	system-config-f
4935	root	20	0	683m	37m	16m	S	0.0	3.8	0:00.81	python
4081	root	20	0	332m	27m	12m	S	0.0	2.8	0:00.95	system-config-d
4105	root	20	0	346m	27m	12m	S	0.0	2.8	0:01.63	system-config-s
4104	root	20	0	389m	27m	14m	S	0.0	2.7	0:00.59	python2
4930	root	20	0	357m	24m	12m	S	0.0	2.5	0:00.47	system-config-a
2431	root	20	0	48124	21m	480	S	0.0	2.2	0:00.15	restorecond
2270	root	20	0	273m	21m	4420	S	0.0	2.2	0:00.76	gnome-screensav
2211	root	20	0	533m	21m	13m	S	0.0	2.1	0:00.81	nautilus
2258	root	20	0	316m	18m	9276	S	0.0	1.9	0:00.25	python
2365	root	20	0	407m	17m	13m	S	0.0	1.8	0:00.15	gnote
4074	root	20	0	309m	15m	11m	S	0.0	1.6	0:00.43	gedit
2367	root	20	0	455m	13m	10m	S	0.0	1.4	0:00.33	clock-applet
2255	root	20	0	324m	13m	9876	S	0.0	1.3	0:00.27	nm-applet
2209	root	20	0	328m	13m	9768	S	0.0	1.3	0:00.57	gnome-panel
2366	root	20	0	390m	12m	8604	S	0.0	1.3	0:00.12	gdm-user-switch

# Handling overcommit

- Use `vm.overcommit_memory` to manage overcommit
- When limiting but not disallowing overcommit, overcommit ratio determines how much can be allocated.



# Let's talk about memory usage

```
MemTotal:      1012352 kB
MemFree:       460432 kB
Buffers:       17656 kB
Cached:        100884 kB
SwapCached:    0 kB
Active:        329748 kB
Inactive:      88312 kB
Active(anon):  299368 kB
Inactive(anon): 5736 kB
Active(file):  30380 kB
Inactive(file): 82576 kB
```

# anon and file memory

- anon memory roughly corresponds to memory that is used by programs
- file memory roughly corresponds to memory in use by cache and buffers
- inactive file memory is memory that is candidate to be dropped from cache
- inactive anon memory is candidate to be swapped out
- it makes no sense to drop active file memory as you need it
- it is very bad for performance to move active anon memory to swap

# How using swap can be good to optimize memory

- Imagine this memory usage (based on a customer case)
  - MemTotal: 8G
  - Swaptotal: 2G
  - MemFree: 12M
  - SwapFree: 0G
  - Active(anon):2G
  - Inactive(anon):6G
  - Active(file): 1900M
  - Inactive(file):90M
- What's happening here?

# Previous scenario analysis

- Swap is completely full
- There's not enough cache, because almost all file memory is active
- There's much Inactive(anon) but not enough swap to put it in
- The solution?
  - Increase swap to at least 8GB
  - (Temporarily) increase `vm.swappiness` to take of the memory stress