

---

Subject: Re: mem leak

Posted by [kapper](#) on Thu, 24 Apr 2008 07:32:17 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Dear Vaverin,

I would absolutely agree - but....

things on the hardware-node look bad after roughly 45 days uptime:

```
free -m
```

	total	used	free	shared	buffers	cached
Mem:	3946	3846	100	0	34	585
-/+ buffers/cache:		415	515			
Swap:	8191	8191	0			

which similar 2.6.9 kernel based host-nodes do not share as a common resource-problem.

this view doesn't change when killing all VEs - swap is still at 8 GB USED - which is for an idle host - definitely nothing normal.

the same box after the reboot yesterday and several hours uptime:

```
free -m
```

	total	used	free	shared	buffers	cached
Mem:	3946	3782	163	0	201	2430
-/+ buffers/cache:		1150	2796			
Swap:	8191	0	8191			

also three different hardware-nodes do experience the same problem running the EL5 kernel, EL4 kernel based boxes do not suffer this very problem.

finally there are some error-messages that might enlighten anyone on this problem:

WARNING: Kernel Errors Present

[<ffffff8005ee39>] error\_exit+0x0/0x84 ...: 14 Time(s)

and on the other box

WARNING: Kernel Errors Present

[<ffffff8005ee39>] error\_exit+0x0/0x84 ...: 30 Time(s)

and the serverlog shows the outofmemory-killing going on as we invoke yum to upgrade the kernel...

```
Apr 23 23:49:22 hostname kernel: yum invoked oom-killer: gfp_mask=0x201d2, order=0, oomkilladj=0
```

```
Apr 23 23:49:22 hostname kernel:
```

```

Apr 23 23:49:22 hostname kernel: Call Trace:
Apr 23 23:49:22 hostname kernel: [<ffffffff800c6672>] out_of_memory+0x9f/0x25e
Apr 23 23:49:22 hostname kernel: [<ffffffff8000eda3>] __alloc_pages+0x242/0x332
Apr 23 23:49:22 hostname kernel: [<ffffffff800124c4>]
__do_page_cache_readahead+0x95/0x1d9
Apr 23 23:49:22 hostname kernel: [<ffffffff8002999f>] sync_page+0x0/0x42
Apr 23 23:49:22 hostname kernel: [<ffffffff8005b8d4>] getnstimeofday+0x10/0x28
Apr 23 23:49:22 hostname kernel: [<ffffffff8009d3c7>] ktime_get_ts+0x1a/0x4d
Apr 23 23:49:22 hostname kernel: [<ffffffff800bf752>] delayacct_end+0x5d/0x86
Apr 23 23:49:22 hostname kernel: [<ffffffff80032e9c>]
blockable_page_cache_readahead+0x53/0xb2
Apr 23 23:49:22 hostname kernel: [<ffffffff8002fdef>] make_ahead_window+0x82/0x9e
Apr 23 23:49:22 hostname kernel: [<ffffffff80013763>] page_cache_readahead+0x17f/0x1af
Apr 23 23:49:22 hostname kernel: [<ffffffff8000b992>] do_generic_mapping_read+0x126/0x3f8
Apr 23 23:49:22 hostname kernel: [<ffffffff8000c917>] file_read_actor+0x0/0x13c
Apr 23 23:49:22 hostname kernel: [<ffffffff8000bdb0>] __generic_file_aio_read+0x14c/0x190
Apr 23 23:49:22 hostname kernel: [<ffffffff80016800>] generic_file_aio_read+0x34/0x39
Apr 23 23:49:22 hostname kernel: [<ffffffff8000c633>] do_sync_read+0xc7/0x104
Apr 23 23:49:22 hostname kernel: [<ffffffff80066051>] do_page_fault+0x4e9/0x7ef
Apr 23 23:49:22 hostname kernel: [<ffffffff8009b432>] autoremove_wake_function+0x0/0x2e
Apr 23 23:49:22 hostname kernel: [<ffffffff80061597>] __sched_text_start+0x177/0xee2
Apr 23 23:49:22 hostname kernel: [<ffffffff8000af5c>] vfs_read+0xaa/0x150
Apr 23 23:49:22 hostname kernel: [<ffffffff8000b125>] fget_light+0x18/0x7c
Apr 23 23:49:22 hostname kernel: [<ffffffff80012c4a>] sys_pread64+0x54/0xc4
Apr 23 23:49:22 hostname kernel: [<ffffffff8005e166>] system_call+0x7e/0x83
Apr 23 23:49:22 hostname kernel:
Apr 23 23:49:22 hostname kernel: Mem-info:
Apr 23 23:49:22 hostname kernel: Node 0 DMA per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: Node 0 DMA32 per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 186, batch 31 used:125
Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 62, batch 15 used:50
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 186, batch 31 used:15
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 62, batch 15 used:14
Apr 23 23:49:22 hostname kernel: Node 0 Normal per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 186, batch 31 used:57
Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 62, batch 15 used:57
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 186, batch 31 used:36
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 62, batch 15 used:46
Apr 23 23:49:22 hostname kernel: Node 0 HighMem per-cpu: empty
Apr 23 23:49:22 hostname kernel: Free pages: 21556kB (0kB HighMem)
Apr 23 23:49:22 hostname kernel: Active:482971 inactive:471530 dirty:500 writeback:0 unstable:0
free:5389 slab:26825 mapped-file:9603 mapped-anon:844193 pagetables:10871
Apr 23 23:49:22 hostname kernel: Node 0 DMA free:10980kB min:20kB low:24kB high:28kB
active:0kB inactive:0kB present:10532kB pages_scanned:0 all_unreclaimable? yes

```

```

Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 3502 4006 4006
Apr 23 23:49:22 hostname kernel: Node 0 DMA32 free:9196kB min:7072kB low:8840kB
high:10608kB active:1727524kB inactive:1716528kB present:3586732kB pages_scanned:3724
all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 504 504
Apr 23 23:49:22 hostname kernel: Node 0 Normal free:1380kB min:1016kB low:1268kB
high:1524kB active:204360kB inactive:169592kB present:516096kB pages_scanned:4
all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 0 0
Apr 23 23:49:22 hostname kernel: Node 0 HighMem free:0kB min:128kB low:128kB high:128kB
active:0kB inactive:0kB present:0kB pages_scanned:0 all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 0 0
Apr 23 23:49:22 hostname kernel: Node 0 DMA: 3*4kB 5*8kB 3*16kB 4*32kB 4*64kB 2*128kB
2*256kB 1*512kB 1*1024kB 0*2048kB 2*4096kB = 10980kB
Apr 23 23:49:22 hostname kernel: Node 0 DMA32: 31*4kB 2*8kB 18*16kB 32*32kB 7*64kB
1*128kB 0*256kB 0*512kB 1*1024kB 1*2048kB 1*4096kB = 9196kB
Apr 23 23:49:22 hostname kernel: Node 0 Normal: 105*4kB 8*8kB 10*16kB 1*32kB 3*64kB
0*128kB 0*256kB 1*512kB 0*1024kB 0*2048kB 0*4096kB = 1380kB
Apr 23 23:49:22 hostname kernel: Node 0 HighMem: empty
Apr 23 23:49:22 hostname kernel: Swap cache: add 2338321, delete 2336791, find
32746461/32778627, race 0+92+75
Apr 23 23:49:22 hostname kernel: Free swap = 0kB
Apr 23 23:49:22 hostname kernel: Total swap = 8388600kB
Apr 23 23:49:22 hostname kernel: Free swap: 0kB
Apr 23 23:49:22 hostname kernel: 1179648 pages of RAM
Apr 23 23:49:22 hostname kernel: 169401 reserved pages
Apr 23 23:49:22 hostname kernel: 73976 pages shared
Apr 23 23:49:22 hostname kernel: 1531 pages swap cached
Apr 23 23:49:22 hostname kernel: Top 10 caches:
Apr 23 23:49:22 hostname kernel: filp : size 1617920 objsize 256
Apr 23 23:49:22 hostname kernel: page_beancounter : size 61296640 objsize 64
Apr 23 23:49:22 hostname kernel: buffer_head : size 1445888 objsize 96
Apr 23 23:49:22 hostname kernel: radix_tree_node : size 3624960 objsize 536
Apr 23 23:49:22 hostname kernel: size-2048 : size 1688960 objsize 2048
Apr 23 23:49:22 hostname kernel: size-128 : size 1941504 objsize 128
Apr 23 23:49:22 hostname kernel: vm_area_struct : size 5017600 objsize 176
Apr 23 23:49:22 hostname kernel: ext3_inode_cache : size 8697216 objsize 808
Apr 23 23:49:22 hostname kernel: task_struct : size 1425408 objsize 2192
Apr 23 23:49:22 hostname kernel: dentry_cache : size 5943296 objsize 248
Apr 23 23:49:22 hostname kernel: Out of memory: Killed process 12866 (apache2).
Apr 23 23:49:22 hostname kernel: Mem-info:
Apr 23 23:49:22 hostname kernel: Node 0 DMA per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 0, batch 1 used:0
Apr 23 23:49:22 hostname kernel: Node 0 DMA32 per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 186, batch 31 used:125

```

```

Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 62, batch 15 used:49
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 186, batch 31 used:15
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 62, batch 15 used:14
Apr 23 23:49:22 hostname kernel: Node 0 Normal per-cpu:
Apr 23 23:49:22 hostname kernel: cpu 0 hot: high 186, batch 31 used:22
Apr 23 23:49:22 hostname kernel: cpu 0 cold: high 62, batch 15 used:61
Apr 23 23:49:22 hostname kernel: cpu 1 hot: high 186, batch 31 used:36
Apr 23 23:49:22 hostname kernel: cpu 1 cold: high 62, batch 15 used:46
Apr 23 23:49:22 hostname kernel: Node 0 HighMem per-cpu: empty
Apr 23 23:49:22 hostname kernel: Free pages:      24136kB (0kB HighMem)
Apr 23 23:49:22 hostname kernel: Active:482668 inactive:471214 dirty:128 writeback:0 unstable:0
free:6034 slab:26851 mapped-file:9603 mapped-anon:844193 pagetables:10871
Apr 23 23:49:22 hostname kernel: Node 0 DMA free:10980kB min:20kB low:24kB high:28kB
active:0kB inactive:0kB present:10532kB pages_scanned:0 all_unreclaimable? yes
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 3502 4006 4006
Apr 23 23:49:22 hostname kernel: Node 0 DMA32 free:11536kB min:7072kB low:8840kB
high:10608kB active:1726548kB inactive:1715316kB present:3586732kB pages_scanned:364
all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 504 504
Apr 23 23:49:22 hostname kernel: Node 0 Normal free:1620kB min:1016kB low:1268kB
high:1524kB active:204124kB inactive:169540kB present:516096kB pages_scanned:6
all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 0 0
Apr 23 23:49:22 hostname kernel: Node 0 HighMem free:0kB min:128kB low:128kB high:128kB
active:0kB inactive:0kB present:0kB pages_scanned:0 all_unreclaimable? no
Apr 23 23:49:22 hostname kernel: lowmem_reserve[]: 0 0 0 0
Apr 23 23:49:22 hostname kernel: Node 0 DMA: 3*4kB 5*8kB 3*16kB 4*32kB 4*64kB 2*128kB
2*256kB 1*512kB 1*1024kB 0*2048kB 2*4096kB = 10980kB
Apr 23 23:49:22 hostname kernel: Node 0 DMA32: 584*4kB 18*8kB 18*16kB 32*32kB 7*64kB
1*128kB 0*256kB 0*512kB 1*1024kB 1*2048kB 1*4096kB = 11536kB
Apr 23 23:49:22 hostname kernel: Node 0 Normal: 165*4kB 8*8kB 10*16kB 1*32kB 3*64kB
0*128kB 0*256kB 1*512kB 0*1024kB 0*2048kB 0*4096kB = 1620kB
Apr 23 23:49:22 hostname kernel: Node 0 HighMem: empty
Apr 23 23:49:22 hostname kernel: Swap cache: add 2338321, delete 2336791, find
32746461/32778627, race 0+92+75
Apr 23 23:49:22 hostname kernel: Free swap = 0kB
Apr 23 23:49:22 hostname kernel: Total swap = 8388600kB
Apr 23 23:49:22 hostname kernel: Free swap:      0kB
Apr 23 23:49:22 hostname kernel: 1179648 pages of RAM
Apr 23 23:49:22 hostname kernel: 169401 reserved pages
Apr 23 23:49:22 hostname kernel: 74036 pages shared
Apr 23 23:49:22 hostname kernel: 1531 pages swap cached
Apr 23 23:49:22 hostname kernel: OOM killed process apache2 (pid=12893, ve=419) exited,
free=6096 gen=147.
Apr 23 23:54:21 hostname shutdown[14605]: shutting down for system reboot
Apr 23 23:54:21 hostname init: Switching to runlevel: 6
Apr 23 23:54:23 hostname smartd[22658]: smartd received signal 15: Terminated
Apr 23 23:54:23 hostname smartd[22658]: smartd is exiting (exit status 0)

```

Apr 23 23:54:33 hostname ntpd[6324]: ntpd exiting on signal 15  
Apr 23 23:54:35 hostname kernel: Kernel logging (proc) stopped.  
Apr 23 23:54:35 hostname kernel: Kernel log daemon terminating.  
Apr 23 23:54:36 hostname exiting on signal 15  
Apr 24 00:43:11 hostname syslogd 1.4.1: restart.

any advise is greatly appreciated  
hk

---