
Subject: URGENT help need! invoked OOM-KILLER with no clear reason

Posted by [prosto](#) on Fri, 24 Aug 2012 09:47:32 GMT

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

Hi list,

We used openvz in production servers from 1-2 years ago, but now really need your suggestions and help. So this is the situation.

We have setup with two identical openvz servers, and they was working perfectly fine until a the problem rise, so this is the problem:

In a random time interval (1-2-3 hours), the system (on the HN) lunch OOM-KILLER and start killing processes in different containers (so far not so unusual) , but the real problem is that 1. on the HN there is a lot of free memory resources and 2. in the VPS containers there is also free memory resources.

From the log analyze it's clear that random processes do PF , which invoke the oom-killer , but it's really must NOT to do it.

So, there are the /var/log/messages messages:

```
Aug 23 21:40:09 vz2 kernel: [13444137.871993] postgres invoked oom-killer: gfp_mask=0x0, order=0, oom_adj=0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.871995] Pid: 653, comm: postgres Not tainted 2.6.32-5-openvz-amd64 #1
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.871997] Call Trace:
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872002] [<ffffffff810b8eec>] ? oom_kill_process+0x75/0x22c
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872004] [<ffffffff810b940e>] ? __out_of_memory+0x93/0xad
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872006] [<ffffffff810b969a>] ? pagefault_out_of_memory+0xba/0xe8
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872009] [<ffffffff81032612>] ? mm_fault_error+0x5c/0x142
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872012] [<ffffffff812ed805>] ? do_page_fault+0x2e0/0x2fc
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872016] [<ffffffff812eb695>] ? page_fault+0x25/0x30
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872017] Mem-Info:
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872018] Node 0 DMA per-cpu:
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872020] CPU 0: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872021] CPU 1: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872022] CPU 2: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872024] CPU 3: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872025] CPU 4: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872026] CPU 5: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872028] CPU 6: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872029] CPU 7: hi: 0, btch: 1 usd: 0
```

```
Aug 23 21:40:09 vz2 kernel: [13444137.872030] Node 0 DMA32 per-cpu:
```

Aug 23 21:40:09 vz2 kernel: [13444137.872031] CPU 0: hi: 186, btch: 31 usd: 165
Aug 23 21:40:09 vz2 kernel: [13444137.872033] CPU 1: hi: 186, btch: 31 usd: 12
Aug 23 21:40:09 vz2 kernel: [13444137.872034] CPU 2: hi: 186, btch: 31 usd: 170
Aug 23 21:40:09 vz2 kernel: [13444137.872035] CPU 3: hi: 186, btch: 31 usd: 156
Aug 23 21:40:09 vz2 kernel: [13444137.872037] CPU 4: hi: 186, btch: 31 usd: 157
Aug 23 21:40:09 vz2 kernel: [13444137.872038] CPU 5: hi: 186, btch: 31 usd: 155
Aug 23 21:40:09 vz2 kernel: [13444137.872039] CPU 6: hi: 186, btch: 31 usd: 157
Aug 23 21:40:09 vz2 kernel: [13444137.872040] CPU 7: hi: 186, btch: 31 usd: 174
Aug 23 21:40:09 vz2 kernel: [13444137.872041] Node 0 Normal per-cpu:
Aug 23 21:40:09 vz2 kernel: [13444137.872043] CPU 0: hi: 186, btch: 31 usd: 108
Aug 23 21:40:09 vz2 kernel: [13444137.872044] CPU 1: hi: 186, btch: 31 usd: 161
Aug 23 21:40:09 vz2 kernel: [13444137.872045] CPU 2: hi: 186, btch: 31 usd: 163
Aug 23 21:40:09 vz2 kernel: [13444137.872046] CPU 3: hi: 186, btch: 31 usd: 127
Aug 23 21:40:09 vz2 kernel: [13444137.872048] CPU 4: hi: 186, btch: 31 usd: 139
Aug 23 21:40:09 vz2 kernel: [13444137.872049] CPU 5: hi: 186, btch: 31 usd: 171
Aug 23 21:40:09 vz2 kernel: [13444137.872050] CPU 6: hi: 186, btch: 31 usd: 133
Aug 23 21:40:09 vz2 kernel: [13444137.872052] CPU 7: hi: 186, btch: 31 usd: 172
Aug 23 21:40:09 vz2 kernel: [13444137.872055] active_anon:127099 inactive_anon:16297
isolated_anon:0
Aug 23 21:40:09 vz2 kernel: [13444137.872056] active_file:668874 inactive_file:2839950
isolated_file:0
Aug 23 21:40:09 vz2 kernel: [13444137.872056] unevictable:0 dirty:3089 writeback:0 unstable:0
Aug 23 21:40:09 vz2 kernel: [13444137.872057] free:225934 slab_reclaimable:166494
slab_unreclaimable:13922
Aug 23 21:40:09 vz2 kernel: [13444137.872058] mapped:13973 shmem:1139 pagetables:8974
bounce:0
Aug 23 21:40:09 vz2 kernel: [13444137.872059] Node 0 DMA free:15800kB min:12kB low:12kB
high:16kB active_anon:0kB inactive_anon:0kB active_file:0kB inactive_file:0kB unevictable:0kB
isolated(anon):0kB isolated(file):0kB present:15188kB mlocked:0kB dirty:0kB writeback:0kB
mapped:0kB shmem:0kB slab_reclaimable:0kB slab_unreclaimable:0kB kernel_stack:0kB
pagetables:0kB unstable:0kB bounce:0kB writeback_tmp:0kB pages_scanned:0
all_unreclaimable? yes
Aug 23 21:40:09 vz2 kernel: [13444137.872067] lowmem_reserve[]: 0 2982 16086 16086
Aug 23 21:40:09 vz2 kernel: [13444137.872069] Node 0 DMA32 free:313340kB min:3008kB
low:3760kB high:4512kB active_anon:844kB inactive_anon:3524kB active_file:259008kB
inactive_file:2010296kB unevictable:0kB isolated(anon):0kB isolated(file):0kB present:3054420kB
mlocked:0kB dirty:32kB writeback:0kB mapped:4344kB shmem:0kB slab_reclaimable:181804kB
slab_unreclaimable:9844kB kernel_stack:408kB pagetables:156kB unstable:0kB bounce:0kB
writeback_tmp:0kB pages_scanned:0 all_unreclaimable? no
Aug 23 21:40:09 vz2 kernel: [13444137.872077] lowmem_reserve[]: 0 0 13104 13104
Aug 23 21:40:09 vz2 kernel: [13444137.872078] Node 0 Normal free:574596kB min:13216kB
low:16520kB high:19824kB active_anon:507552kB inactive_anon:61664kB active_file:2416488kB
inactive_file:9349504kB unevictable:0kB isolated(anon):0kB isolated(file):0kB
present:13418496kB mlocked:0kB dirty:12324kB writeback:0kB mapped:51548kB
shmem:4556kB slab_reclaimable:484172kB slab_unreclaimable:45844kB kernel_stack:4464kB
pagetables:35740kB unstable:0kB bounce:0kB writeback_tmp:0kB pages_scanned:0
all_unreclaimable? no
Aug 23 21:40:09 vz2 kernel: [13444137.872086] lowmem_reserve[]: 0 0 0 0

```

Aug 23 21:40:09 vz2 kernel: [13444137.872088] Node 0 DMA: 2*4kB 2*8kB 0*16kB 1*32kB
2*64kB 0*128kB 1*256kB 0*512kB 1*1024kB 1*2048kB 3*4096kB = 15800kB
Aug 23 21:40:09 vz2 kernel: [13444137.872093] Node 0 DMA32: 20691*4kB 8508*8kB
4799*16kB 1613*32kB 405*64kB 67*128kB 0*256kB 0*512kB 0*1024kB 0*2048kB 0*4096kB =
313724kB
Aug 23 21:40:09 vz2 kernel: [13444137.872099] Node 0 Normal: 71219*4kB 27806*8kB
1748*16kB 1207*32kB 0*64kB 0*128kB 0*256kB 0*512kB 1*1024kB 0*2048kB 0*4096kB =
574940kB
Aug 23 21:40:09 vz2 kernel: [13444137.872104] 3512939 total pagecache pages
Aug 23 21:40:09 vz2 kernel: [13444137.872105] 2963 pages in swap cache
Aug 23 21:40:09 vz2 kernel: [13444137.872106] Swap cache stats: add 58675, delete 55712, find
294564024/294568150
Aug 23 21:40:09 vz2 kernel: [13444137.872108] Free swap = 31207600kB
Aug 23 21:40:09 vz2 kernel: [13444137.872109] Total swap = 31247348kB
Aug 23 21:40:09 vz2 kernel: [13444137.909766] 4194288 pages RAM
Aug 23 21:40:09 vz2 kernel: [13444137.909768] 86418 pages reserved
Aug 23 21:40:09 vz2 kernel: [13444137.909769] 1510860 pages shared
Aug 23 21:40:09 vz2 kernel: [13444137.909770] 2659168 pages non-shared
Aug 23 21:40:09 vz2 kernel: [13444137.909845] Mem-Info:
Aug 23 21:40:09 vz2 kernel: [13444137.909847] Node 0 DMA per-cpu:
Aug 23 21:40:09 vz2 kernel: [13444137.909850] CPU 0: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909852] CPU 1: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909855] CPU 2: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909857] CPU 3: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909859] CPU 4: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909862] CPU 5: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909864] CPU 6: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909866] CPU 7: hi: 0, btch: 1 usd: 0
Aug 23 21:40:09 vz2 kernel: [13444137.909868] Node 0 DMA32 per-cpu:
Aug 23 21:40:09 vz2 kernel: [13444137.909871] CPU 0: hi: 186, btch: 31 usd: 165
Aug 23 21:40:09 vz2 kernel: [13444137.909873] CPU 1: hi: 186, btch: 31 usd: 12
Aug 23 21:40:09 vz2 kernel: [13444137.909876] CPU 2: hi: 186, btch: 31 usd: 170
Aug 23 21:40:09 vz2 kernel: [13444137.909878] CPU 3: hi: 186, btch: 31 usd: 156
Aug 23 21:40:09 vz2 kernel: [13444137.909880] CPU 4: hi: 186, btch: 31 usd: 157
Aug 23 21:40:09 vz2 kernel: [13444137.909883] CPU 5: hi: 186, btch: 31 usd: 155
Aug 23 21:40:09 vz2 kernel: [13444137.909885] CPU 6: hi: 186, btch: 31 usd: 157
Aug 23 21:40:09 vz2 kernel: [13444137.909887] CPU 7: hi: 186, btch: 31 usd: 174
Aug 23 21:40:09 vz2 kernel: [13444137.909889] Node 0 Normal per-cpu:
Aug 23 21:40:09 vz2 kernel: [13444137.909892] CPU 0: hi: 186, btch: 31 usd: 108
Aug 23 21:40:09 vz2 kernel: [13444137.909894] CPU 1: hi: 186, btch: 31 usd: 161
Aug 23 21:40:09 vz2 kernel: [13444137.909897] CPU 2: hi: 186, btch: 31 usd: 163
Aug 23 21:40:09 vz2 kernel: [13444137.909899] CPU 3: hi: 186, btch: 31 usd: 127
Aug 23 21:40:09 vz2 kernel: [13444137.909901] CPU 4: hi: 186, btch: 31 usd: 139
Aug 23 21:40:09 vz2 kernel: [13444137.909903] CPU 5: hi: 186, btch: 31 usd: 171
Aug 23 21:40:09 vz2 kernel: [13444137.909906] CPU 6: hi: 186, btch: 31 usd: 133
Aug 23 21:40:09 vz2 kernel: [13444137.909908] CPU 7: hi: 186, btch: 31 usd: 173
Aug 23 21:40:09 vz2 kernel: [13444137.909912] active_anon:127099 inactive_anon:16297
isolated_anon:0

```

Aug 23 21:40:09 vz2 kernel: [13444137.909914] active_file:668874 inactive_file:2839950
isolated_file:0
Aug 23 21:40:09 vz2 kernel: [13444137.909915] unevictable:0 dirty:3089 writeback:0 unstable:0
Aug 23 21:40:09 vz2 kernel: [13444137.909916] free:225934 slab_reclaimable:166494
slab_unreclaimable:13922
Aug 23 21:40:09 vz2 kernel: [13444137.909917] mapped:13973 shmem:1139 pagetables:8974
bounce:0
Aug 23 21:40:09 vz2 kernel: [13444137.909920] Node 0 DMA free:15800kB min:12kB low:12kB
high:16kB active_anon:0kB inactive_anon:0kB active_file:0kB inactive_file:0kB unevictable:0kB
isolated(anon):0kB isolated(file):0kB present:15188kB mlocked:0kB dirty:0kB writeback:0kB
mapped:0kB shmem:0kB slab_reclaimable:0kB slab_unreclaimable:0kB kernel_stack:0kB
pagetables:0kB unstable:0kB bounce:0kB writeback_tmp:0kB pages_scanned:0
all_unreclaimable? yes
Aug 23 21:40:09 vz2 kernel: [13444137.909931] lowmem_reserve[]: 0 2982 16086 16086
Aug 23 21:40:09 vz2 kernel: [13444137.909936] Node 0 DMA32 free:313340kB min:3008kB
low:3760kB high:4512kB active_anon:844kB inactive_anon:3524kB active_file:259008kB
inactive_file:2010296kB unevictable:0kB isolated(anon):0kB isolated(file):0kB present:3054420kB
mlocked:0kB dirty:32kB writeback:0kB mapped:4344kB shmem:0kB slab_reclaimable:181804kB
slab_unreclaimable:9844kB kernel_stack:408kB pagetables:156kB unstable:0kB bounce:0kB
writeback_tmp:0kB pages_scanned:0 all_unreclaimable? no
Aug 23 21:40:09 vz2 kernel: [13444137.909948] lowmem_reserve[]: 0 0 13104 13104
Aug 23 21:40:09 vz2 kernel: [13444137.909953] Node 0 Normal free:574596kB min:13216kB
low:16520kB high:19824kB active_anon:507552kB inactive_anon:61664kB active_file:2416488kB
inactive_file:9349504kB unevictable:0kB isolated(anon):0kB isolated(file):0kB
present:13418496kB mlocked:0kB dirty:12324kB writeback:0kB mapped:51548kB
shmem:4556kB slab_reclaimable:484172kB slab_unreclaimable:45844kB kernel_stack:4464kB
pagetables:35740kB unstable:0kB bounce:0kB writeback_tmp:0kB pages_scanned:0
all_unreclaimable? no
Aug 23 21:40:09 vz2 kernel: [13444137.909965] lowmem_reserve[]: 0 0 0 0
Aug 23 21:40:09 vz2 kernel: [13444137.909967] Node 0 DMA: 2*4kB 2*8kB 0*16kB 1*32kB
2*64kB 0*128kB 1*256kB 0*512kB 1*1024kB 1*2048kB 3*4096kB = 15800kB
Aug 23 21:40:09 vz2 kernel: [13444137.909975] Node 0 DMA32: 20691*4kB 8508*8kB
4799*16kB 1613*32kB 405*64kB 67*128kB 0*256kB 0*512kB 0*1024kB 0*2048kB 0*4096kB =
313724kB
Aug 23 21:40:09 vz2 kernel: [13444137.909983] Node 0 Normal: 71219*4kB 27806*8kB
1749*16kB 1207*32kB 0*64kB 0*128kB 0*256kB 0*512kB 1*1024kB 0*2048kB 0*4096kB =
574956kB
Aug 23 21:40:09 vz2 kernel: [13444137.909991] 3512939 total pagecache pages
Aug 23 21:40:09 vz2 kernel: [13444137.909993] 2963 pages in swap cache
Aug 23 21:40:09 vz2 kernel: [13444137.909995] Swap cache stats: add 58675, delete 55712, find
294564024/294568150
Aug 23 21:40:09 vz2 kernel: [13444137.909997] Free swap = 31207600kB
Aug 23 21:40:09 vz2 kernel: [13444137.909998] Total swap = 31247348kB
Aug 23 21:40:10 vz2 kernel: [13444137.959017] 4194288 pages RAM
Aug 23 21:40:10 vz2 kernel: [13444137.959019] 86418 pages reserved
Aug 23 21:40:10 vz2 kernel: [13444137.959020] 1510878 pages shared
Aug 23 21:40:10 vz2 kernel: [13444137.959021] 2659168 pages non-shared
Aug 23 21:40:10 vz2 kernel: [13444137.964021] OOM killed process apache2 (pid=13259,

ve=202) exited, free=230863 gen=22391.
Aug 23 21:40:10 vz2 kernel: [13444137.972302] OOM killed process apache2 (pid=13239, ve=202) exited, free=235866 gen=22392.
Aug 23 21:40:10 vz2 kernel: [13444137.976072] OOM killed process apache2 (pid=27304, ve=201) exited, free=237075 gen=22393.
Aug 23 21:40:10 vz2 kernel: [13444137.983820] OOM killed process apache2 (pid=13247, ve=202) exited, free=242072 gen=22394.
Aug 23 21:40:10 vz2 kernel: [13444137.987171] OOM killed process udevd (pid=29123, ve=211) exited, free=242167 gen=22395.
Aug 23 21:40:10 vz2 kernel: [13444137.991245] OOM killed process rsyslogd (pid=29416, ve=211) exited, free=242167 gen=22396.
Aug 23 21:40:10 vz2 kernel: [13444137.995244] OOM killed process sshd (pid=29429, ve=211) exited, free=242332 gen=22397.
Aug 23 21:40:10 vz2 kernel: [13444137.999140] OOM killed process xinetd (pid=29436, ve=211) exited, free=242408 gen=22398.
Aug 23 21:40:10 vz2 kernel: [13444138.003141] OOM killed process saslauthd (pid=29443, ve=211) exited, free=242408 gen=22399.
Aug 23 21:40:10 vz2 kernel: [13444138.007302] OOM killed process sendmail (pid=29459, ve=211) exited, free=242856 gen=22400.
Aug 23 21:40:10 vz2 kernel: [13444138.011214] OOM killed process sendmail (pid=29466, ve=211) exited, free=243228 gen=22401.
Aug 23 21:40:10 vz2 kernel: [13444138.015316] OOM killed process httpd (pid=29474, ve=211) exited, free=243414 gen=22402.
Aug 23 21:40:10 vz2 kernel: [13444138.023087] OOM killed process crond (pid=29482, ve=211) exited, free=243600 gen=22403.
Aug 23 21:40:10 vz2 kernel: [13444138.027811] OOM killed process radiusd (pid=3940, ve=211) exited, free=244904 gen=22404.
Aug 23 21:40:10 vz2 kernel: [13444138.039797] OOM killed process apache2 (pid=13234, ve=202) exited, free=249923 gen=22405.
Aug 23 21:40:10 vz2 kernel: [13444138.043952] OOM killed process apache2 (pid=29056, ve=201) exited, free=251318 gen=22406.
Aug 23 21:40:10 vz2 kernel: [13444138.051801] OOM killed process apache2 (pid=13271, ve=202) exited, free=256367 gen=22407.
Aug 23 21:40:10 vz2 kernel: [13444138.059855] OOM killed process apache2 (pid=27702, ve=202) exited, free=263823 gen=22408.
Aug 23 21:40:10 vz2 kernel: [13444138.067754] OOM killed process apache2 (pid=27701, ve=202) exited, free=271078 gen=22409.
Aug 23 21:40:10 vz2 kernel: [13444138.075214] OOM killed process apache2 (pid=32672, ve=202) exited, free=276172 gen=22410.
Aug 23 21:40:10 vz2 kernel: [13444138.083472] OOM killed process apache2 (pid=16198, ve=202) exited, free=283240 gen=22411.
Aug 23 21:40:10 vz2 kernel: [13444138.090803] OOM killed process apache2 (pid=19018, ve=202) exited, free=288169 gen=22412.
Aug 23 21:40:10 vz2 kernel: [13444138.099506] OOM killed process apache2 (pid=6824, ve=202) exited, free=293052 gen=22413.
Aug 23 21:40:10 vz2 kernel: [13444138.107363] OOM killed process apache2 (pid=21691, ve=202) exited, free=297671 gen=22414.
Aug 23 21:40:10 vz2 kernel: [13444138.114606] OOM killed process apache2 (pid=26866,

```
ve=202) exited, free=302015 gen=22415.
Aug 23 21:40:10 vz2 kernel: [13444138.122742] OOM killed process apache2 (pid=29596,
ve=202) exited, free=306201 gen=22416.
Aug 23 21:40:10 vz2 kernel: [13444138.129498] OOM killed process apache2 (pid=32155,
ve=202) exited, free=306736 gen=22417.
Aug 23 21:40:10 vz2 kernel: [13444138.133860] OOM killed process mysqld (pid=18656, ve=202)
exited, free=314273 gen=22418.
Aug 23 21:40:10 vz2 kernel: [13444138.135274] OOM killed process httpd (pid=29476, ve=211)
exited, free=314981 gen=22419.
Aug 23 21:40:10 vz2 kernel: [13444138.140485] OOM killed process postgres (pid=653, ve=206)
exited, free=316294 gen=22420.
```

as you see there in that particular situation the postgres process (now is postgres, but it's random process I mean, not everytime this process invokes OOM) invokes the killer, do call trace, and then with no meaning the HN do OOM killing on a lot of processes in different containers. (because the kernel thinks that this process should be killed - ok) , BUT in the time frames of oom killing on the system there is a lot of free memory resources (no swap used, a lot of cached ram , etc.). I am attaching a munin memory graph , so you can see what i'm talking about:

There is the HN OS information:

```
root@vz2:~# uname -a
Linux vz2 2.6.32-5-openvz-amd64 #1 SMP Mon Oct 3 05:12:50 UTC 2011 x86_64 GNU/Linux
root@vz2:~# cat /proc/version
Linux version 2.6.32-5-openvz-amd64 (Debian 2.6.32-38) (ben@decadent.org.uk) (gcc version
4.3.5 (Debian 4.3.5-4) ) #1 SMP Mon Oct 3 05:12:50 UTC 2011
```

Please reply if you need more information , or anything I can provide.
The problem is urgent, because the kernel kill some of the critical services for us.

OK, of course there is NO failcnt in the beancounters , there is no overcommit or something like this on vzmemcheck and other vz tools.

And Not least, we have two absolute identical servers (same hardware, same install, same kernels, just different VEs in that), but the problem rise only on one of them.

File Attachments

1) [Picture 20.png](#), downloaded 329 times
