

---

Subject: /usr/sbin/httpd: page allocation failure?

Posted by [HHawk](#) on Wed, 23 Feb 2022 08:15:08 GMT

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

---

Hi all,

I noticed, on almost all nodes with OpenVZ 7 with kernel: Linux servername  
3.10.0-1160.42.2.vz7.184.10 #1 SMP Fri Dec 31 04:05:16 MSK 2021 x86\_64 x86\_64 x86\_64  
GNU/Linux

The following (kind) of error: /usr/sbin/httpd: page allocation failure: order:5, mode:0x2040d0

```
[ +0.000170] /usr/sbin/httpd: page allocation failure: order:5, mode:0x2040d0
[ +0.000003] CPU: 2 PID: 63709 Comm: /usr/sbin/httpd ve: 10744 Kdump: loaded Not tainted
3.10.0-1160.42.2.vz7.184.10 #1 184.10
[ +0.000002] Hardware name: Dell Inc. PowerEdge R300/0TY179, BIOS 1.5.2 11/02/2010
[ +0.000002] Call Trace:
[ +0.000004] [<ffffffff831c04ce>] dump_stack+0x19/0x1b
[ +0.000002] [<ffffffff82bdd970>] warn_alloc_failed+0x110/0x170
[ +0.000003] [<ffffffff831bb8b8>] __alloc_pages_slowpath+0x684/0x799
[ +0.000004] [<ffffffff82be23af>] __alloc_pages_nodemask+0x6af/0x6e0
[ +0.000003] [<ffffffff831bc8ba>] kmalloc_large_node+0x5f/0x80
[ +0.000003] [<ffffffff82c4a59a>] __kmalloc_node_track_caller+0x26a/0x2c0
[ +0.000003] [<ffffffff82c8a76f>] ? __pollwait+0x7f/0xf0
[ +0.000005] [<ffffffff83079a9d>] ? __alloc_skb+0x8d/0x2d0
[ +0.000003] [<ffffffff83078a34>] __kmalloc_reserve.isra.32+0x44/0xa0
[ +0.000002] [<ffffffff83079a6d>] ? __alloc_skb+0x5d/0x2d0
[ +0.000003] [<ffffffff83079a9d>] __alloc_skb+0x8d/0x2d0
[ +0.000003] [<ffffffff830edaf2>] sk_stream_alloc_skb+0x52/0x1b0
[ +0.000002] [<ffffffff830ee9d8>] tcp_sendmsg+0x788/0xc60
[ +0.000003] [<ffffffff82ae64d8>] ? __enqueue_entity+0x78/0x80
[ +0.000003] [<ffffffff8311acc9>] inet_sendmsg+0x69/0xb0
[ +0.000003] [<ffffffff8306e50d>] sock_aio_write+0x15d/0x180
[ +0.000003] [<ffffffff82c73397>] do_sync_readv_writev+0x77/0xd0
[ +0.000003] [<ffffffff82c74f9e>] do_readv_writev+0xce/0x260
[ +0.000003] [<ffffffff82c751c5>] vfs_writev+0x35/0x60
[ +0.000002] [<ffffffff82c7537f>] Sys_writev+0x7f/0x110
[ +0.000003] [<ffffffff831d3f95>] ? system_call_after_swapgs+0xa2/0x13a
[ +0.000003] [<ffffffff831d4052>] system_call_fastpath+0x25/0x2a
[ +0.000003] [<ffffffff831d3f95>] ? system_call_after_swapgs+0xa2/0x13a
[ +0.000001] Mem-Info:
[ +0.000004] active_anon:579695 inactive_anon:93072 isolated_anon:0
active_file:537858 inactive_file:312461 isolated_file:0
unevictable:0 dirty:1031 writeback:4 wbtmpt:0 unstable:0
slab_reclaimable:372480 slab_unreclaimable:16200
mapped:94651 shmem:56310 pagetables:17328 bounce:0
free:39530 free_pcp:4 free_cma:0
[ +0.000007] Node 0 DMA free:15904kB min:128kB low:160kB high:192kB active_anon:0kB
inactive_anon:0kB active_file:0kB inactive_file:0kB unevictable:0kB isolate
```

```

                d(anon):0kB isolated(file):0kB present:15996kB managed:15904kB
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 free_pcp:0kB local_pcp:0kB free_cma:0kB
writeback_tmp:0kB pages_scanned:0 all_unreclaimable? yes
[ +0.000006] lowmem_reserve[]: 0 3081 7799 7799
[ +0.000004] Node 0 DMA32 free:72260kB min:25836kB low:32292kB high:38748kB
active_anon:890440kB inactive_anon:56080kB active_file:848188kB inactive_file:51839
                6kB unevictable:0kB isolated(anon):0kB isolated(file):0kB
present:3385984kB managed:3155552kB mlocked:0kB dirty:2316kB writeback:12kB
mapped:122296kB shmem:4553                                6kB
slab_reclaimable:710976kB slab_unreclaimable:21344kB kernel_stack:3088kB
pagetables:19944kB unstable:0kB bounce:0kB free_pcp:4kB local_pcp:0kB free_cma:0kB
                writeback_tmp:0kB pages_scanned:0 all_unreclaimable?
no
[ +0.000006] lowmem_reserve[]: 0 0 4718 4718
[ +0.000003] Node 0 Normal free:69956kB min:39560kB low:49448kB high:59336kB
active_anon:1428340kB inactive_anon:316208kB active_file:1303244kB inactive_file:7
                31448kB unevictable:0kB isolated(anon):0kB
isolated(file):0kB present:4980736kB managed:4831284kB mlocked:0kB dirty:1808kB
writeback:4kB mapped:256308kB shmem:1
79704kB slab_reclaimable:778944kB slab_unreclaimable:43456kB kernel_stack:6160kB
pagetables:49368kB unstable:0kB bounce:0kB free_pcp:164kB local_pcp:0kB free_cm
                a:0kB writeback_tmp:0kB pages_scanned:48
all_unreclaimable? no
[ +0.000006] lowmem_reserve[]: 0 0 0 0
[ +0.000003] Node 0 DMA: 0*4kB 0*8kB 0*16kB 1*32kB (U) 2*64kB (U) 1*128kB (U) 1*256kB (U)
0*512kB 1*1024kB (U) 1*2048kB (M) 3*4096kB (M) = 15904kB
[ +0.000012] Node 0 DMA32: 8531*4kB (UEM) 1709*8kB (UEM) 1011*16kB (UEM) 187*32kB
(UEM) 35*64kB (UEM) 3*128kB (UM) 0*256kB 0*512kB 0*1024kB 0*2048kB 0*4096kB =
                72580kB
[ +0.000013] Node 0 Normal: 11231*4kB (UEM) 2687*8kB (UEM) 122*16kB (UEM) 34*32kB
(UEM) 12*64kB (EM) 0*128kB 0*256kB 0*512kB 0*1024kB 0*2048kB 0*4096kB = 70228
                kB
[ +0.000012] Node 0 hugepages_total=0 hugepages_free=0 hugepages_surp=0
hugepages_size=2048kB
[ +0.000001] 906866 total pagecache pages
[ +0.000002] 313 pages in swap cache
[ +0.000002] Swap cache stats: add 144423, delete 144109, find 6879297/6880886
[ +0.000001] Free swap = 8378108kB
[ +0.000002] Total swap = 8388604kB
[ +0.000001] 2095679 pages RAM
[ +0.000002] 0 pages HighMem/MovableOnly
[ +0.000001] 94994 pages reserved

```

Since I am seeing these kinds of errors on various nodes, I guess it's not a big issue. However I am wondering why this is happening? And if it's something to worry about in general?

Similar error examples:

```
[Feb23 03:37] /usr/sbin/httpd: page allocation failure: order:4, mode:0x2040d0  
[ +0.000029] /usr/sbin/httpd: page allocation failure: order:5, mode:0x2040d0  
[ +0.000776] php-fpm7.2: page allocation failure: order:5, mode:0x2040d0  
[ +0.000005] docker-proxy: page allocation failure: order:4, mode:0x2040d0  
[ +0.000154] php: page allocation failure: order:4, mode:0x2040d0
```

Mostly it's happening with "httpd", but also other occasions happen e.g. php, perl, etc.

Sidenote: it seems also to happen on an older kernel e.g. Linux servername

```
3.10.0-1160.41.1.vz7.183.5 #1 SMP Thu Sep 23 18:26:47 MSK 2021 x86_64 x86_64 x86_64  
GNU/Linux
```

Bean statistics are fine though with no failcnt.

Maybe someone can enlighten me on the above. Is it a serious? Can I do something about it?  
Can someone explain the issue to me?

Thanks in advance.

Regards

---

Subject: Re: /usr/sbin/httpd: page allocation failure?  
Posted by [HHawk](#) on Wed, 16 Mar 2022 18:59:09 GMT  
[View Forum Message](#) <> [Reply to Message](#)

\* BuMP! \*

---

Subject: Re: /usr/sbin/httpd: page allocation failure?  
Posted by [vaverin](#) on Sat, 26 Mar 2022 11:40:49 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Dear HHawk,  
first of all I would like to ask you to address any noticed troubles to [bugs.openvz.org](https://bugs.openvz.org).  
Unlike the forum it is checked by developers more or less regularly.

According to "page allocation failure: order:X, mode:Y" messages: it is not an error.  
It is warning messages controlled by `sysctl vm.warn_high_order`.  
It is used in testing to detect places in code where high-order memory allocations can happen.  
Such allocations can cause significant memory allocation delays in case of memory shortage or in case of huge memory fragmentation on the node.  
We're looking at all found places and trying to handle it somehow.  
In some cases it allows us to detect huge non-accounted memory allocation,  
in other we can replace huge `kmallo` allocation by per-pages `kvmalloc` allocation,  
sometimes we can do nothing here and just hide the warning in the place.

However anyway, it is not an error, and this should not be a cause for your concern.  
Please check state of sysctl vm.warn\_high\_order, obviously it was set to 3, either from /etc/sysctl.conf-like files, or due some mistake in openvz scripts.  
Please set it to 11, it disables the message.

Thank you,  
Vasily Averin

---

---

Subject: Re: /usr/sbin/httpd: page allocation failure?  
Posted by [HHawk](#) on Tue, 05 Apr 2022 12:59:35 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi Vasily,

Didn't expect to see an answer here to be honest, but thank you for your explanation.  
Will make the necessary changes! Thank you.

Sidenote: checked the /etc/sysctl.conf but it had no setting there like vm.warn\_high\_order. Will add it nevertheless.

---