Subject: can't allocate memory
Posted by sander19 on Mon, 23 Apr 2007 16:15:44 GMT
View Forum Message <> Reply to Message

I think I have something wrong, I'm constantly getting memory exhausted errors and stuff, can't allocated memory blabla.

I just try to do some silly commmands but all have too little memory.

I understood that the privvmpages were the amount that a process can take. I've put this high enough but still it's hitting the limits. What am I doing wrong ?

Below you can find my specs:

client w:

root@*****:/# w
Inconsistency detected by ld.so: dl-minimal.c: 84: __libc_memalign: Assertion `page != ((void *) -1)' failed!

server vzctl:

*******:/home/sander# vzctl enter 101
entered into VE 101
Inconsistency detected by ld.so: dl-minimal.c: 84: __libc_memalign: Assertion `page != ((void *) -1)' failed!

client top:

```
  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
26225 clamav    16   0 51016  38m  928 S  0.0  3.8   0:02.04 clamd
26223 amavis    16   0 36032  30m 2188 S  0.0  3.1   0:00.41 amavisd-new
26224 amavis    16   0 35876  30m 2188 S  0.0  3.0   0:00.23 amavisd-new
26207 amavis    16   0 35112  29m 2156 S  0.0  3.0   0:00.27 amavisd-new
26347 mysql     16   0  117m  24m 4812 S  0.0  2.4   0:18.13 mysqld
32526 www-data  16   0 24148  13m 2948 S  0.0  1.3   0:08.88 apache2
32177 www-data  16   0 23472  12m 3056 S  0.0  1.3   0:12.82 apache2
28585 www-data  15   0 23204  12m 3180 S  0.0  1.3   0:06.16 apache2
30533 www-data  16   0 23156  12m 3200 S  0.0  1.3   0:03.45 apache2
32525 www-data  16   0 23172  12m 2956 S  0.0  1.2   0:00.45 apache2
32529 www-data  15   0 23052  12m 2832 S  0.0  1.2   0:00.19 apache2
32528 www-data  16   0 18440 8172 2996 S  0.0  0.8   0:00.34 apache2
```

server beancounters:

Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt
101:	kmemsize	4252389	5210322	7752512	8936012	0
	lockedpages	0	5	32	32	0
	privvmpages	96595	105237	98304	131000	2672
	shmpages	702	1998	8192	8192	0
	dummy	0	0	0	0	0
	numproc	58	73	128	128	6
	physpages	48643	54437	0	2147483647	0
	vmguarpages	0	0	98304	2147483647	0
	oomguarpages	48643	54437	100196	2147483647	0
	numtcpsock	16	59	80	80	0
	numflock	10	21	100	110	0
	numpty	1	3	16	16	0
	numsiginfo	0	26	256	256	0
	tcpsndbuf	6708	297388	319488	524288	0
	tcprcvbuf	4472	175312	319488	524288	0
	othersockbuf	172172	425560	1585883	2937563	137
	dgramrcvbuf	0	20124	132096	132096	0
	numothersock	122	144	520	520	307
	dcachesize	0	0	1048576	1097728	0
	numfile	1423	1731	2048	2048	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	30	90	128	128	0

Free command server:

	total	used	free	shared	buffers	cached
Mem:	1029844	1017256	12588	0	115040	176528
-/+ buffers/cache:		725688	304156			
Swap:	4000168	2072	3998096			

Subject: Re: can't allocate memory
Posted by Vasily Tarasov on Mon, 23 Apr 2007 17:45:17 GMT
View Forum Message <> Reply to Message

You have a resource shortage in your VE. Read more at
http://wiki.openvz.org/Resource_shortage

Vasily.

Subject: Re: can't allocate memory
Posted by sander19 on Mon, 23 Apr 2007 18:23:47 GMT
View Forum Message <> Reply to Message

Yes I can see that too.

The reason is why ! The privvmpages show how much a process my consume ? right ?

I have it set at about 400 MB, and still my processes aren't consuming that much of memory, so why am I always against the limit.

Subject: Re: can't allocate memory
Posted by Vasily Tarasov on Tue, 24 Apr 2007 07:39:32 GMT
View Forum Message <> Reply to Message

Read more about what is _privvmpages_ (http://wiki.openvz.org/Privvmpages#privvmpages). It is not used physical memory, it is _allocated_ _virtual_ memory.

Vasily.

Subject: Re: can't allocate memory
Posted by hvdkamer on Tue, 24 Apr 2007 08:51:05 GMT
View Forum Message <> Reply to Message

As said, the privvmpages is allaocated memory. Most of the times this isn't used at all. For example let's take one of my machines:

```
hoefnix2:~# ubc | grep -E 'priv' | sed -re 's/^(.{40}).*/\1/'
        privvmpages       6240
        privvmpages        913
        privvmpages        360
        privvmpages        378
        privvmpages      46847
        privvmpages      86436
        privvmpages      89116
        privvmpages      56970
        privvmpages      76104
        privvmpages      47447
        privvmpages      34170
        privvmpages      47919
        privvmpages      39446
```

It looks likes this machine is currently using 532.346 pages of 4 KiB is 2.079 MiB of memory. However:

```
hoefnix2:~# free
             total       used       free     shared    buffers     cached
Mem:       2053724    2020396      33328          0      51568    1475556
-/+ buffers/cache:     493272    1560452
Swap:      2104472     265072    1839400
```

excluding buffers, cache and shared it is only consuming 493.272 KiB of memory. You can roughly find this amount through physpages:

```
hoefnix2:~# ubc | grep -E 'phys' | sed -re 's/^(.{40}).*/\1/'
          physpages         1181
          physpages          202
          physpages          199
          physpages          232
          physpages         9384
          physpages         5168
          physpages        40408
          physpages        15069
          physpages         2171
          physpages         7003
          physpages         1595
          physpages          523
          physpages         3843
```

or 71.909 pages of 4 KiB is 287.636 KiB. The difference between this and 493.272 is allocated memory which is really in use. The difference between 2.129.384 and 493.272 is allocated memory which isn't used. Or at least not at this moment. That is why OpenVZ is so eeficient with hardware. You can use more memory than is strictly available .

Memory usage in modern OSses is a very difficult thing. For example an image manipulation program could ask for say three times the memory needed for reading an image. The reserve is then used for undo buffers and things like that. For a programmer it is nice to ask for more than is needed. If given you can use it as you want and is in general more efficient. The drawback is that most times al lot of memory is alloacted but not used. When in need for more mmemory, it can be efficintly swapped.

Now for your problem:

```
          privvmpages      96595     105237      98304     131000       2672
          physpages        48643      54437          0 2147483647          0
```

As you can see your machine uses at this moment 48.643 pages of 4 KiB and in the past has used up to 54.437 pages. That is more than the most consuming VE on my machine. And that one can ask up till 131.072 as barrier and 147.456 as limit. So I think that the solution to your problem is lowering the physical amount used or raise the amount that can be allocated. In my experience the factor between these two parameters is between 3 and 5 with te lower in more consuming VE's. So with 54.437 I think you must raise the privvmpages to 163.840 (is 640 MiB). You could try 131.072 (512 MiB) and the optimizing the amount of started processes...
