
Subject: Where has my memory gone?

Posted by [gnutered](#) on Fri, 25 Jul 2008 14:49:40 GMT

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

If I add up my memory usage in each VE, it adds up to a lot less than what is (apparently) actually being consumed.

This is causing me grief on a production server, in that I'm running out of RAM, and going into swap, which is really hurting my web server - it slows to a crawl, then processes back up, and memory usage really spirals up.

Evidence:

```
# for v in ` cd /proc/bc ; echo * `; do echo $v; vzctl exec $v free -m; done
```

```
---- VE: 100
```

	total	used	free	shared	buffers	cached
Mem:	256	9	246	0	0	0
-/+ buffers/cache:		9	246			
Swap:	0	0	0			

```
---- VE: 101
```

	total	used	free	shared	buffers	cached
Mem:	921	493	427	0	0	0
-/+ buffers/cache:		493	427			
Swap:	0	0	0			

```
---- VE: 102
```

	total	used	free	shared	buffers	cached
Mem:	460	58	402	0	0	0
-/+ buffers/cache:		58	402			
Swap:	0	0	0			

```
---- VE: 106
```

	total	used	free	shared	buffers	cached
Mem:	345	20	324	0	0	0
-/+ buffers/cache:		20	324			
Swap:	0	0	0			

```
---- VE: 107
```

	total	used	free	shared	buffers	cached
Mem:	115	37	77	0	0	0
-/+ buffers/cache:		37	77			
Swap:	0	0	0			

Add up the "used" columns, and I get around 617MB

On the HN:

```
# free -m
```

	total	used	free	shared	buffers	cached
Mem:	1945	1901	43	0	40	100
-/+ buffers/cache:		1761	184			
Swap:	6143	281	5862			

It is showing 1761MB in use.

If I do some maths on the output of `/proc/bc/*`, I get similar results:

```
# ~/beano < /proc/user_beancounters
VE 107: priv = 36MB, kmem = 0MB
VE 106: priv = 20MB, kmem = 1MB
VE 102: priv = 58MB, kmem = 1MB
VE 101: priv = 497MB, kmem = 11MB
VE 100: priv = 9MB, kmem = 0MB
VE 0: priv = 39MB, kmem = 7MB
```

This totals 679MB (note, this was run a few minutes after the above maths, so numbers differ slightly).

So, why the large discrepancy? Where's my missing 1GB of memory gone?

Caveats: I know I'm probably not doing the maths 100% accurately, but it's near enough to know that something isn't right.

Detail:

Here's `/proc/user_beancount`:

```
# cat /proc/user_beancounters
```

```
Version: 2.5
```

uid	resource	held	maxheld	barrier	limit	failcnt
107:	kmemsize	903172	2714806	14372700	14790164	
0						
	lockedpages	0	0	256	256	0
	privvmpages	9457	20561	29491	32768	
0						
	shmpages	0	1296	21504	21504	0
	dummy	0	0	0	0	0
	numproc	9	22	240	240	0
	physpages	493	8769	0	9223372036854775807	
0						
	vmguarpages	0	0	33792	9223372036854775807	
0						
	oomguarpages	1165	8957	26112	9223372036854775807	
0						
	numtcpsock	6	12	360	360	0
	numflock	1	6	188	206	0
	numpty	0	3	16	16	0
	numsiginfo	0	2	256	256	0
	tcpsndbuf	104640	165944	1720320	2703360	

0	tcprcvbuf	98304	401464	1720320	2703360	
0	othersockbuf	11560	22328	1126080	2097152	
0	dgramrcvbuf	0	34880	262144	262144	
0	numothersock	7	16	360	360	0
0	dcachesize	98595	193317	3409920	3624960	
0	numfile	221	471	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	14	14	128	128	0
106:	kmemsize	1798197	3523117	14372700	14790164	
0	lockedpages	0	0	256	256	0
0	privvmpages	5189	19476	88473	98304	
0	shmpages	18	674	21504	21504	0
	dummy	0	0	0	0	0
	numproc	8	22	240	240	0
	physpages	116	5050	0	9223372036854775807	
0	vmguarpages	0	0	29491	32768	0
	oomguarpages	1588	6541	26112	9223372036854775807	
0	numtcpsock	3	11	360	360	0
	numflock	1	2	188	206	0
	numpty	0	1	16	16	0
	numsiginfo	0	2	256	256	0
0	tcpsndbuf	52320	338296	1720320	2703360	
0	tcprcvbuf	49152	0	1720320	2703360	0
0	othersockbuf	2312	9808	1126080	2097152	
0	dgramrcvbuf	0	8456	262144	262144	0
	numothersock	1	9	360	360	0
	dcachesize	131460	201016	3409920	3624960	
0	numfile	202	430	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	14	14	128	128	0
102:	kmemsize	1725889	4340821	14372700	14790164	
0						

	lockedpages	0	0	256	256	0
	privvmpages	14717	39041	117964	131072	
0	shmpages	0	3216	21504	21504	0
	dummy	0	0	0	0	0
	numproc	9	26	240	240	0
	physpages	9404	25165	0	9223372036854775807	
0	vmguarpages	0	0	29491	32768	0
	oomguarpages	20698	28347	26112	9223372036854775807	
0	numtcpsock	25	360	360	360	2775
	numflock	1	3	188	206	0
	numpty	3	5	16	16	0
	numsiginfo	0	2	256	256	0
	tcpsndbuf	213768	2029080	1720320	2703360	
14887573	tcprcvbuf	202448	2002936	1720320	2703360	
2513133	othersockbuf	2312	322376	1126080	2097152	
0	dgramrcvbuf	0	8456	262144	262144	0
	numothersock	1	35	360	360	0
	dcachesize	181125	444357	3409920	3624960	
0	numfile	771	1103	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	14	14	128	128	0
101:	kmemsize	11635745	46486929	134217728	147639500	
0	lockedpages	0	0	256	256	0
	privvmpages	128074	263359	235929	262144	
165044	shmpages	8512	15568	21504	21504	
0	dummy	0	0	0	0	0
	numproc	65	159	400	400	0
	physpages	36906	85441	0	9223372036854775807	
0	vmguarpages	0	0	58982	65536	0
	oomguarpages	89615	141348	58982	65536	
0	numtcpsock	32	173	360	360	0
	numflock	9	21	188	206	0
	numpty	1	2	16	16	0
	numsiginfo	0	78	256	256	0

113	tcpsndbuf	215952	1724304	1720320	2703360	
0	tcprcvbuf	210216	710136	1720320	2703360	
0	othersockbuf	18496	179904	1126080	2097152	
0	dgramrcvbuf	0	23648	262144	262144	
0	numothersock	20	93	360	360	0
0	dcachesize	452598	620679	3409920	3624960	
0	numfile	1569	2534	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	14	14	128	128	0
100:	kmemsize	788463	2142311	14372700	14790164	
0	lockedpages	0	0	256	256	0
0	privvmpages	2372	16843	58982	65536	
0	shmpages	0	656	21504	21504	0
	dummy	0	0	0	0	0
	numproc	6	19	240	240	0
0	physpages	133	7774	0	9223372036854775807	
0	vmguarpages	0	0	14745	16384	0
0	oomguarpages	1915	8764	26112	9223372036854775807	
	numtcpsock	4	10	360	360	0
	numflock	1	5	188	206	0
	numpty	0	1	16	16	0
	numsiginfo	0	3	256	256	0
0	tcpsndbuf	69760	353160	1720320	2703360	
0	tcprcvbuf	65536	484104	1720320	2703360	
0	othersockbuf	11560	34840	1126080	2097152	
0	dgramrcvbuf	0	25168	262144	262144	
0	numothersock	6	14	360	360	0
0	dcachesize	17841	59157	3409920	3624960	
0	numfile	153	388	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0

numiptent	10	10	128	128	0
0: kmemsize	8083585	9508188	9223372036854775807		
9223372036854775807	0				
lockedpages	0	5209	9223372036854775807		
9223372036854775807	0				
privvmpages	10142	20783	9223372036854775807		
9223372036854775807	0				
shmpages	1927	2599	9223372036854775807		
9223372036854775807	0				
dummy	0	0	9223372036854775807	9223372036854775807	
0					
numproc	78	88	9223372036854775807		
9223372036854775807	0				
physpages	1885	15036	9223372036854775807		
9223372036854775807	0				
vmguarpages	0	0	9223372036854775807		
9223372036854775807	0				
oomguarpages	6012	25423	9223372036854775807		
9223372036854775807	0				
numtcpsock	8	33	9223372036854775807		
9223372036854775807	0				
numflock	2	8	9223372036854775807	9223372036854775807	
0					
numpty	3	4	9223372036854775807	9223372036854775807	
0					
numsiginfo	1	5	9223372036854775807	9223372036854775807	
0					
tcpsndbuf	160128	177568	9223372036854775807		
9223372036854775807	0				
tcprcvbuf	131072	330576	9223372036854775807		
9223372036854775807	0				
othersockbuf	32368	116888	9223372036854775807		
9223372036854775807	0				
dgramrcvbuf	0	65688	9223372036854775807		
9223372036854775807	0				
numothersock	43	56	9223372036854775807		
9223372036854775807	0				
dcachesize	1936469	2800110	9223372036854775807		
9223372036854775807	0				
numfile	1766	2793	9223372036854775807		
9223372036854775807	0				
dummy	0	0	9223372036854775807	9223372036854775807	
0					
dummy	0	0	9223372036854775807	9223372036854775807	
0					
dummy	0	0	9223372036854775807	9223372036854775807	
0					
numiptent	306	321	9223372036854775807		

Here's my beano script:

```
#!/usr/bin/awk -f
$1 ~ /[1-9][0-9]*:/ { veid = substr($1, 0, length($1)) }
$1 ~ /0:/ { veid = substr($1, 0, length($1)) }
$2 ~ /kmemsize/ { kmem = int($3 / 1024 / 1024) ; kmem_mh = int($3 / 1024 / 1024) }
$1 ~ /privvmpages/ { priv = int($2 / 256) ; priv_mh = int($2 / 256) }
$1 ~ /numiptent/ { printf ("VE %s: priv = %dMB, kmem = %dMB\n", veid, priv, kmem) }
```

Operating system is Debian etch. Kernel is:

```
# uname -a
```

```
Linux alpha 2.6.18-12-fza-amd64 #1 SMP Sun May 18 09:57:08 UTC 2008 x86_64 GNU/Linux
```

Packages:

```
# dpkg -l \*vz\*
```

```
Desired=Unknown/Install/Remove/Purge/Hold
```

```
| Status=Not/Installed/Config-files/Unpacked/Failed-config/Half-installed
```

```
|| Err?=(none)/Hold/Reinst-required/X=both-problems (Status,Err: uppercase=bad)
```

```
||/ Name Version Description
```

```
+++=====
=====
=====
un ovzkernel <none> (no description available)
ii vzctl 3.0.22-1dso1 OpenVZ - server virtualization solution -
control tools
un vzctl-ostemplate <none> (no description available)
ii vzctl-ostmpl-debian 4.0-20071003 OpenVZ - OS Template - Debian 4.0
un vzctl-template-debian <none> (no description available)
ii vzdump 1.0-2 OpenVZ backup scripts
ii vzprocps 2.0.11-1dso2 OpenVZ - server virtualization solution -
procps utilities
ii vzquota 3.0.11-1dso1 OpenVZ - server virtualization solution -
quota tools
```

Subject: Re: Where has my memory gone?

Posted by [kenjy](#) on Sat, 26 Jul 2008 16:23:26 GMT

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

Please try using `vzmemcheck` and `vzmemcheck -vA` in HN, take care of `VMGUARPAGES`, `OOMGUARPAGES` and `PRIVVMPAGES`

Subject: Re: Where has my memory gone?
 Posted by [gnutered](#) on Sun, 27 Jul 2008 12:19:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks.

My output is thus:

```
$ sudo cat /proc/bc/101/resources && sudo vzmemcheck -v
    kmemsize          12755129          46486929          134217728          147639500
    0
    lockedpages        0                  0                 256                256                0
    privvmpages        133586            263359            235929             262144
208779
    shmpages           9152              15568             21504              21504
0
    numproc            71                159               400                400                0
    physpages          36617             85441             0 9223372036854775807
0
    vmguarpages        0                 0                 58982              65536              0
    oomguarpages       92833            151046            58982              65536
0
    numtcpsock         34                173               360                360                0
    numflock           9                 21                188                206                0
    numpty             0                 2                 16                 16                 0
    numsignifo         0                 78                256                256                0
    tcpsndbuf          378856            1724304           1720320            2703360
113
    tcprcvbuf          214840            710136            1720320            2703360
0
    othersockbuf       27744             179904            1126080            2097152
0
    dgramrcvbuf        0                 25960             262144             262144
0
    numothersock       29                93                360                360                0
    dcachesize         458232            620679            3409920            3624960
0
    numfile            1688              2534              9312               9312              0
    numiptent          14                14                128                128                0
```

Output values in %

veid	LowMem util	LowMem commit	LowMem util	RAM util	MemSwap commit	MemSwap util	MemSwap commit	Alloc limit	Alloc	Alloc
107	0.14	2.76	0.14	0.07	1.53	0.47	1.90	1.85		
106	0.23	2.76	0.11	0.09	1.53	0.27	1.69	5.01		
102	0.31	2.76	2.10	1.05	1.53	0.76	1.69	6.59		
101	1.64	19.04	8.01	4.64	4.68	6.61	4.68	14.49		
100	0.11	2.76	0.06	0.10	1.53	0.12	0.98	3.43		

 Summary: 2.42 30.10 10.42 5.95 10.79 8.23 10.94 31.38

I use this page to interpret these results: http://wiki.openvz.org/UBC_systemwide_configuration

Ratios on 101 (my big server) look dire. Ignore lowmem, as this is an AMD64.

RAM on VE 101 is 8.01. If I add up the numbers $\text{physpages} * 4096 + \text{kmemsize} + \text{tcpsndbuf} + \text{tcpvbuf} + \text{dgramrcvbuf} + \text{othersockbuf}$, I get 163359801, which is 155.8MB. For the ratio to be 8.01, the system thinks (or treats) my server like it has 19.45MB of RAM. Maths on VE 106 (arbitrarily chosen) shows similar.

Something isn't right.

I mention again that I'm running the FZA kernel (from the debian-systs repository), which I chose for reasons that no longer apply. I can go back to the slightly more official kernel if something is broken inside the kernel, but I'm keen to avoid reboots, as this is a production server on the other side of the world.

Other than that, what can I try?

Subject: Re: Where has my memory gone?
Posted by [kenjy](#) on Mon, 28 Jul 2008 05:23:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Well I can see that you have a big problem with `privvmpages` and a small one with `tcpsndbuf`, you need to increase these values in your VE 101, please use `vzcfgvalidate /PATH/VE.conf` to check the config, don't think in this way "I have 10gb of ram and this VE is going to have 2gb of ram ...", you can't set a specific amount of RAM for a VE because the RAM for a computer is not a piece of something.