
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
208779	privvmpages	133586	263359	235929	262144
0	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
	numflock	9	21	188	206
	numpty	0	2	16	16
	numsiginfo	0	78	256	256
113	tcpsndbuf	378856	1724304	1720320	2703360
0	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
	numiptent	14	14	128	128

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 $\text{physpages} \times 4096 + \text{kmemsize} + \text{tcpndbuf} + \text{tcpvbuf} + \text{dgramrcvbuf} + \text{othersockbuf}$, 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?