
Subject: Containers not using vswap

Posted by [huan](#) on Fri, 14 Sep 2012 17:59:09 GMT

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

I've moved some containers to CentOS6 hardware node and decided to use VSWAP. Apparently it is not being used by containers.

Here's an example. One container configured with 4GB RAM and 2GB vswap but when it starts to use 4GB the oom killer starts to kill processes and swap is 0kb.

System is CentOS 6. Kernel 2.6.32-042stab059.7 #1 SMP x86_64

vzctl version 3.3

Errors from /var/log/messages:

Sep 14 05:21:44 vps3 kernel: [574167.843920] UB-134-Mem-Info:

Sep 14 05:21:44 vps3 kernel: [574167.843967] Node 0 DMA scan:0 a_anon:0 i_anon:0 a_file:0 i_file:0 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574167.844061] Node 0 DMA32 scan:0 a_anon:1754 i_anon:817 a_file:4 i_file:9 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574167.844174] Node 0 Normal scan:0 a_anon:1012882 i_anon:4865 a_file:0 i_file:20 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574167.844270] Total 1020351 anon:1020318 file:33 a_anon:1014636 i_anon:5682 a_file:4 i_file:29 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574167.844384] RAM: 1042872 / 1048576 [0] SWAP: 0 / 536576 [0] KMEM: 91000832 / inf [0] DCSZ: 70918284 / inf [0] OOMG: 1020197 / inf [0] Dirty 325 Wback 0 Dche 62544 Prnd 1544378

Sep 14 05:21:44 vps3 kernel: [574167.844648] Out of memory in UB: Kill process 244305 (mysqld) score 948 or sacrifice child

Sep 14 05:21:44 vps3 kernel: [574167.844822] OOM killed process 244305 (mysqld) vm:4264920kB, rss:3971912kB, swap:0kB

Sep 14 05:21:44 vps3 kernel: [574168.091829] OOM killed process mysqld (pid=244321, ve=134) exited, free=1019989.

Sep 14 05:21:44 vps3 kernel: [574168.091927] UB-134-Mem-Info:

Sep 14 05:21:44 vps3 kernel: [574168.091974] Node 0 DMA scan:0 a_anon:0 i_anon:0 a_file:0 i_file:0 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574168.092077] Node 0 DMA32 scan:0 a_anon:1712 i_anon:14 a_file:4 i_file:9 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574168.092180] Node 0 Normal scan:0 a_anon:20757 i_anon:4865 a_file:0 i_file:43 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574168.092275] Total 27404 anon:27348 file:56 a_anon:22469 i_anon:4879 a_file:4 i_file:52 unevictable:0

Sep 14 05:21:44 vps3 kernel: [574168.092377] RAM: 58527 / 1048576 [0] SWAP: 0 / 536576 [0] KMEM: 123826176 / inf [0] DCSZ: 112364344 / inf [0] OOMG: 1020197 / inf [1] Dirty 325 Wback 0 Dche 107859 Prnd 1544378

Container config file /etc/vz/134.conf:

```
ONBOOT="yes"
# UBC parameters (in form of barrier:limit)
KMEMSIZE="unlimited"
LOCKEDPAGES="unlimited"
PRIVVMPAGES="unlimited"
SHMPAGES="unlimited"
NUMPROC="unlimited"
PHYSPAGES="0:1048576"
VMGUARPAGES="unlimited"
OOMGUARPAGES="unlimited"
NUMTCPSOCK="unlimited"
NUMFLOCK="unlimited"
NUMPTY="unlimited"
NUMSIGINFO="unlimited"
TCPSNDBUF="unlimited"
TCPRCVBUF="unlimited"
OTHERSOCKBUF="unlimited"
DGRAMRCVBUF="unlimited"
NUMOTHERSOCK="unlimited"
DCACHESIZE="unlimited"
NUMFILE="unlimited"
AVNUMPROC="1000:1000"
NUMIPTENT="unlimited"

# Disk quota parameters (in form of softlimit:hardlimit)
DISKSPACE="157286400:157286400"
DISKINODES="78643200:78643200"
QUOTATIME="0"

# CPU fair scheduler parameter
CPUUNITS="1000"
VE_ROOT="/vz2/root/$VEID"
VE_PRIVATE="/vz2/private/$VEID"
OSTEMPLATE="debian-6.0-x86_64"
ORIGIN_SAMPLE="basic"
IP_ADDRESS="xx.xx.xx.xx"
HOSTNAME="xx.pl"
NAMESERVER="213.133.99.99 8.8.8.8"
MEMINFO="pages:2097152"
QUOTAUGIDLIMIT="9223372036854775807"
IOPRIO="4"
CPULIMIT="147"
CPUS="2"
SWAPPAGES="0:536576"
```

What am I doing wrong?

Subject: Re: Containers not using vswap
Posted by [huan](#) on Sat, 15 Sep 2012 14:39:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

vzubb 134

CT 134 | HELD Bar% Lim%| MAXH Bar% Lim%| BAR | LIM | FAIL

-----+-----+-----+-----+-----
kmemsize|52.2M - -| 963M - -| - | - | -
lockedpages| - - -| - - -| - | - | -
privvmpages|3.02G - -| 4.4G - -| - | - | -
shmpages| 130M - -| 134M - -| - | - | -
numproc| 60 - -| 96 - -| - | - | -
physpages|2.73G - 68%| 4G - 100%| - | 4G| -
vmguarpages| - - -| - - -| - | - | -
oomguarpages|2.54G - -| 3.9G - -| - | - | 1
numtcpsock| 29 - -| 98 - -| - | - | -
numflock| 7 - -| 13 - -| - | - | -
numpty| - - -| 9 - -| - | - | -
numsiginfo| 4 - -| 36 - -| - | - | -
tcpndbuf| 553K - -| 1.71M - -| - | - | -
tcprcvbuf| 464K - -| 1.53M - -| - | - | -
othersockbuf|24.1K - -| 453K - -| - | - | -
dgramrcvbuf| - - -| 4.52K - -| - | - | -
numothersock| 58 - -| 206 - -| - | - | -
dcachesize|37.7M - -| 948M - -| - | - | -
numfile|10.5K - -| 39.3K - -| - | - | -
numiptent| 24 - -| 24 - -| - | - | -
swappages| - - -| - - -| - | 2.05G| -

Apparently there is no limit set for 'oomguarpages' but it did fail in the past. Strange.
