
Subject: PAE support inside VE
Posted by [eastitaly](#) on Fri, 08 Jan 2010 08:56:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

I've installed OpenVZ (64 bit kernel - CentOS 5.4) on some servers with 8+ GB of RAM.

I've also set up some VEs with a CentOS 5.4 32 bit template and all worked fine until now.

Now the available memory needed for some VEs is increased (for traffic and load) and I need to allocate more than 4 GB for them.

I was wondering if the VE could use more than 4 GB of RAM, then I tried and got a positive answer.

The problem comes as the maximum memory available inside that container seems to be limited to exactly 4613120 Kb (as per /proc/meminfo), even if I try to set the VE's memory allocation to 5 or even 6 GB.

The kernel I'm using for the HN is 2.6.18-128.2.1.el5.028stab064.8 and I've checked the release notes for the subsequent releases and didn't find anything related to this problem.

Any hint?

Subject: Re: PAE support inside VE
Posted by [nuno](#) on Sat, 09 Jan 2010 09:24:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

I just made a test in one host (amd64, 8GB ram) and set everything to:

9223372036854775807:9223372036854775807

```
root@proxy:~# head /proc/meminfo
```

```
MemTotal: 8169084 kB
MemFree: 8154528 kB
Buffers: 0 kB
Cached: 0 kB
SwapCached: 0 kB
Active: 0 kB
Inactive: 0 kB
SwapTotal: 0 kB
SwapFree: 0 kB
Dirty: 0 kB
```

```
root@proxy:~#
```

```
root@proxy:~# free -m
```

	total	used	free	shared	buffers	cached
Mem:	7977	14	7963	0	0	0

```
-/+ buffers/cache:    14    7963
Swap:                0      0      0
root@proxy:~#
```

hmmm...

Subject: Re: PAE support inside VE
Posted by [eastitaly](#) on Mon, 11 Jan 2010 08:53:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Please, could you post some more info about your system (HN kernel version, distro and the VE's template used on that server) ?

Subject: Re: PAE support inside VE
Posted by [nuno](#) on Mon, 11 Jan 2010 09:38:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

2.6.24-26-openvz #1 SMP Tue Dec 1 20:30:40 UTC 2009 x86_64 GNU/Linux

The templates don't matter because that won't influence /proc/meminfo's output.

Are you sure you have a 64bit kernel? `uname -a`

Subject: Re: PAE support inside VE
Posted by [eastitaly](#) on Mon, 11 Jan 2010 10:07:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

the HN reports

Linux xxxxxxx 2.6.18-128.2.1.el5.028stab064.8 #1 SMP Fri Nov 6 12:26:59 MSK 2009 x86_64
x86_64 x86_64 GNU/Linux

and the VE reports

Linux yyyyyyy 2.6.18-128.2.1.el5.028stab064.8 #1 SMP Fri Nov 6 12:26:59 MSK 2009 i686 i686
i386 GNU/Linux

Subject: Re: PAE support inside VE
Posted by [eastitaly](#) on Thu, 14 Jan 2010 09:04:44 GMT

Update:

Moved that VE on another HN with a new kernel, same distribution (CentOS 5.4 x86_64) with 24 GB of RAM

uname -a

```
Linux zzzzzz 2.6.18-164.2.1.el5.028stab066.10 #1 SMP Sat Dec 12 18:52:53 MSK 2009 x86_64
x86_64 x86_64 GNU/Linux
```

/proc/meminfo

```
MemTotal: 24624960 kB
MemFree: 288368 kB
Buffers: 1713548 kB
Cached: 20170624 kB
SwapCached: 0 kB
Active: 5424216 kB
Inactive: 17507444 kB
HighTotal: 0 kB
HighFree: 0 kB
LowTotal: 24624960 kB
LowFree: 288368 kB
SwapTotal: 2097144 kB
SwapFree: 2096976 kB
Dirty: 1052500 kB
Writeback: 140 kB
AnonPages: 1048552 kB
Mapped: 174452 kB
Slab: 1310100 kB
PageTables: 27180 kB
NFS_Unstable: 0 kB
Bounce: 0 kB
CommitLimit: 14409624 kB
Committed_AS: 7696144 kB
VmallocTotal: 34359738364 kB
VmallocUsed: 286208 kB
VmallocChunk: 34359451720 kB
HugePages_Total: 0
HugePages_Free: 0
HugePages_Rsvd: 0
Hugepagesize: 2048 kB
```

VE reports:

/proc/meminfo

```
MemTotal: 4613120 kB
MemTotal: 4613120 kB
MemFree: 2745300 kB
Buffers: 0 kB
Cached: 0 kB
```

```

SwapCached:      0 kB
Active:          0 kB
Inactive:        0 kB
HighTotal:       0 kB
HighFree:        0 kB
LowTotal:        4613120 kB
LowFree:         2745300 kB
SwapTotal:       0 kB
SwapFree:        0 kB
Dirty:           0 kB
Writeback:       0 kB
AnonPages:       0 kB
Mapped:          0 kB
Slab:            0 kB
PageTables:      0 kB
NFS_Unstable:    0 kB
Bounce:          0 kB
CommitLimit:     0 kB
Committed_AS:   0 kB
VmallocTotal:    0 kB
VmallocUsed:     0 kB
VmallocChunk:    0 kB
HugePages_Total: 0
HugePages_Free:  0
HugePages_Rsvd:  0
Hugepagesize:    2048 kB

```

/proc/user_beancounters

resource	held	maxheld	barrier	limit	failcnt	
kmemsize	32875448	47104131	95000000	95000000	0	0
lockedpages	0	7	256	256	0	
privvmpages	540234	824747	1310720	1310720	0	0
shmpages	2576	2912	21504	21504	0	
dummy	0	0	0	0	0	
numproc	186	248	880	880	0	
physpages	148452	256702	0	9223372036854775807	0	0
vmguarpages	0	0	524288	9223372036854775807	0	0
oomguarpages	148452	256702	26112	9223372036854775807	0	0
numtcpsock	149	278	1000	1000	0	
numflock	6	30	188	206	0	
numpty	1	1	16	16	0	
numsiginfo	0	49	384	384	0	
tcpsndbuf	1231840	3433336	8800000	11360000	0	0
tcprcvbuf	1245280	1990232	10000000	12560000	0	0
othersockbuf	148384	798960	2831155	3752755	0	0
dgramrcvbuf	0	133424	330000	330000	0	
numothersock	87	122	360	360	0	
dcachesize	2986196	3179234	6446000	11520000	0	0

numfile	6381	9414	30000	30000	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	29	29	128	128	0

1310720 4K-pages equals exactly 5G (or 5242880 kB), instead of the 4613120 kB reported by /proc/meminfo.

Problem NOT solved...

Subject: Re: PAE support inside VE
Posted by [dedicados](#) on Fri, 15 Jan 2010 03:05:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

VE still being 32bit, why you dont use a 64bit template, maybe that helps

Subject: Re: PAE support inside VE
Posted by [eastitaly](#) on Fri, 15 Jan 2010 09:13:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

That VE is in production and can't be changed.
Understanding the origin of this issue will be helpful for avoiding further mistakes in choosing the template for new VEs.

Quote:VE still being 32bit, why you dont use a 64bit template, maybe that helps

VE's template is 32 bit, correct, but currently seems it's supporting something like PAE (4613120 Kb is more than the limit for the 32 bit architecture).
This is, btw, the title of this post.

So the question remains...

Subject: Re: PAE support inside VE
Posted by [nuno](#) on Fri, 15 Jan 2010 09:32:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Show us the VE.conf file. I suspect you're setting the MEMINFO option.

Regards,

Subject: [SOLVED] Re: PAE support inside VE
Posted by [eastitaly](#) on Fri, 15 Jan 2010 09:48:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
# Copyright (C) 2000-2008, Parallels, Inc. All rights reserved.
#
# This program is free software; you can redistribute it and/or modify
# it under the terms of the GNU General Public License as published by
# the Free Software Foundation; either version 2 of the License, or
# (at your option) any later version.
#
# This program is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
# GNU General Public License for more details.
#
# You should have received a copy of the GNU General Public License
# along with this program; if not, write to the Free Software
# Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
#
```

```
ONBOOT="yes"
```

```
# UBC parameters (in form of barrier:limit)
KMEMSIZE="95000000:95000000"
LOCKEDPAGES="256:256"
PRIVVMPAGES="1310720:1310720"
SHMPAGES="21504:21504"
NUMPROC="880:880"
PHYSAPAGES="0:9223372036854775807"
VMGUARPAGES="524288:9223372036854775807"
OOMGUARPAGES="26112:9223372036854775807"
NUMTCPSOCK="1000:1000"
NUMFLOCK="188:206"
NUMPTY="16:16"
NUMSIGINFO="384:384"
TCPSNDBUF="8800000:11360000"
TCPRCVBUF="10000000:12560000"
OTHERSOCKBUF="2831155:3752755"
DGRAMRCVBUF="330000:330000"
NUMOTHERSOCK="360:360"
DCACHESIZE="6446000:11520000"
NUMFILE="30000:30000"
AVNUMPROC="180:180"
NUMIPTENT="128:128"
```

```
# Disk quota parameters (in form of softlimit:hardlimit)
DISKSPACE="483183820:483183820"
DISKINODES="120795955:134217728"
```

```
QUOTATIME="0"

# CPU fair sheduler parameter
CPUUNITS="1000"

# Network customization section
CONFIG_CUSTOMIZED="yes"
VETH_IP_ADDRESS="aaa.aaa.aaa.aaa/bb"
VE_DEFAULT_GATEWAY="ccc.ccc.ccc.ccc"
BRIDGEDEV="br0"
NAMESERVER="ddd.ddd.ddd.ddd"
HOSTNAME="name.domain.tld"
NAME=""
DESCRIPTION=""

VE_ROOT="/vz/root/ct$VEID/data"
VE_PRIVATE="/vz/private/ct$VEID/data"
OSTEMPLATE="centos-5-i386-minimal"
ORIGIN_SAMPLE="vps.basic"
NETIF="ifname=eth0,bridge=br0,mac=xx:xx:xx:xx:xx:xx,host_ifname=veth101.0,host_mac=yy:yy:yy:yy:yy:yy"
QUOTAUGIDLIMIT="15000"
CAPABILITY="SYS_PACCT:on "
MEMINFO="pages:1153280"
NOATIME="yes"
#DEVNODES="loop0:rw "
CPUS="4"
```

Found the MEMINFO option as you suggested, nuno.
I haven't ever set it, was an addition from an installation of HyperVM.

Removed it and all started magically work as expected.
This, btw, confirms that a 32 bit template is PAE-like enabled.

Thank you so much for addressing me to the right direction!

Subject: Re: PAE support inside VE
Posted by [nuno](#) on Fri, 15 Jan 2010 10:01:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

No problem.

Please note that those values are merelly information. You can define that to 1000 pages, /proc/meminfo will report 4MB of memory, free will report 0% free memory, etc, but your applications will continue to work. The reverse is also true: you can set a huge value and

application will fail after exceeding kmemsize, vmguar* or some other limit.
