

---

Subject: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [snoozer](#) on Thu, 09 Aug 2007 18:00:25 GMT

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

---

hi,

i am completely puzzled by the configuration options for memory. i am looking for a way to use 6 VM's sharing the resources of an P4 2GHz with 1 GB Ram and 2 GB Swap. i keep having trouble with error messages like "cannot allocate memory". but only sometimes and not other times, its really frustrating not to find a setting which keeps everything going. to make this clear, i am NOT looking for a all-resources divided by 6 solution, i am more looking for a nice way to SHARE resources between all 6 VM's. the very strange thing is that the VM's do not use any swap.

i am looking forward to any help.

PS: if i do everything on an old cobalt raq with 32MB Ran what i am doing on the VM, i never get a "cannot allocate memory", everything just takes ages. can that behavior somehow be replicated on the VM's ?

regards

Jan

---

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [vaverin](#) on Fri, 10 Aug 2007 06:31:21 GMT

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

---

He Jan

snoozer wrote on Thu, 09 August 2007 22:00i am completely puzzled by the configuration options for memory. i am looking for a way to use 6 VM's sharing the resources of an P4 2GHz with 1 GB Ram and 2 GB Swap.

You can try to use vzsplitt utility -- it splits the Hardware Node into equal parts and generate a sample VE configuration file

Then you can tune this file four your requirements.

snoozer wrote on Thu, 09 August 2007 22:00 i keep having trouble with error messages like "cannot allocate memory".

it looks like a shortage for privvmpages resource, this parameter limits the memory allocations. You can try to increase its value: a lot of applications do not use the memory efficiently and, sometimes, allocated memory will newer be used later.

snoozer wrote on Thu, 09 August 2007 22:00the very strange thing is that the VM's do not use any swap.

VE shares system memory -- and therefore some part of memory used by applications inside VE can be swapped.

However do you probably mean that swap is not used on your HW node? It is sign that your VE are not used all node resources -- probably due too low values in VE config files. I hope that vzsplite helps you to create proper VE configuration files for your node.

Thank you,  
Vasily Averin

---

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap  
Posted by [snoozier](#) on Fri, 10 Aug 2007 06:41:51 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

thanks a lot,

i will try vzsplite and give feedback.

regards  
Jan

---

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap  
Posted by [locutius](#) on Sat, 18 Aug 2007 01:49:16 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

ve-vps.basic.conf-sample appears to have an insane values here:-  
VMGUARPAGES="6144:2147483647"  
PRIVVMPAGES="49152:53575"

vzsplite on a server with 2GB RAM split 16 ways gives:-  
VMGUARPAGES="310949:2147483647"  
PRIVVMPAGES="310949:342043"

vzsplite on a server with 2GB RAM split 32 ways gives:-  
VMGUARPAGES="45759:2147483647"  
PRIVVMPAGES="274554:302009"

the default config cannot run apache, mysql and ssh without constant '... fork: Cannot allocate memory' errors

i dont know how you guys are using your VE's? for playing pong?

i have to say after my experience of using it for a week i love OpenZV it is one of the best pieces of software i have ever happened upon, right up there with nx it is going to drive a revolution

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [dowdle](#) on Sat, 18 Aug 2007 18:53:31 GMT

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

---

Just to clarify about the vps.basic sample. It is my understanding that the values you will find there were meant to split a reasonably fast machine with 2GB of RAM into 120 VPSes. The vps.light one was designed for 320 VPSes.

The basic will run most of the default OS Templates just fine... that include sshd and apache but not much else. If you add PHP and/or a database... and/or other services, you'll definitely need to up the resources... and yes, vzsplite is a good place to start with creating custom config files... but you'll still need to monitor the VPSes you create for a while and up any values as needed... until you get to a place where you don't have any failures. This might be problematic on systems that have periodic spikes, but since resources are dynamic, resources can be increased during a spike and reduce after it is over. Currently it is a very manual process.

The basic and light samples are very lean on purpose. It makes sense to start light and increase resources as needed... rather than starting with arbitrarily high values and expecting users to scale back resources. How much memory should we assume a machine has and how fast of a CPU? vzsplite will take a look and generate a config sample specifically for your hardware.

The good thing is that OpenVZ really excels in DYNAMIC resource management... and you can tweak all of the values without having to restart your VPS. None of the other virtualization products out there offer much in the way of resource management much less dynamic resource management.

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [locutius](#) on Sat, 18 Aug 2007 19:22:12 GMT

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

---

thanks Scott. what i am thinking is the memory config sent me so deep into the config that it is not a bad thing for my learning.

i split the server 24 ways, then manually edited these lines in the config to set the VE limits at 256MB:

```
VMGUARPAGES="65536:2147483647"  
OOMGUARPAGES="65536:2147483647"  
PRIVVMPAGES="65536:75536"  
SHMPAGES="6553:6553"
```

mainly all guess work from studying the patterns in the numbers

this is the whole config when done for 256MB

```
# Primary parameters
```

```
NUMPROC="333:333"  
AVNUMPROC="93:93"  
NUMTCP SOCK="333:333"  
NUMOTHERSOCK="333:333"  
VMGUARPAGES="65536:2147483647"
```

# Secondary parameters

```
KMEMSIZE="7628151:8390966"  
TCPSNDBUF="1178749:2542717"  
TCPRCVBUF="1178749:2542717"  
OTHERSOCKBUF="589374:1953342"  
DGRAMRCVBUF="589374:589374"  
OOMGUARPAGES="65536:2147483647"  
PRIVVMPAGES="65536:75536"
```

# Auxiliary parameters

```
LOCKEDPAGES="372:372"  
SHMPAGES="6553:6553"  
PHYSPAGES="0:2147483647"  
NUMFILE="2976:2976"  
NUMFLOCK="297:326"  
NUMPTY="33:33"  
NUMSIGINFO="1024:1024"  
DCACHESIZE="1664248:1714176"  
NUMIPTENT="83:83"  
DISKSPACE="6291456:6291456"  
DISKINODES="157911:173703"  
CPUUNITS="7173"
```

vzcfgvalidate validates the config ok

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [dowdle](#) on Sat, 18 Aug 2007 20:43:55 GMT

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

---

Just so you know, I believe they have made many of the parameters more friendly now and you can use more human readable notations... so for example:

```
--diskspace 10G:10G
```

or

```
--vmguarpages=nnM:nG
```

vzsplit doesn't use the simpler notation though.

I must admit I think all of my configs have the older style notation... so it wouldn't be a bad idea to test that out before converting all of your configs.

---

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [locutius](#) on Sat, 18 Aug 2007 21:09:19 GMT

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

---

good to know Scott

the system is not unfriendly when you know it. but you must RTFM (pdf) to the last page. and i finally got there with some hand-holding from yourself, thanks

the wiki gives the impression of a fast start up which you can then run with and tune as you go. but that is not the reality because of the default configs which are conservative rather than generous. today i am editing my own templates and making my own configs but only by necessity, and if must be necessary, then the wiki must drive the n00b to it as fast as possible

the single blocker i hit was deep in the HN network config that could not have been solved from the documentation, the solution required knowledge from experience. i maintain this is a hole in the wiki because it is intuitive to map IPs to a HN NIC

ON TOPIC:

is this sane?

```
SHMPAGES="6553:6553"
```

i guessed to reduce the PRIVVMPAGES by a factor 10

---

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [dowdle](#) on Sun, 19 Aug 2007 01:27:53 GMT

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

---

locutius wrote on Sat, 18 August 2007 15:09

the wiki gives the impression of a fast start up which you can then run with and tune as you go. but that is not the reality because of the default configs which are conservative rather than generous.

As stated previously, for a default we want conservative rather than generous. That is a judgement call. If it was the other way around there would be a group of people asking why. There isn't really a way to choose and make everyone happy.

Quote:today i am editing my own templates and making my own configs but only by necessity

Good. Using vzsplrit? It is pretty darn handy... but then again... you STILL have to monitor your UBCs and adjust for failures until it is tuned. Again, this is the nature of the beast. All of the other virtualization products don't give you much choice... you absolutely must dedicate to each virtual machine as many resources (especially with regards to memory) as you would a physical machine... and that means you can only fit a handful of virtual machines per physical machine. With OpenVZ you can easily fit dozens and even hundreds if you machines are only a couple of lighter weight apps. And again, OpenVZ let's you manage all of the resources dynamically... so you can adjust anything on the fly. With that flexibility does come complexity.

Quote:the single blocker i hit was deep in the HN network config that could not have been solved from the documentation, the solution required knowledge from experience. i maintain this is a hole in the wiki because it is intuitive to map IPs to a HN NIC

Let me ask you this... if you had been installing VMware or Xen would you have assumed that you need to pre-configure the host machine for all of the network interfaces you plan to use in virtual machines? I'm guessing not.

Quote:is this sane?

```
SHMPAGES="6553:6553"
```

i guessed to reduce the PRIVVMPAGES by a factor 10

It might be sane if your virtual machine will work with those values. I can only tell you what the VPSes I run need. I run a few fairly heavy VPSes, a few pretty light ones... and a couple of absolute pigs. The configs between those machines vary greatly.

The pig VPSes are running Zimbra (only one Zimbra VPS per physical host though). Zimbra's system requirements basically say that Zimbra should be installed on a standalone server as the only application.

With OpenVZ I've been able to run Zimbra in a VPS along with 6 other VPSes on the hardware node... and while it did take me about a week to get it tweaked, I haven't had to change any of the resource settings since... and the two boxes I have running Zimbra have been doing so for about a year and a half.

If your VPSes are all a like, then vzsplrit will do almost all of your work for you.

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap

Posted by [locutius](#) on Mon, 20 Aug 2007 18:10:06 GMT

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

---

nice reply

i investigated VMware or Xen 18 months ago when this project was started and decided then i

could not operate VPS under the conditions of their use

the VPS are intended to run MySQL intensive high-traffic web applications (hence the multi dual-core cpu HN). security is also an issue with these sites and the HN has to run antidos parsing and high volume ipfiltering (2.6million+ ips) to protect the VPS services

i have 2 x VPS now operational on the HN but the admins have not yet cranked up the traffic, so we will have to wait and see what the real limits are going to be and what needs to be done to the config

vzsplit is the definitive application for anyone setting up VZ. the first thing to do on your box is run vzsplit to determine the true HN limits. then edit the template config and adjust the settings to your expectations for the VPS (set RAM and Disk). of course this is only good when all VPS are the same, but it is a better place to start than the default config because the default config does not reflect your true HN capabilities

vzsplit is also useful because it educates how VZ is treating the i/o level of the o/s. some settings are more important than others, while there are upper/lower limits to others, and many of the settings are directly interrelated. i discovered this running vzsplit for 8, 16, 24, 32, 40, 48, 54 (61 is the limit) ways then comparatively reading the configuration output. i recommend this to everyone to get a grasp of how VZ is working and how to configure your VPS

---

Subject: Re: configuration needed for P4 2GHz 1GB RAM and 2GB swap  
Posted by [mvdb](#) on Mon, 20 Aug 2007 22:14:18 GMT

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

Would you mind sharing your config file for zimbra ? (still learning openvz here and first thing on the list is moving zimbra from vmware to openvz)

---

Subject: Configuration for Zimbra  
Posted by [dowdle](#) on Tue, 21 Aug 2007 02:11:11 GMT

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

There is a Zimbra howto in the OpenVZ wiki that you might want to check out. If memory serves me, it doesn't say too much though.

Basically I installed CentOS 4 into a VPS. Removed apache and turned off sendmail (chkconfig sendmail off).

I started out with about 8x the values of the basic config. The install went pretty well... and I just monitored the machine tuning up the values whenever I saw a failure.

After a while I learned of vzsplit and used it to chop the machine I was using into 8 machines...

and oddly enough, the values provided by vzsplite were well over the values I had set by hand... but in an effort to improve performance, I upped them again. There were only a few parameters that I had set that were larger than what vzsplite provided... but the values below should work... but your mileage may vary.

```
Quote:AVNUMPROC="320:320"  
NUMPROC="1000:1000"  
NUMTCPSOCK="1000:1000"  
NUMOTHERSOCK="1000:1000"  
VMGUARPAGES="310487:2147483647"  
KMEMSIZE="22820454:25102499"  
TCPSNDBUF="3510818:7606818"  
TCPRCVBUF="3510818:7606818"  
OTHERSOCKBUF="1755409:5851409"  
DGRAMRCVBUF="1755409:1755409"  
OOMGUARPAGES="310487:2147483647"  
LOCKEDPAGES="1114:1114"  
SHMPAGES="31048:31048"  
PRIVVMPAGES="1572864:1714400"  
NUMFILE="8896:8896"  
NUMFLOCK="889:977"  
NUMPTY="100:100"  
NUMSIGINFO="1024:1024"  
DCACHESIZE="4974850:5124096"  
PHYSPAGES="0:2147483647"  
NUMIPTENT="200:200"
```

You are probably better off using the method I used rather than the specific values I have set above... as your hardware is going to be different from mine. The critical parameters are memory related... and zimbra usually has about 320 processes running (most of which are Java threads). Everything else is really overtuned... but in a way that doesn't hurt my other VPSes... since I split it up evenly with vzsplite.

---

Subject: Re: Configuration for Zimbra  
Posted by [locutius](#) on Thu, 23 Aug 2007 20:04:08 GMT  
[View Forum Message](#) <> [Reply to Message](#)

i discover a strange configuration outcome. please can you check

i have a 256MB RAM VPS giving failcnt on privvmpages

this was the memory settings:

```
VMGUARPAGES="65536:2147483647"  
OOMGUARPAGES="65536:2147483647"  
PRIVVMPAGES="65536:75536"  
SHMPAGES="6553:6553"
```



the solution according to the manual is to increase the RAM available for allocation ie. increase PRIVVMPAGES. VMGUARPAGES will remain at the 256MB service level and OOMGUARPAGES will account the actual RAM usage

```
vzctl set 102 --privvmpages 384m:512m --save
```

and now the VPS reports it has 384MB RAM available and not 256MB. i have sent an email to the VPS customer telling them to ignore the 384MB RAM and keep in mind that the extra is only for process allocation and that their real RAM is 256MB!

before i further convince my customers that i am crazy, plz confirm i am so, thanks

---