
Subject: Swap Usage

Posted by [hello-world](#) on Wed, 21 Jun 2006 14:00:37 GMT

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

Hi Folks,

i apologize in advance, if this question is not directly relevant to the forum. Thanks.

I have rented an openvz system from a third party, and i am impressed by its performance. Anyway, my problem is this: My vps doesn't have any swap. It shows a total memory usage of around 412MB, and as u know, it will get filled up, and i routinely get a 'memory exhausted' error. I contacted my provider, but they say that vps does not have swap, and we are supposed to 'LIMIT' our memory usage. Now i can only laugh at such an idea. I told them that swap is a concept basic to computing, and u can never lock you computer to that small a memory, and if the memory usage goes up, the system should slwo down--owing to swap usage--and not crash like it is happening now. But they are still adamant.

Anywa, i would like to know the basic swapping principle used in openvz? Can u have per vps swap, or is it a global swap? How do i enable swap?

Thanks a lot in advance.

Subject: Re: Swap Usage

Posted by [kir](#) on Thu, 22 Jun 2006 09:19:06 GMT

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

0. This is actually your provider which should help you with that. Probably you need to buy more advanced hosting plan, with raised resource limits.

1. Swap is global for all the VEs. It is used the same way as in ordinary Linux system -- if there are not enough memory, cache gets shrinked, then some pages are swapped out.

2. OpenVZ has a system of per-VE limits and guarantees called User Beancounters. You can check the current limits and usage values in `/proc/user_beancounters` file in your VE. For more info, see OpenVZ user's guide (from <http://openvz.org/documentation/guides/>) and http://wiki.openvz.org/Resource_management

Subject: Re: Swap Usage

Posted by [hello-world](#) on Thu, 22 Jun 2006 09:56:13 GMT

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

Hey, thanks for the reply.

Could you please tell me how i can find out how much is swap, and how much is RAM? i already have this script to print out the memory usage, but i do not not the actual contribution of Swap, and RAM.

```
-----  
beans=`cat /proc/user_beancounters | grep priv`  
max=`echo $beans | awk '{ print $4;}'`  
use=`echo $beans | awk '{ print $2;}'`  
let "per=$use*100/$max"  
let "mb=$use/256"  
echo "privvmpages usage: $mb MB ($per%)"  
-----
```

Does the memory usage as displayed by the above program include both RAM AND Swap?

Thanks a lot

Subject: Re: Swap Usage
Posted by [hello-world](#) on Thu, 22 Jun 2006 10:01:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

Also if I get a 512MB memory, doesn't it mean 512MB Ram? And shouldn't i be actually be able use to more memory, since i will have swap too. Of course, the machine will slow down owing to swap usage, but it shouldn't crash right?

Thanks

Subject: Re: Swap Usage
Posted by [kir](#) on Thu, 22 Jun 2006 10:13:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

It's not the machine which crashes, it's your application.

All the other questions should really be directed to your provider -- it's their job not ours, hope you understand that.

Subject: Re: Swap Usage
Posted by [hello-world](#) on Thu, 22 Jun 2006 13:10:36 GMT

kir wrote on Thu, 22 June 2006 06:13: It's not the machine which crashes, it's your application.

That's not entirely correct. It is indeed the machine that crashes. It is killing off my application.

Quote:

All the other questions should really be directed to your provider -- it's their job not ours, hope you understand that.

I contacted my provider, and spoke to them, but they said that per vps swap is not possible, and that swap is global.

My question is: If so, how can we ensure quality of service. A provider can use large amount of swap and sell it as memory. If you don't distinguish between RAM and Swap, how can a company make sure that the Service Levels Agreements are met?

Anyway, thanks a lot for your time.

Subject: Re: Swap Usage
Posted by [kir](#) on Thu, 22 Jun 2006 13:26:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote: That's not entirely correct. It is indeed the machine that crashes. It is killing off my application.

Most probably your application is not killed, but just crashes because it can not cope well with the out-of-memory situation. In some cases, though, it's indeed the kernel which kills your application in order for your VE to stay within the boundaries specified by User Beancounters limits.

Quote: I contacted my provider, and spoke to them, but they said that per vps swap is not possible, and that swap is global.

That's correct.

Quote: My question is: If so, how can we ensure quality of service. A provider can use large amount of swap and sell it as memory. If you don't distinguish between RAM and Swap, how can a company make sure that the Service Levels Agreements are met?

Quality of service is implemented via User Beancounters. A physical box is divided into multiple partitions (VEs), that are granted a subset of its resources. There are some explicit guarantees (see *guar* parameters in /proc/user_beancounters) and some implicit guarantees (if the physical

resources are not oversold you'll be able to use resources up to the limits specified in `/proc/user_beancounters`).

What your provider can do is to use `vzmemcheck` and `vzcpucheck` utilities to check if they are overselling or not, and `vzcfgvalidate` utility to check if a set of User Beancounters for a VE is consistent.

What you can do in this situation is to check `/proc/user_beancounters` output, and see if you have some non-zero values in `failcnt` column (as described in wiki: Resource management). If there are some, then you require more than the current hosting plan offers you -- and in that case you either tune your apps to ask for less resources, or ask your provider to give more resources to you (probably by switching to a better hosting plan, or tuning the individual resources).

So, at the end, this is a question of money, not the technology. Believe me, the technology (OpenVZ and its User Beancounters) is doing its best both for you as a VE owner, and for other VEs on the same box (by not letting you to abuse the physical server resources).

Hope that helps.

Subject: Re: Swap Usage

Posted by [hello-world](#) on Thu, 22 Jun 2006 15:28:44 GMT

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

Quote:

Most probably your application is not killed, but just crashes because it can not cope well with the out-of-memory situation. In some cases, though, it's indeed the kernel which kills your application in order for your VE to stay within the boundaries specified by User Beancounters limits.

It is trying to allocate memory, it fails, and thus `malloc` returns zero, and it crashes. `Ps` is the program the segfaulted due to this.

Quote:

physical box is divided into multiple partitions (VEs), that are granted a subset of its resources. There are some explicit guarantees (see `*guar*` parameters in `/proc/user_beancounters`) and some implicit guarantees (if the physical resources are not oversold you'll be able to use resources up to the limits specified in `/proc/user_beancounters`).

Hope that helps.

Thanks. My point is this: As far as monetary considerations go, RAM and swap r two entirely different things. So how can a VPS *USER* make sure that he gets a particular amount of RAM and certain fixed amount of swap.

Anyway, if you don't have proper differentiation between RAM and swap, i think that is a small

drawback, especially if you are selling VPSes to external vendors. What we need is actually a per vps swap, and both the RAM and the swap usage should be separately controllable.

Again thanks for your response.

Subject: Re: Swap Usage
Posted by [kir](#) on Thu, 22 Jun 2006 15:37:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote: It is trying to allocate memory, it fails, and thus malloc returns zero, and it crashes. Ps is the program the segfaulted due to this.

The program written right should handle this situation somehow.

At the very least, it can write something like "Failed to allocate %d bytes of memory, aborting". If it is segfaulting -- the program has a bug.

Subject: Re: Swap Usage
Posted by [hello-world](#) on Thu, 22 Jun 2006 15:44:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote:
The program written right should handle this situation somehow.

It is happening in the procs library, and actually it is irrelevant. The problem is that there are services that will not function without the resources.

When memory usage goes up, the machine should slow down, due to the use of swap, and not crash. I personally don't like this model at all. Each vps should be left to decide how much swap it wants.

Anyway, i would like to know if you will implement per vps swapping in future versions. Also, i would like to know if virtuozzo supports this particular feature.

Thanks again for your patience.

Subject: Re: Swap Usage
Posted by [kir](#) on Thu, 22 Jun 2006 16:11:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote:When memory usage goes up, the machine should slow down, due to the use of swap, and not crash.

1. Again, the machine is not crashing. It's your application that is crashing, because it *probably* hits the memory limit specified for your VE (I am not sure since you never shown or told me what's in your /proc/user_beancounters).

2. Note that this is not a machine but a VE. The limits of memory which you can is is specified in User Beancounters for your VE. It can not eat all RAM and swap.

Quote:I personally don't like this model at all. Each vps should be left to decide how much swap it wants.

Each VE is limited to a certain amount of memory. Some of that memory could be swapped out, because of RAM shortage. RAM shortage can happen because of yours or other VEs using RAM.

The fact that your VE memory is not swapp is just because the physical server is not loaded that high to use swap (and you should help your provider for such a good service). If physical server owner will put more VEs on that very box, it will start to swap. If this is what you need, ask them to put more VEs on your box -- and you will see some memory pages will be swapped out. Still, it will not increase the amount of memory available for your VE, because the total amount of RAM+swap is limited per VE by your hoster using User Beancounters.

Said that, I can not really do anything about it.

Quote:Anyway, i would like to know if you will implement per vps swapping in future versions. Also, i would like to know if virtuoizzo supports this particular feature.

No. No.

Subject: Re: Swap Usage

Posted by [frankfischer](#) on Sun, 25 Jun 2006 19:50:18 GMT

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

I think the main problem here is that you misunderstand the concept of openvz. IMHO openVz is not about giving you a virtual Server in the sense of a mainframe's hardware partition or the like Xen is doing. The concept of openVz is more the like of a very very (do I said very?) sophisticated changeroot. It's about having all critical resources of a servers hardware controlled by the UBC and beeing able to allocate these resources to virtual environments. This concept is best compared to Solaris zones. Maybe you find some better explanation of this concept on the solaris pages at www.sun.com.

This concept has advantages and disadvantages - but this is an other topic.

Regards,
Frank