Subject: Memory & SWAP
Posted by devonblzx on Mon, 22 Jan 2007 21:08:50 GMT
View Forum Message <> Reply to Message

All right,

After reading a few discussions on these forums, I think I'm getting a partial hang of it...but I still want some clarification.

free -m, when --meminfo privvmpages:1 is set returns the "physpages", right? So why is there such a gap between my physpages and whats actually being used in privvmpages?

vzmemcheck -vA returns this:

Output values in Mbytes

veid	LowMem util	LowMem commit	RAM util	MemSwap util	MemSwap commit	Alloc util	Alloc commit	Alloc limit
1	6.15	13.61	90.50	90.50	37.61	250.94	1037.61	1037.61
Summary:	6.15	13.61	90.50	90.50	37.61	250.94	1037.61	1037.61
	1579.00	1579.00	3949.00	10093.00	10093.00	10093.00	10093.00	10093.00

I realize the kernel memory isn't included, but free -m returns 84MB being used and in privvmpages, it says I'm using ~250MB.

Also where does SWAP come into play? I have 4GB of RAM and 6GB swap partition but it seems that the SWAP is never used even when a VPS is maxed out on RAM usage. I started getting fork errors but I still had 0/6GB Swap usage.

Is there a setting I should set for how much SWAP a VPS can use? Or do I set privvmpages for SWAP + RAM and then maybe vmguarpages for just RAM?

Any help would be greatly appreciated!

Thanks,
Devon

[EDIT]
Maybe someone can give me some exact values...

I want a VPS with 1024MB RAM and 1024MB SWAP available, right now I have privvmpages and vmguarpages to 262144.

Subject: Re: Memory & SWAP
Posted by rickb on Tue, 23 Jan 2007 01:15:05 GMT

View Forum Message <> Reply to Message

privvm and phys will never be equal or even close in most cases. As the manual/wiki states, privvm is memory allocation whereas phys is memory usage. Most applications allocate large gobs of memory and actually use a fraction of it. you can look at privvm as the potential memory usage of the running applications and phys as the actual memory usage.

Subject: Re: Memory & SWAP
Posted by devonblzx on Tue, 23 Jan 2007 02:15:39 GMT
View Forum Message <> Reply to Message

All right I understand what your saying. Why does it give me fork messages when I run out of privvmpages?

And do you know anything about my SWAP problem (above)?

Subject: Re: Memory & SWAP
Posted by xemul on Tue, 23 Jan 2007 09:17:52 GMT
View Forum Message <> Reply to Message

When you fork your new task inherits all memory from its parent. Thus if your task has allocated 128M of private anonymous mappings and forks new system can potentially eat 256M of memory - if this exceeds privvmpages we don't allow forking.

SWAP isn't limited now - it's just accounted into oomguarpages. When system becomes low of memory it starts to oom kill tasks taking oomguarpages into account.

When VE eats all provvmpages no swapout starts - only when the whole system runs out of mem.

Also note, that freeing memory consists of 3 stages:
1. shrink kernel caches
2. sysncing file caches to disk
3. swapping anonymous pages out

so if kernel caches were shrinked no user pages go away.

Subject: Re: Memory & SWAP
Posted by devonblzx on Wed, 24 Jan 2007 03:00:35 GMT
View Forum Message <> Reply to Message

So there is no way to set how much SWAP a VPS can use without affecting the amount of allocated RAM with it?

Lets say this scenario:

I have 4 clients on a 2GB RAM box with 2GB SWAP, can I give each of my clients privvmpages 262144 vmguarpages 131072 oomguarpages 262144 so they will have 512MB guaranteed RAM but have the ability to use the extra 512MB in SWAP? Is there another way to do this if that doesn't work?

Subject: Re: Memory & SWAP
Posted by rickb on Wed, 24 Jan 2007 04:14:11 GMT
View Forum Message <> Reply to Message

As xemul said.. the VE has no concept of the difference between ram and swap. Its just "memory" at the VE level. So, your HN can have 64MB of ram and 64GB of swap. Or, your HN can have 64GB of ram and 64MB of swap. In the VE, the privvm, oomguar, vmguar are uneffected by this. Only the performance changes between the two scenarios. In either case, you can guarantee and limit each VE to an arbitrary amount of memory. This drastically differs from the conventional mem management of vmware/xen (isoloation). VZ is a shared kernel model and no resources are "hard dedicated" or "hard assigned".

-not possible to give/assign/guarantee a VE XX swap and YY physical memory
-not possible to restrict a VE from using physical memory and use swap instead

Think of the VZ virtual servers like resource containers, not individual linux boxes- VZ is a shared kernel approach to virtualization.

Hope this clears it up for you.

Rick Blundell

Subject: Re: Memory & SWAP
Posted by devonblzx on Wed, 24 Jan 2007 05:03:27 GMT
View Forum Message <> Reply to Message

Yes thank you.

So when I give 1GB burstable RAM that will give them the ability to use that but they won't necessarily be using the SWAP but when the physical RAM does runs out the new memory left over memory will be SWAP, right? So I can give 512MB to 1GB burst even though I don't have the physical RAM to support it?

Subject: Re: Memory & SWAP

Posted by rickb on Wed, 24 Jan 2007 14:27:19 GMT
View Forum Message <> Reply to Message

This is your same question reworded.

Think of a normal linux kernel with two instances of Apache. Now, can you tell the kernel to allow each to allocate XX MB of physical mem and YY MB of swap?- no you cannot. Now, think of each of those apaches being a VZ VE. Bottom line, in this simple scenario, the same memory manager is used on the HN vs. a conventional kernel. The only differences is that the kernel can reap back memory to enforce guarantees (vmguar+oomguar) or apply an upper alloc limit (privvm).

Quote:So when I give 1GB burstable RAM that will give them the ability to use that but they won't necessarily be using the SWAP but when the physical RAM does runs out the new memory left over memory will be SWAP, right?

I already answered this with:
Quote:
not possible to give/assign/guarantee a VE XX swap and YY physical memory

---
Quote:
So I can give 512MB to 1GB burst even though I don't have the physical RAM to support it?
I already answered this with:
Quote:
you can guarantee and limit each VE to an arbitrary amount of memory.

Subject: Re: Memory & SWAP
Posted by pzYsTorM on Fri, 21 Aug 2009 09:12:53 GMT
View Forum Message <> Reply to Message

Sorry, but I have some questions to that.

What happens if I assign (f.ex.) /dev/sdb12 with 4GB swap space to a VE (vzctl --devnodes) and type swapon /dev/sdb12 inside the VE?
It is useless? Does it work like I hope? Is there a workaround?

Or would it be better to give 20GB swap space (5 VEs, each one should have 4GB swap) to the host system and no swap to the VE?

Thanks.

Subject: Re: Memory & SWAP
Posted by poige on Sun, 18 Oct 2009 13:10:17 GMT
View Forum Message <> Reply to Message

Quote:Think of a normal linux kernel with two instances of Apache. Now, can you tell the kernel to allow each to allocate XX MB of physical mem and YY MB of swap?- no you cannot. Now, think of each of those apaches being a VZ VE.

Ok, but OpenVZ is kernel level software and comparing it to Apache's looking wrong. The real problem is that OpenVZ doesn't support swap space partitioning (i. e., one can not say "ok, this instance can have as much as N MiBs/GiBs of swap"), though it could be really cool. And Apache isn't around at all.
