

---

Subject: Re: Calculating kmemsize

Posted by [dragos2](#) on Thu, 03 Sep 2009 13:34:01 GMT

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

---

damswf wrote on Wed, 02 September 2009 21:17: I agree, I just installed a web server in an openvz container, and I frequently have this problem. Either I can't do ps anymore, or I don't have a shell when I ssh, etc... instead I got "fork: Cannot allocate memory".

This is a pain for me since I wanted openvz to be more protected against hardware problems, and I have this headache instead.

Every time I grow up the kmemsize barrier limit, I see the maxheld increasing until being greater to the barrier limit then I got some failcnt and errors.

This server is not at all under load now, but once it will be, it will be because of a lot of customers hitting the site (I cannot reproduce the same thing myself to test), and I cannot take the risk to have databases crashing nor losing data because of being unable to allocate kernel memory on a 8GB 64bits system (VC have 4GB).

And I have no clue what the values should be! I have  $\text{kmemsize} \geq 40\text{kB} * \text{avnumproc} + \text{dcache size}$  but still the problem.

This is a serious issue for production servers.

So what do I do? I explain to my boss I just virtualized our servers and we have more problems although I explained him a few times ago that it will work better?

You can try other virtualization types like Citrix FreeXenServer or even the opensource Xen (I don't have experience using VMware ESX).

But using OpenVZ as a webserver in full load I think the main things to consider are:

- 1) number of processes and files opened  
openvz parameter for processes number 66535 is a huge number  
also the number of open file descriptors should be at least  
twice the threads your webserver spawns
- 2) tcp/ip buffers size, number of sockets. Pay attention to these both on HN and as new container parameters. Be sure  
to reuse them fast
- 3) cpu power if you make use of cgi, fcgi etc.

The upper values should be configured both on HN and in containers too.

If you set them only on one side, like only in HN or only in container you won't gain anything.

This was from my little experience.

I hope it helps.

---