Subject: vmguarpages != oomguarpages ?
Posted by divB on Sun, 06 Sep 2009 20:29:31 GMT
View Forum Message <> Reply to Message

Hi,

Let us assume a (hypothetical) HN with 1GB RAM and 10 VEs. For the sake of simplicity we do not care about kernel mem, buffers, hardware node...so we can give every hardware node 100 MB. And I want to guarantee each container 100 MB.

If I want to make sure that every container is guaranteed 100 MB of RAM I have to set vmguarpages to 100 MB, haven't I?!

Because let's assume that all memory malloced is actually used we would face a problem and we could not guarantee 100 MB of RAM?!

But on the other hand, oomguarpages tells the guaranteed pages in an out of memory contition or in a different way: The guarantee of actual used pages (i.e. malloced and used).

So I must conclude that I have to always set oomguarpages = vmguarpages in order to ensure each VE is guaranteed 100 MB?

Is this correct or do I get something wrong? Is there a good reason to set oomguarpages to 100 MB and vmguarpages to 120 MB in my hypothetical example?

When hosting providers talk about memory guarantee: Do they mean the guaranteed amount of malloc-able memory (i.e. vmguarpages) or do they mean the guaranteed amount of useable memory (i.e. oomguarpages)?

Regards,
divB

Subject: Re: vmguarpages != oomguarpages ?
Posted by maratrus on Fri, 11 Sep 2009 15:45:36 GMT
View Forum Message <> Reply to Message

Hello Niki,

sorry but I really couldn't catch you.
If you don't mind I'd rather not pass through each of your logical steps.
I just want to pay my attention to
Quote:
Is there a good reason to set oomguarpages to 100 MB and vmguarpages to 120 MB in my hypothetical example?

Please keep in mind the following statement which is taken from\
http://wiki.openvz.org/Privvmpages#oomguarpages
Quote:
The total out-of-memory guarantees given to the containers should not exceed the physical capacity of the computer, as discussed in UBC systemwide configuration#Memory and swap space. If guarantees are given for more than the system has, in out-of-memory situations applications in containers with guaranteed level of service and system daemons may be killed.

http://wiki.openvz.org/UBC_systemwide_configuration#Memory_a nd_swap_space
Quote:
Commitment levels more than 1 means that the containers are guaranteed more memory than the system has. Such overcommitment is strongly not recommended, because in that case if all the memory is consumed, random applications, including the ones belonging to the host system, may be killed and the system may become inaccessible by ssh(1) and lose other important functionality.

Use the formula (or vzcfgvalidate utility) to make sure that  Commitment level doesn't exceed 1.