
Subject: physpages vs oomguarpages

Posted by [nethubonline](#) on Tue, 08 May 2012 04:09:32 GMT

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

Hi all,

In my limited knowledge, for vSwap-enabled kernels, physpages can limit total size of physical RAM used by a container, for old kernels physpages is just for accounting purpose.

oomguarpages is to limit current amount of RAM + swap + kernel memory (kmemsize) + socket buffers used by container, if container exceeds this limit, its process may be killed.

Hence, for old kernels, if there is no memory being swapped, physpages will be equal or similar to oomguarpages.

However for vSwap-enabled kernels, oomguarpages seems are always less than physpages, sometimes oomguarpages is just 10% of physpages.....

Anyone can let me know why?

Secondly, while physpages reaches 1G, but oomguarpages is still 400MB, how can we limit the container 1G RAM, so that this container's process being killed while exceed 1G RAM?
