
Subject: vmprivpage and vm.overcommit_ratio
Posted by [jreb](#) on Tue, 31 May 2011 09:38:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

I want to know if there is a relation between limit in term of sum of vmprivpages (all vps) on a host, and available memory (RAM + SWAP).

In particular is kernel overcommit behavior on the node is impacting the behavior of OpenVz, that is to say, can I control OpenVz overcommit of vmpages, by modifying vm.overcommit_memory and vm.overcommit_ratio (as reported by syctl -a in the kernel node).

Default linux kernel is overcommit_memory=1 and overcommit_ratio=50. Imagine I set all vmprivpage limit above this ratio. Could a process be blocked by kernel limitation, before I reach an UBC parameter limit (that is before reaching a vmprivpage limit) ?
