Hi,

I have made this question to ServerFault, but no-one seems to be able to answer:

When the failcnt of tcpsndbuf increases, what does this mean? Does it mean the system had to go past the barrier, or past the limit? Or, maybe, that the system failed to provide enough buffers, either because it needed to go past the limit, or because it needed to go past the barrier but couldn't because other VMs were using too many resources?

I understand the difference between barrier and limit only for disk space, where you can specify a grace period for which the system can exceed the barrier but not the limit. But in resources like tcpsndbuf, which have no such thing as a grace period, what is the meaning of barrier vs. limit?

Why does the difference between barrier and limit in tcpsndbuf have to be at least 2.5KB times tcpnumsock? I could understand it if, e.g., tcpsndbuf should be at least 2.5KB times tcpnumsock (either the barrier or the limit), but why should I care about the difference between the barrier and the limit?

(You can reply at serverfault if you prefer, serverfault.com/questions/395759/)

Thanks!

Page 1 of 1 ---- Generated from OpenVZ Forum