
Subject: Caches on 32 kernel

Posted by [non7top](#) on Thu, 26 Jun 2014 20:50:15 GMT

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

I faced a strange behavior of the memory inside the VE. VE has 760Mb ram and 512 vswap. At some point the oom-killer kicked in killing some random processes, at that time there were 250Mb ram and 250 vswap used, rest(500Mb) was utilized for caches. I don't have bean counters recorded, but i believe pshyspages failcnt had non-zero value.

In this case I would expect the VM to first use remaining vswap, then free up some caches and only when those are all used oom-killer should have kicked in. It appears to me that caches are not freed or freed very slowly, so that if at some point caches+used reaches total memory kernel considers the oom situation instead of dropping some caches.

Here is the image with red arrow pointing to moment when oom-killer was invoked.

Here is the dmesg

[xttp://pastebin.com/HAcNwnsS](http://pastebin.com/HAcNwnsS)

Any input would be appreciated, could be that this behavior is expected and it's me who is getting ths wrong, or could be some sort of bug.

File Attachments

1) [image2.png](#), downloaded 1173 times
