

---

Subject: Re: TCP: time wait bucket table overflow - memory leak?

Posted by [nksupport](#) on Wed, 15 Jul 2009 14:36:47 GMT

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

---

maratrus wrote

But there is no such a killer in OpenVZ. Parallels Virtuozzo Containers contains subsystem which does what you have said but in OpenVZ you can adjust the behavior of the container via `user_beancounters` i.e. if barriers and limits are really huge there is no a restrictive force which prevents a containers from consuming a lot of resources.

Yes there is! A process that failed to allocate memory will obviously be killed, isn't that right? I am talking exactly about UBC in my post. I'm using UBC to limit the VEs. Like i said, the current UBC limits are draconian. I want all processes that try to overuse memory to be killed and OpenVZ is supposed to do it.

maratrus wrote

Could you please say what means "the entire node just dies"? Are you able to ping the node? Are you able to invoke commands?

What is the reason of LA being so big?

Does the node perform CPU-consuming operations? What is the CPU state when "the node is being died"?

Or there are a lot of input-output operations being invoked?

You know, it's hard to debug a server at LA of 500 I can still ping it, but it's just too slow. Random processes from VEs are in top, it behaves just like it would with unlimited UBC, i.e. instead of killing or throttling the runaway process it gives it more and more CPU. Eventually the node's ssh dies and this ends the show. Anyway - i attached my `/proc/user_beancounters` and performance stats, please have a look if i missed anything.

---