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

Posted by [maratrus](#) on Mon, 20 Jul 2009 06:23:18 GMT

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

Hi,

Quote:

I am now nearly sure it's a kernel bug - i can't imagine any other reason.

You can file a new bugreport whenever you want. Please don't hesitate doing it. I just want to clarify the situation because I don't understand clearly what your problem consists of.

The last "top" output doesn't contain anything frightening from my point of view.

A "huge" load average is just a consequence of a "great" number of process. A good idea is to obtain the status of these processes. You can do it with help of "ps" utility. Please, examine their states.

Quote:

I can't restart the failed VE, can't raise the kmemsize or anything.

Why do you want to increase kmemsize? The only failcounters I can see on provided vzstats output are those concerning with privvmpages not kememsize.

What process are running in a failed VE?

You have mentioned that ssh dies when this occurs. Don't you think that this problem relates to network? Do you have a direct access to that server to confirm that the node completely dies?

Do you have any specific settings, for example nfs inside VE?

Please the next time it occurs please gather alt-sysrq-* information from the problem node.

Alt-sysrq-

- 1) "m" - for memory info dump
- 2) "p" - for registers - several times, please, twice the number of CPUs
- 3) "a" - for scheduler stat - 3 times
- 4) "w" - another scheduler info - 3 times
- 5) "t" - for all processes calltraces. Warning - this is a resource consuming operation. At least - twice.

To gather it you probably have to install serial console

http://wiki.openvz.org/Remote_console_setup#Serial_console
