Subject: TCP: time wait bucket table overflow - memory leak? Posted by nksupport on Tue, 14 Jul 2009 08:45:35 GMT

View Forum Message <> Reply to Message

Hi guys.

We've set up a brand new openVZ server. It's hardly ever loaded in normal mode.

At times when one of the VEs gets overloaded we start getting "TCP: time wait bucket table overflow" and the kmemsize bean counter grows.

What i expect from the node in this case is to kill the failed process or even the entire VE.

Instead, the node's LA spikes up to hundreds and then the entire node just dies.

vzctl stop times out. vzctl stop --fast times out. kill -9 `init of the VE` fails.

This looks like a memory leak to me.

I tried the solution to the error message described at http://bugzilla.openvz.org/show_bug.cgi?id=460 - it did not work for me. Anyway, i doubt it's the root cause of our problem - looks like just one of the symptoms.

The main problem is that when a VE hits a counter (probably only kmemsize, but i'm not sure whether other limits trigger the same problem), the node dies itself instead of killing a VE. That's not what you'd expect from an encapsulated virtual server.

The node's normal production load is 0.05 to 0.40.

I've tried setting both loose and strict memory UBC limits - it didn't change anything. The current UBC limits are drakonian.

The server's 2xXeon L5410 (8 cores total) with 8G RAM running Centos 5.3 x64.

The kernel's 2.6.18-128.1.1.el5.028stab062.3 #1 SMP Sun May 10 18:54:51 MSD 2009 x86_64

rpm -qa | grep vz ovzkernel-2.6.18-128.1.1.el5.028stab062.3 vzyum-2.4.0-11 vzrpm43-python-4.3.3-7_nonptl.6 vzquota-3.0.12-1 ovzkernel-devel-2.6.18-128.1.1.el5.028stab062.3 vzctl-3.0.23-1 vzctl-lib-3.0.23-1 vzpkg-2.7.0-18 The VEs are different, centos and debian. I have confirmed the same behaviour on three different VEs running different OS releases: one of them hits the limit, node dies.

I've attached some debug output, hope someone can find a clue - so far i could not. I could really use a hand, thanks!

File Attachments

- 1) vzstats, downloaded 670 times
- 2) vzctl_strace, downloaded 642 times
- 3) sysctl, downloaded 637 times
- 4) vzctl_fast_strace, downloaded 663 times
- 5) dmesg, downloaded 660 times