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

Posted by [seanfulton](#) on Thu, 27 Aug 2009 13:46:43 GMT

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

I am not sure if this will help but here is my \$.02:

I have had something similar happen when we over-allocated the RAM on the HN. We had 6G in the machine and had four VEs that collectively had access to 18G. The HN can't stuff 18G of crap into a 6G bag so it keeps going until it dies. The solution for us was to run ve-split for four (in your case five) containers and start with that. Then adjust resources *carefully*.

Another issue we have had (and are still having) has to do with I/O activity on the HN during nightly backups. We tar each VE nightly (used to use vldump, but it uses tar as well, so there was no functional difference in this respect). The tar used to be done to an NFS mount, now we use ssh/dd to pipe it to the backup server.

Anyway, tar -czv * knocks the machine to its knees. The VE it is backing up will effectively lock. On a busy mail server (inside a VE), this causes all sorts of damage to ongoing mail connections, so bad that we now have to restart the VE after every nightly backup.

In my experience, OpenVZ is very good about enforcing limits of activities within a VE, but can not seem to contain activities on the HN itself. It has been suggested that this is a kernel bug in the 2.6.18 kernel, but no solutions have been found or suggested.

One solution recommended was to change the scheduler on the drive from cpq to deadline like this:

For a drive that is /dev/sda, use:

```
# cat /sys/block/sda/queue/scheduler
```

```
# cat /sys/block/sda/queue/scheduler
```

This still caused the load to grow, but the VEs no longer completely lock up during the process, which is a plus.

I am still looking for a more lasting solution to this last issue, but that's not your problem.

Hopefully this experience will point you in a useful direction. Good luck. I know it's frustrating.

sean