
Subject: HW Node Lockup Issue

Posted by [shushry](#) on Mon, 26 Feb 2007 21:30:52 GMT

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

Have a number of hardware nodes, all configured identically (hardware/OS). Running into a strange problem that has been occurring quite frequently, where the hw node will suddenly stop accepting any kind of tcp connection (ssh, smtp, etc). Existing sessions just freeze up, and even my direct connection to the serial port (system configured to allow com port console login), freezes up.

The freeze up will happen for anywhere from 30 seconds to 15 minutes (or as long as I can tolerate before power-bouncing the box).

This is a high-bandwidth system, but not very heavy in memory or CPU (other than what supports the traffic).

One question - there is alot of inter-vps traffic occurring; for example http to proxy servers and back...upwards of 10-50 mbs. Could the tcp communication amongst virtual servers along the venet0 interface be somehow "overloading" and freezing up the box?

Output of user_beancounters attached.

In general, I have all 63,000+ ports open for connectivity on the system. Result of "cat /proc/net/sockstat" is generally never more than 7000 - 10000 sockets in use.

HW node system specs:

64bit Centos 4.4

4GB Ram (fully utilized per allocation, but no swap dipping)

2.6.18-ovz028test015.1-smp

vzctl version 3.0.13

Any other details I can provide?

Anyone that can solve this gets lunch on me.

File Attachments

1) [beancouter_output.txt](#), downloaded 299 times

Subject: Re: HW Node Lockup Issue

Posted by [kir](#) on Mon, 26 Feb 2007 22:05:17 GMT

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

Sounds like bug #460. It is already fixed in GIT and will be released in 028test018 (or later) kernel.

For now, these two settings should solve the problem:
echo 16536 > /proc/sys/net/ipv4/tcp_max_tw_buckets_ub
echo 384 > /proc/sys/net/ipv4/tcp_max_tw_kmem_fraction

You can either set it from the command line, or via /etc/sysctl.conf. In the latter case, add the following:

```
net.ipv4.tcp_max_tw_buckets_ub=16536  
net.ipv4.tcp_max_tw_kmem_fraction=384
```

and then run sysctl -p

Subject: Re: HW Node Lockup Issue
Posted by [shushry](#) on Mon, 26 Feb 2007 22:14:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks Kir -

I found reference to that bug in several other posts, and took liberty to make adjustments.

I've been running with:

```
echo 20000 > /proc/sys/net/ipv4/tcp_max_tw_buckets_ub  
echo 256 > /proc/sys/net/ipv4/tcp_max_tw_kmem_fraction
```

...and still experienced the issue. Do you think the 384 value will improve this? And would my 20,000 value for _ub be sufficient, or is the 16536 significant?

One other question - would there be some way to monitor the onset of this issue? I.e, some value in /proc that would indicate this problem is about to occur?

Subject: Re: HW Node Lockup Issue
Posted by [dev](#) on Tue, 27 Feb 2007 13:07:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

shushry wrote on Tue, 27 February 2007 00:30 Have a number of hardware nodes, all configured identically (hardware/OS). Running into a strange problem that has been occurring quite frequently, where the hw node will suddenly stop accepting any kind of tcp connection (ssh, smtp, etc). Existing sessions just freeze up, and even my direct connection to the serial port (system configured to allow com port console login), freezes up.

the freeze of the console means a lot... Thanks for such details.

Can you please collect the output of Alt-SysRq-P (press 5-8 times plz) and Alt-SysRq-T when it hangs?

shushry wrote on Tue, 27 February 2007 00:30The freeze up will happen for anywhere from 30 seconds to 15 minutes (or as long as I can tolerate before power-bouncing the box).

This is a high-bandwidth system, but not very heavy in memory or CPU (other than what supports the traffic).

One question - there is alot of inter-vps traffic occurring; for example http to proxy servers and back...upwards of 10-50 mbs. Could the tcp communication amongst virtual servers along the venet0 interface be somehow "overloading" and freezing up the box?

there should be no problem with lots of inter-VE traffic. it there is - it is a bug.

shushry wrote on Tue, 27 February 2007 00:30
Output of user_beancounters attached.

In general, I have all 63,000+ ports open for connectivity on the system. Result of "cat /proc/net/sockstat" is generally never more than 7000 - 10000 sockets in use.

HW node system specs:

64bit Centos 4.4
4GB Ram (fully utilized per allocation, but no swap dipping)
2.6.18-ovz028test015.1-smp
vzctl version 3.0.13

Any other details I can provide?

Anyone that can solve this gets lunch on me.

yes, see above the info which can help us to resolve.

at first I though it is TCP/IP problem, but the fact the your console freezes up makes me believe it is something else - e.g. long loop somewhere in the kernel or similar stuff.

BTW, is it a binary kernel from openvz.org or you build it yourself?

Subject: Re: HW Node Lockup Issue

Posted by [shushry](#) on Tue, 27 Feb 2007 16:45:38 GMT

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

We've reduced some of the load, and certainly the inter-ve heavy bandwidth and the problem hasn't occurred again; I'm trying to recreate the scenario in our lab where I'll try out your suggestions. If I can't reproduce it, I'll re-engage the load on the original server.

Also, these are stock binary kernel's from openvz.org.

More soon!

Thanks!
