

Re: netstat – negative number of queues free

Source: <http://unix.derkeiler.com/Newsgroups/comp.unix.sco.misc/2003-10/1188.html>

From: Mike Brown (*mike_at_tkg.ca*)

Date: 10/31/03

Date: Fri, 31 Oct 2003 15:35:02 GMT

"Stephen M. Dunn" wrote:

- >
- > *In article <3FA1C0C9.ACFEDEC@tkg.ca> Mike Brown <mike@tkg.ca> writes:*
- > *\$What network card and driver are you using?*
- >
- > *name=bcme0 vec=5 dma=- chip=BCM5702 mem=F5FE0000 phy=BCM5703 addr=00:0b:cd:4f:c0:72*
- >
- > *This is the integrated Broadcom gigabit Ethernet card. The*
- > *driver would either have been included in 5.0.7 or, more likely,*
- > *come from EFS 5.58a (I initially set up the server using something*
- > *older than 5.58a – 5.48a sticks in my mind for some reason – and*
- > *then upgraded to 5.58a, which was the latest at the time I*
- > *upgraded it).*
- >

The current driver (which may be the same binary) would come from "hpninstall bcme" after installing EFS5.60a.

- > *\$What happens to the server if you pull and reconnect the network*
- > *\$cable a few times under high network load. I would be interested*
- > *\$to see if the system panics and goes down in a few hours.*
- >
- > *Haven't tried it, and I'm not sure my client would be enthusiastic*
- > *about me trying to cause further disruption. Unless we've run*
- > *out of other ideas, in which case they'll go for anything because*
- > *we really need to stop this from happening ...*
- >

I am concerned about problems with the cards based on the n100c driver, particularly when two cards are configured for fail over. So far I have not seen the same issue with the bcme driver.

- > *\$Whats it the original problem you are seeing?*
- >
- > *BBx (Pro5) data files get corrupted. It happens occasionally,*
- > *sometimes a few times a week, sometimes not for a couple of weeks.*
- > *When a file gets corrupted, almost all of it can be read, and the*

comp.unix.sco.misc: Re: netstat – negative number of queues free

> method they've been using to recover is to read the file sequentially
> from the beginning, writing the data to a new file, until they hit
> an error; then they start reading sequentially from the end forwards,
> until they hit an error. They tell me this recovers almost all of the
> data, but of course it's time-consuming; many of their files are
> tens of megabytes, and quite a few are hundreds of megs. Even with
> a 642 array card with five 15k rpm Ultra320 drives, it takes a
> while to rebuild hundreds of megs of data ...
>
> The users access the data by telnetting to the server and running
> programs, so the actual data isn't going across the network. However,
> there is access to the data via ODBC; I'm trying to find out from
> the folks who wrote the applications what this entails (whether it's
> read-only or read/write, and whether it's something that might
> possibly be involved in causing corruption). Unfortunately, when
> the problem happens, their IT guy usually reboots the server,
> rebuilds the data, and lets the users back on before letting
> me know "It happened again" and that makes it hard to do any
> troubleshooting; I only got the netstat -m output because I
> happened to be there today when it happened.
>
> They started having this problem* with an older ProLiant running
> 5.0.5 so they bought a new box and had me install 5.0.7 on it and
> transfer their user accounts, config files, data, etc. Both this and
> the old server have been on two different UPSes – including one which
> powers all the rest of their servers, which do not have data
> corruption issues. All of the patch cables in the server room have
> been replaced; the hubs are soon to be replaced with brand new
> switches, and they are considering having an electrician test all of
> their cabling. They've added an air conditioner to the room (and
> it's powered off a separate electrical feed) because it did get
> kinda warm in there sometimes.
>
> *: and others as well – on the old server, sometimes any process
> that tried to access the /u filesystem would hang and become
> unkillable, and then when they rebooted the server they'd get
> data corruption; they also got nasty performance problems in
> which %sys would be very high when accessing that filesystem.
> The hangs and the %sys cleared up with the new server.
> --
> Stephen M. Dunn <stephen@stevedunn.ca>
> >>>-----> <http://www.stevedunn.ca/> <-----<<<
>

> Say hi to my cat -- <http://www.stevedunn.ca/photos/toby/>

Tough one, particularly if it happens once a week. I would recommend doing some packet sniffing just to see if anything suspicious is happening. It may be that we need to see up a loaner machine that can grab all of the network packets in a looping log, and have someone onsite freeze the capture when they see a problem. It is

comp.unix.sco.misc: Re: netstat – negative number of queues free

important to grab the corrupted file so we can dump out the bad area and see what happened.

The only time I have seen strange problems like this, which I assumed are fixed now, was when the database did not handle inode numbers larger than 64K. 99% of the files had inode numbers < 64K, but a few temp files would be allocated above that. When the program wrote out a block of info it truncated the inode number and overwrote a different file. Happened only very randomly when the work file area had lots of files.

Mike

--

Michael Brown
The Kingsway Group
Voice: 905 669 8101