

Re: TCP parameters and interpreting tcpdump output

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2006-11/msg01235.html>

- *From:* Dieter <frebsd@xxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Sat, 18 Nov 2006 23:42:31 +0000
-

Dan writes:

Dan> A shrinking window and no packet loss is an indication that the program
Dan> the socket is connected to isn't reading data fast enough. If you're
Dan> locally gzipping the output of a remote backup, for example, you'll see
Dan> this.

Just a tight loop reading the socket and writing to stdout, which is directed into a file on disk.

Dan> The completely duplicated data packets from the sender, even before any
Dan> perceived packet loss, are troubling. Either the sender decided to
Dan> resend that data on its own, or the packet was duplicated by a router
Dan> or switch in transit. Dumps of the same stream from both sender and
Dan> receiver would help, as would enabling rfc 1323 extensions on both
Dan> systems (which will put a timestamp value on each packet and enable
Dan> SACK. It's enabled by default on FreeBSD).

No router or switch, just a piece of wire.

net.inet.tcp.rfc1323: 1

Bill writes:

Bill> My guess would be that your process blocked on stdout.
Bill> You don't mention what you're doing with stdout from the program, are
Bill> you just letting it scroll on the terminal, or redirecting it to a file?

Just redirected to a file. FFS, soft updates, 7200 rpm SATA drive with the disk's write cache turned off. Input data rate is less than 20 M bits/sec. I can write to the disk at approx 6 M Bytes/sec sustained. (or 10x that with disk write cache turned on, but I don't like trashed filesystems after the machine goes down hard) The machine and the disk are plenty fast enough, AMD64, 2 GB main memory. CPU is 90-something percent idle.

Sometimes it works fine for extended periods, 30-40 minutes. Other times the src box reports thousands of network errors. So far I haven't figured

Re: TCP parameters and interpreting tcpdump output

out what the difference is between the working tests and the failing tests.
The crontab directory is empty, so it shouldn't be cron jobs.

As an experiment, try running the process and redirecting
stdout to /dev/null — if it doesn't exhibit the problem, then you
need to look at where you're actually storing the data and speed that
part up.

I've thought of trying /dev/null but haven't yet. It might provide
a clue.

I would expect that the filesystem should be buffering the write
from short term disk latency. Surely FreeBSD 6.0 provides the
classic Unix write-behind?

The disk activity LED flashes constantly, so it doesn't appear to be
saving up disk writes and then doing a bunch at once,

Is the data coming in at a fairly constant rate?

Yes.

you've got plenty of RAM

The machine has 2 GB. I wonder if the process is getting its fair share?
I have been observing other problems where disk activity to one disk
will make an unrelated process reading data from a different disk *very*
unresponsive.

freebsd-questions@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-questions>

To unsubscribe, send any mail to "freebsd-questions-unsubscribe@xxxxxxxxxxx"