

RE: suffering from an apparently broken tcp

Source: <http://linux.derkeiler.com/Mailing-Lists/Debian/2004-06/0495.html>

From: Kim Sparrow (ksparrow_at_lightpointe.com)

Date: 06/04/04

Date: Thu, 3 Jun 2004 20:05:56 -0700

To: "Paul Galbraith" <paul@paulgalbraith.net>

Well, I'm really starting to convince myself that this is a hardware problem.

1) I think this may be the computer that was experiencing similar problems when running Win2k. It's a 50/50 chance that it was this box.

2) I ran ethereal on it, and it frequently (but not always) reported that outgoing packets had a checksum error at the TCP layer. This can be fixed by setting the `hw_checksums=0` option for the 3c59x module, which forces software calculation of the FCS (I seem to recall that much of the 3c59x can calculate TCP checksums in hardware). Strangely enough, as far as I can tell the Windows boxes didn't seem to mind these errors. It doesn't seem to affect throughput.

3) `ifconfig` reports a really large number of receive errors. Running `ifconfig` before and after a large file transfer, there were 419 frames received, and 129 frame errors!

4) I've seen a few reports of somewhat similar problems on the 3c920, apparently a pretty common NIC chipset in Dells. Inexplicable slow transfers in one direction.

One thing I figured out is that the baby switch in my office is crappy. Cutting that out at least makes the link usable (transfers no longer break after 64k) though it's still marginally unusably slow, at ~50kB/s. Considering that this will be a revision control server, it needs to be a bit snappier than that!

One of the curious things I'm seeing is that data transfer occurs in bursts with a period of .32 seconds, which would explain the 50kB/s. Most of the time three TCP continuation frames come in back-to-back, then there's that .32 second gap... and then three more frames. I'm no TCP or SMB expert, but it looks to me like one of the ACK frames is getting lost in there. That might be corroborated by the unusually high frame error count in `ifconfig`. (I can make a `libpcap` dump if anybody's really that interested.)

Debian–User: RE: suffering from an apparently broken tcp

The thing that still gets me is that downloading from the Internet is blazingly fast, it's only on the local network that's dreadfully slow. I don't know. I've already tried swapping ports to our main switch, which didn't make a difference. So at this point I'm inclined to stick this hard drive in a different box. We've got a handful of these Precision 420s sitting around, so I can hope that one of them will work nicely!

Thanks for the help! Well, it didn't exactly "help", but it's nice to have some moral support. Anyways, I haven't tried netperf, but ethereal is pretty sweet. If the motherboard swap doesn't help, I may have to hook it up our Ixia network analyzer (does 100base–T, OC–3, OC–12, GigE, and is also pretty sweet, despite the steep learning curve). Still, I'd rather just have everything work!

— Kim

-----Original Message-----

From: Paul Galbraith [mailto:paul@paulgalbraith.net]
Sent: Tuesday, June 01, 2004 19:45
To: Kim Sparrow
Cc: debian–user@lists.debian.org
Subject: Re: suffering from an apparently broken tcp

Kim Sparrow wrote:

- > *So I managed to set up a Debian Woody box with Tomcat + Scarab,*
- > *Apache + Subversion, winbind authentication, Mailman, and a few other*
- > *goodies. I thought that everything was fine, until I tried to move the*
- > *existing Subversion repository over to the new system via SMB. I then*
- > *found that files larger than 64k would transfer at pitiful rates --*
- > *essentially, chunks (64k or smaller) of the files would float over*
- > *with*
- > *gaps of many seconds between them. At first I thought the problem was*
- > *essentially a Samba problem, but I achieved similar (lack of) results*
- > *with FTP and HTTP. This behavior is limited to the local network; file*
- > *transfer from the Internet moves at a good clip. Additionally, pulling*
- > *a*
- > *file from the Linux box to another computer on the network works just*
- > *fine.*
- >
- > *Now I'm at a loss for what's going on, and Linux system administration*
- > *isn't at all my specialty. I've looked all over the Internet, and only*
- > *found one message thread noting similar behavior: gaps in transmission*
- > *from the Linux box to Win2k, but good receive behavior. The*
- > *resolution:*
- > *it went away by itself! Anybody have a clue? The thing is essentially*
- > *unusable as it is!*
- >
- > *Relevant (?) specs:*
- > *Dell Precision 420 – Dual 800MHz P3, 512MB RAM*
- > *Integrated 3com 3c920 (3c905C compatible, according to the Dell site)*

RE: suffering from an apparently broken tcp

Debian–User: RE: suffering from an apparently broken tcp

- > *Kernel: 2.6.5.1 (I started out with 2.4.19; switching was an act of*
- > *desperation).*
- >
- > *Any help would be greatly appreciated!*
- >
- >
- > *Kim Sparrow*
- > *Sr. Software Engineer*
- > *www.LightPointe.com*
- > *Speed of fiber. Flexibility of wireless.*
- >

I suffered from a similar problem on a woody box. After a lot of frustration and testing, I found out that a lot of UDP packets were getting dropped in local network communications during high volume connections. I still don't know exactly what was going on, but I believe that it was at least partly faulty drivers for my nic. Upgrading my kernel to 2.4.x solved my problems. You're already ahead of me there, having upgraded your kernel a few times. I can only suggest grabbing a good high–volume network performance analyzer to see what's going on. I **think** the tool I used was called netperf.

Good luck!

Paul