

Re: Interesting TCP issue

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2007-01/msg00314.html>

- *From:* Uwe Doering <gemini@xxxxxxxxxxxx>
 - *Date:* Mon, 22 Jan 2007 09:15:13 +0100
-

Steve Watt wrote:

In <459AD141.5010502@xxxxxxxxxxxx>, Julian Elischer wrote:

[Snip discussion of symptoms of window scaling broken when talking to at least the skype mail servers.]

we have seen this since 4.x
I think a fix may be in 7.0 but I'm not sure..
I thin kthere is a problem when the far end sets the window down to 1 but scales it by a factor of 2^{big number}.

Andre, can you check out this problem and MFC the correct fix if it is indeed the same problem in 6.2?

It is the same problem; I took the (one-line) fix as indicated by

http://cvs.ironport.com/cgi-bin/viewcvs.cgi/freebsd/src/sys/netinet/tcp_syncache.c.diff?r1=1.84&r2=1

(well, not cvs.ironport.com, which doesn't seem to exist at the moment), and applied the diff from 1.84 to 1.85 and to a 6.2-PRERELEASE box updated around 25 Dec 06. It works like a charm.

I would vote to MFC 1.85 now that 6.2 is out.

The diff, for those who are following along at home, is:

```
=====  
RCS file: /usr/local/www/cvsroot/FreeBSD/src/sys/netinet/tcp_syncache.c,v  
retrieving revision 1.84  
retrieving revision 1.85  
diff -u -p -r1.84 -r1.85  
--- src/sys/netinet/tcp_syncache.c 2006/02/09 21:29:02 1.84  
+++ src/sys/netinet/tcp_syncache.c 2006/02/28 23:05:59 1.85
```

Re: Interesting TCP issue

```
@@ -682,7 +682,7 @@ syncache_socket(sc, lso, m)
tp->t_flags |= TF_NOOPT;
if (sc->sc_flags & SCF_WINSSCALE) {
tp->t_flags |= TF_REQ_SCALE|TF_RCVD_SCALE;
- tp->requested_s_scale = sc->sc_requested_s_scale;
+ tp->snd_scale = sc->sc_requested_s_scale;
tp->request_r_scale = sc->sc_request_r_scale;
}
if (sc->sc_flags & SCF_TIMESTAMP) {
```

I wonder whether it is that easy. As far as I can tell the commit to HEAD actually comprised changes to three files:

http://www.freebsd.org/cgi/cvsweb.cgi/src/sys/netinet/tcp_input.c.diff?r1=1.290&r2=1.291
http://www.freebsd.org/cgi/cvsweb.cgi/src/sys/netinet/tcp_syncache.c.diff?r1=1.84&r2=1.85
http://www.freebsd.org/cgi/cvsweb.cgi/src/sys/netinet/tcp_var.h.diff?r1=1.127&r2=1.128

How about the modifications in 'tcp_input.c'? Are they relevant to the problem this thread is about? If so, assessing the correctness of an MFC might prove to be a little harder.

Uwe

--

Uwe Doering | EscapeBox – Managed On-Demand UNIX Servers
geminix@xxxxxxxxxxxxx | <http://www.escapebox.net>

freebsd-hackers@xxxxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxxxxxxx"