Re: read() returns ETIMEDOUT on steady TCP connection

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/net/2008−04/msg00232.html

- *From*: Andre Oppermann <andre@xxxxxxxxxxx>
- *Date*: Mon, 21 Apr 2008 15:46:46 +0200

Mark Hills wrote:

> On Mon, 21 Apr 2008, Andre Oppermann wrote:
>
> > Mark Hills wrote:
> >
> > > On Sun, 20 Apr 2008, Peter Jeremy wrote:
> > >
> > > > I can't explain the problem but it definitely
> > > > looks like a resource
> > > > starvation issue within the kernel.
> > >
> > > I've traced the source of the ETIMEDOUT within the kernel
> > > to tcp_timer_rexmt() in tcp_timer.c:
> > >
> > > if (++tp−>t_rxtshift > TCP_MAXRXTSHIFT) {
> > > tp−>t_rxtshift = TCP_MAXRXTSHIFT;
> > > tcpstat.tcps_timeoutdrop++;
> > > tp = tcp_drop(tp, tp−>t_softerror ?
> > > tp−>t_softerror : ETIMEDOUT);
> > > goto out;
> > > }
> >
> > Yes, this is related to either lack of mbufs to create a segment
> > or a problem in sending it. That may be full interface queue, a
> > bandwidth manager (dummynet) or some firewall internally rejecting
> > the segment (ipfw, pf). Do you run any firewall in stateful mode?
>
> There's no firewall running.

> I'm new to FreeBSD, but it seems to implies that it's reaching a limit of a number of retransmits of sending ACKs on the TCP connection receiving the inbound data? But I checked this using tcpdump on the server and could see no retransmissions.

When you have internal problems the segment never makes it to the wire and thus you wont see it in tcpdump.

Please report the output of 'netstat −s −p tcp' and 'netstat −m'.

Posted below. You can see it it in there: "131 connections dropped by rexmit timeout"