

Re: drift value very large and very unstable

Source: <http://newsgroups.derkeiler.com/Archive/Comp/comp.protocols.time.ntp/2008-03/msg00066.html>

- *From:* andy.helten@xxxxxxxxxxxxx (Andy Helten)
 - *Date:* Thu, 6 Mar 2008 16:23:59 GMT
-

The good news is that "new ntp.conf" appears to work! This is the first configuration that has produced reasonable results, granted it could still be a fluke since the drift was rather unpredictable (but *_always_* settled near +/-500ppm). The bad news is that we *_require_* some of the commands removed from ntp.conf (at least burst and step). After letting ntp run with the "new ntp.conf" for at least 16 hours, the drift had stabilized around 33ppm:

```
sbc1 root 1->ntpq -crv
assID=0 status=0444 leap_none, sync_uhf_clock, 4 events,
event_peer/strat_chg,
version="ntpd 4.2.4p0@xxxxxx Tue Jan 8 16:23:44 UTC 2008 (1)",
processor="i686", system="Linux/2.6.18.8-RedHawk-4.2-trace", leap=00,
stratum=1, precision=-20, rootdelay=0.000, rootdispersion=0.272,
peer=13451, refid=BTFP,
reftime=c0320bd4.c1843a15 Thu, Mar 7 2002 10:55:00.755, poll=4,
clock=c0320bd5.6dfc379d Thu, Mar 7 2002 10:55:01.429, state=4,
offset=0.029, frequency=33.562, jitter=0.002, noise=0.002,
stability=0.001
```

This test ran with the previously problematic Redhawk kernel and all of the same hardware. To further isolate the problem, I've added the 'burst' command back into ntp.conf, removed the drift file, and restarted ntp.

Andy

Andy wrote:

Rob Neal wrote:

On Mon, 3 Mar 2008, Andy Helten wrote:

-- snippage --

Lose the 'iburst burst' on 16.

Re: drift value very large and very unstable

With the two tinker commands above you give ntpd the requirement to amortize the offset entirely with frequency control.

Are you giving it long enough to do so?

If possible, toss those tinker options and try again.

ntpq -p, ntpq -c as -c "rv &x" (where x is the association index for the refclock 16) and ntpq -crv would be useful.

Rob

Rob,

In this case, the purpose of 'iburst burst' is to decrease startup so that ntp will begin servicing sync requests within a reasonable amount of time. I'm not sure that both are necessary, but definitely one of them (along with minpoll 4) decreases startup time from several minutes to about 20 seconds. I seem to recall reading somewhere in the NTP docs that burst and iburst have no effect on reference clocks -- it simply isn't true for the BC635 (refclock_bancomm.c). Removing them is still worth a try and I will run like that overnight. In fact, I started running ntpd with the ntp.conf below (after making the suggested ntp.conf changes) and the ntpq output below is after only about 25 minutes of ntp operation. This is running the Redhawk 2.6.18 linux kernel on the same exact hardware as was used last night on the Redhat 2.6.9-42 kernel (the relevance of this kernel is mentioned below).

I think I have been giving it enough time to stabilize -- any test I consider legitimate was allowed to run for at least 8 hours. Most tests ran overnight for 18-24 hours and some tests ran over weekends for nearly 72 hours. Results were always the same (very large drift). In fact, if allowed to run long enough, the drift almost always reached the +/-500 max.

The tinker commands are also necessary (at least disabling the step) due to some commercial software that has serious problems with backward time steps. This problem should be fixed in a future version, but that may not be soon enough for us. Even then, we may not want time to step backwards.

I should also provide an update for a test that ran last night in which the base RedHat EL4 Update 4 distribution (2.6.9-42 kernel) was used with ntp 4.2.4p0 and the exact same single board computer and exact same BC635 hardware. This test stabilized at a drift of -35ppm with a very small offset (0.021 milliseconds). This test ran overnight and by late morning the drift was changing only by a few hundredths at a time. In

Re: drift value very large and very unstable

other words, everything was working as expected. So, whatever the problem, it almost definitely is software related (and most likely is a problem with the kernel?).

Regarding the kernel's HZ value and its relation to time loss/gain, is there a way to determine the actual value at runtime? I want the value of HZ that is actually in use in the running kernel. I wasn't able to find a way to do this. By the HZ macro in /usr/include, I get a value of 100 and by the "/boot/config-*" file I see a value of 250. This is why I would like a sysctl type value or /proc entry with the actual HZ value, not a macro or config file. Any ideas?

Thanks,
Andy

```
/******  
new ntp.conf  
/******  
# Debug stuff  
statistics clockstats peerstats loopstats  
statsdir /var/lib/ntp/log/  
filegen clockstats file stats.clock type pid link enable  
filegen peerstats file stats.peer type pid link enable  
filegen loopstats file stats.loop type pid link enable  
  
restrict default nomodify notrap noquery  
restrict 127.0.0.1  
  
driftfile /var/lib/ntp/drift  
  
server 127.127.16.0 prefer mode 2 minpoll 4 # Symmetricom BC635  
tos orphan 6
```

```
/******  
ntpq output  
/******
```

```
sbc1 root 31->ntpq  
ntpq> pe  
remote refid st t when poll reach delay offset  
jitter
```

```
=====  
*GPS_BANC(0) .BTFP. 0 1 4 16 377 0.000 9.121  
3.489  
ntpq> as
```

```
ind assID status conf reach auth condition last_event cnt  
=====  
1 13451 9614 yes yes none sys.peer reachable 1
```

Re: drift value very large and very unstable

```
ntpq> rv &1
assID=13451 status=9614 reach, conf, sel_sys.peer, 1 event, event_reach,
srcadr=GPS_BANC(0), srcport=123, dstadr=127.0.0.1, dstport=123, leap=00,
stratum=0, precision=-21, rootdelay=0.000, rootdispersion=0.000,
refid=BTFFP, reach=377, unreach=0, hmode=3, pmode=4, hpoll=4, ppoll=10,
flash=00 ok, keyid=0, ttl=64, offset=9.121, delay=0.000,
dispersion=0.236, jitter=3.489,
reftime=c0311460.c183a17a Wed, Mar 6 2002 17:19:12.755,
org=c0311460.c183a17a Wed, Mar 6 2002 17:19:12.755,
rec=c0311460.c18428f8 Wed, Mar 6 2002 17:19:12.755,
xmt=c0311460.c1831775 Wed, Mar 6 2002 17:19:12.755,
filtdelay= 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00,
filtoffset= 9.12 9.76 10.44 11.20 12.02 12.93 13.86 14.90,
filtdisp= 0.00 0.24 0.48 0.74 0.99 1.26 1.52 1.79
ntpq> cv
assID=0 status=0000 clk_okay, last_clk_okay,
type=16, timecode="065 22:19:27.764471000 0", poll=110, noreply=0,
badformat=0, baddata=0, fudgetime1=0.000, stratum=0, refid=BTFFP,
flags=0
ntpq>
```

questions mailing list
questions@xxxxxxxxxxxxxx
<https://lists.ntp.org/mailman/listinfo/questions>