

Re: Time Synchronizing Between Two Servers

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2007-05/msg00421.html>

- *From:* "jdow" <jdow@xxxxxxxxxxxxxxx>
 - *Date:* Mon, 7 May 2007 16:45:13 -0700
-

From: "Chuck Swiger" <cswiger@xxxxxxx>

On May 4, 2007, at 9:10 AM, RW wrote:

On Thu, 03 May 2007 11:07:34 -0400
Chuck Swiger <cswiger@xxxxxxx> wrote:

Sun SPARC machines have good HW clocks, and also some
of the newer
Macs also seem to have consistently low values in ntp.drift
and
handle timekeeping well.

Does that matter?

A good question— the answer seems to be that it depends.

The RTC time is almost immediately overridden by ntpdate. The
drift is a systematic error that ntpd allows for. I would
have thought that the only significant issue, is whether the system
loses timer interrupts under load.

There are limits to how rapidly ntpd will slew the clock via adjtime (); the smaller the
intrinsic drift of the HW clock, the sooner any adjustment (beyond the initial stepping at
system boot via ntpdate) will complete. This only matters to stratum-2 and higher systems—
anything with a primary reference clock (GPS/WWV/ACTS/etc) is going to sync to that and
ignore the local HW clock entirely.

On good operating systems, that is to say ones in which the NTP code can
get in and slightly alter the HZ clock timing, it implements a phase
locked loop with a lot of filtering on the data returned in the NTP polls.
So the concept of "good oscillator" is a little different from what you
might presume. It basically means one that is stable with time and with

Re: Time Synchronizing Between Two Servers

temperature variations seen as installed. Wide temperature excursions (in this context that could be as little as 10 degrees C) can cause errors. The loop tries to adapt. But time setting might fall back on frequent correction to be satisfactory.

Most PCs have floor sweepings for their crystal oscillators. Those with good AT cut crystals are, more or less coincidentally, working at a temperature that is fairly benign for temperature variations. The frequency error versus temperature curve for the AT cut crystal is an S curve. It starts negative at low frequencies (maybe -25 to -50 ppm). It runs up to a peak, still at relatively low temperature, of about $+25$ to $+50$ ppm. Then it runs back down to the -25 to -50 ppm range at around 40 to 50 degrees C depending on the precise angles at which the crystal blank was cut. Then it swings back upwards again. The basic error of (only) -25 to -50 ppm can be tweaked out with software pretty easily. (In the bad old days trimming the capacitance across the crystal served the same purpose. And when a precise quartz standard is needed this technique still prevails in various forms.)

So basically if you either get lucky or the motherboard was built with a quality (but uncompensated) quartz oscillator you may get relatively good performance over machine room temperatures. It might be as good as a small number of parts per million. About 11 ppm is a second a day for reference. As long as NTP can get in and modify the division ratio, ideally on a tick per tick basis, it can compensate out time variances to the point that you remain remarkably accurate even after a day of being disconnected from the network.

The second major problem is aging. Oscillators change frequency with age. Precisely how they change depends on the care with which they are made, the evacuation or leakage of the crystal can, whether the crystal is well baked out, and how long it has been turned on. Once the oscillator has been on for a few days NTP can get a decent estimate of the long term aging characteristics of the oscillator and compensate for it as well. This is probably why "server class" motherboards have their good reputation. The oscillators are still probably floor sweepings or perhaps slightly better quality "premium" floor sweepings. But if they are on 24/7 their aging characteristics pretty much settle down and become predictable. As long as it is predictable NTP can correct for it.

(FWIW the oscillators on the old block II GPS satellites were intended to include one Cs standard (as an experiment) and three Rb standards. Rb standards are not primary standards. But their phase noise qualities are superior to Cs. Cs is a primary standard. But they do show some variance.

(The military decided they did not want enemies to be able to easily plunk a cruise missile down Jimmy Carter's toity in the White House. So they implemented a means to deny accuracy to the enemy. They corrupted the clock frequencies. I built the frequency synthesizer which did this and made comments back up the chain that this is also a prime way to ensure maximum GPS constellation accuracy. When you have a synthesizer capable

Re: Time Synchronizing Between Two Servers

of correcting an oscillator to parts per ten to the thirteenth accuracy you can move the correction up or down to get it to about 2–4 parts per 10^{13} th. Then you can move up and down one count to cause the average over time to be something into the ten to the fifteenth range if the oscillator's accuracy over that interval is good enough. The comments I heard come back down the chain amounted to "yum yum."

(This is essentially what NTP does at tens to hundreds of parts per billion levels.)

{^_^} Joanne

freebsd-questions@xxxxxxxxxxx mailing list

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

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