Synchronizing the computer clocks at SUSI

W. J. Tango

22 January 2007

Abstract

This brief note explains how the clocks on the SUSI computers are synchronized to UTC and how to check that the clocks are OK. These notes currently apply to the following computers: `arthur`, `caradoc`, `gaheris`, `gareth` and `peleas`.

1. Network Time Protocol (NTP)

SUSI uses the Network Time Protocol (NTP) to synchronize the clocks on the various (Linux) computers. NTP is a TCP/IP protocol that queries remote "peers" (reference clocks). It corrects for propagation delays and jitter and is the current standard for time distribution over networks. It is built into more recent Linux kernels, providing extremely good time synchronization; on older systems it is implemented as a daemon (this is the case for `gaheris` and `gareth`).

NTP was developed for banking, telecommunications and commercial applications where accurate time synchronization between remote computers is essential.

Standard PCs have notoriously poor internal clocks. Clock drifts of many seconds per day is not unusual. In most cases this is merely an annoyance but it can be a serious problem. NTP is ideal for synchronizing PC clocks and most local area networks now include an NTP server to synchronize the PCs attached to the LAN.

1.1 The choice of peers

The primary clock for SUSI is `susintp.physics.usyd.edu.au`. It is a specialized computer that only does one thing: distribute time. Its reference clock is our own GPS receiver. The GPS receiver and the `susintp` time server are located in the left hand bay in the North control rack.

NTP documentation encourages users to employ more than one peer. This provides backup in case one of the peers is off-line or is malfunctioning.

A recent development has been the establishment of NTP "pools." An NTP peer drawn from the Australian pool has the IP address

n`.au.pool.ntp.org`

where n is an integer (0, 1 or 2). The pool manager randomly assigns participating peers to these addresses, and automatically switches them every hour. The pool manager excludes peers that may be temporarily out of service and distributes the load equitably over the peers in the pool.

1.2 ntpdate

Another NTP tool is `ntpdate`. This is like the Linux `date` command, except that it sets the computer clock to the time read from an NTP peer. Unlike the NTP daemon or the kernel module it does not actively synchronize the local clock to the remote peer; it simply sets the time to agree with the remote clock. It is mainly used as a step in the boot procedure in order to set the local clock to within a second or so of the correct time prior to starting the NTP daemon. This is important, since the daemon (or kernel equivalent) will panic and stop working if the discrepancy between the local clock and NTP time is too great.

2. Clock diagnostics

On `arthur` the time is displayed in the toolbar at the bottom of the screen. Before starting an observing session you should always check that the displayed time is correct. If not you will need to run some diagnostic tests to find out what the problem is (see below).

When you first open a remote `ssh` session on one of the other computers (e.g., `peleas`) you should check the time by running the `date` command. This will display the current time and date on the remote computer.

If the date or time is incorrect then you will need to do some tests to identify the problem. The main diagnostic tool is the NTP query function `ntpq`.

2.1 ntpq

This tool is modeled on the familiar `lpq` used to query the printer.

To run the query tool, type in the command:
verb!/usr/sbin/ntpq!
(you do not need to be superuser to do this). A special prompt will appear. Type the command "`peers`" at the prompt and a list of the peers currently in use will be listed. An asterisk in front of a name indicates that the peer in question is deemed to be the most reliable. A "+" indicates that a peer is being used by NTP while a "−" indicates that the peer has been rejected for one reason or another.

If you get a strange message about peers being inaccessible, or no peers are listed, the most likely explanation is a networking problem. In this case one should check that the computer(s) can access off-site hosts (e. g., `silliac`). The command
`ping silliac.physics.usyd.au`
can be used, for example, to determine if the computer can access `silliac`. One can also ping off-site NTP servers. If, for example, one can access sites using TCP/IP but not with NTP there is probably a networking problem. SUSI's access to the outside world is through the CSIRO router so you will need to get help from the CSIRO IT people.

The other possibility is that NTP is either not running or has stopped because the local time is grossly wrong. If you suspect that this is the case see the next section.

The peer listing includes diagnostic information about each peer. The most important are the delay, offset and jitter (the last three columns).

The delay is essentially the propagation time between the computer and the peer (in microseconds). The offset is the most important parameter: it is the offset in microseconds between the peer and the computer clock. The final column (jitter) is an indication of the

dispersion in the propagation times. For obvious reasons `susintp` will usually (but not always) have the smallest delay and jitter, and will normally be shown with an asterisk.

The first peer should be `susintp`. The remaining three peers are randomly selected from the Australian NTP pool (see 1.1).

To exit the query program, type `q` or `quit`.

2.2 ntpdate

As noted in 1.2 `ntpdate` runs at boot time. If NTP does not appear to be working correctly it may be that the local clock is grossly wrong. The `ntpdate` service cannot be run while the NTP daemon (or kernel module) is operating; the only way to check the local clock with `ntpdate` is to reboot the computer.

The ntpdate message appears towards the end of the list of services and will indicate the peer that it has used and the time offset.

A warning message will appear if `ntpdate` cannot find a peer. If this occurs there is something seriously wrong. The two most likely causes are

1. There is a problem with the network.
2. `susintp` is off-line.

Another sign of trouble is if the clock offset is extremely large. The first thing is to check that `susintp` is broadcasting the correct time. After an extended power failure it can get confused.

If `susintp` is correct then it is a good idea to check the BIOS clock on the machine in question. This can be done by rebooting the computer and stopping the boot procedure when the BIOS has loaded. This can be done by hitting the DEL key at the right point. There is usually a little message at the bottom of the screen (e. g. “Hit ESCAPE to enter the BIOS setup.”) and you may also need to hit another key such as TAB.

The time is always set by going to the first option on the main BIOS setup page, but setting the BIOS clock can be a little tricky. Follow the on-screen instructions and save the settings. The BIOS time will be displayed while the computer boots. You may take more than one attempt to get the time correctly set. Once it is more or less right you will see that `ntpdate` reports only a small correction (typically a couple of seconds).

As noted above, the clocks on PCs are quite often very poor timekeepers, so they may eventually drift out of sync. The fact that the BIOS clocks occasionally need resetting is not indicative of a fault.

However, if a PC consistently shows the wrong time after it has been turned off, and particularly if the time appears to default to some arbitrary epoch like 1 January 2000, the chances are the backup battery needs replacing. The battery is located on the motherboard and can be easily changed.

3. Other information

To change the peers you must edit (as superuser) the file `ntp.conf` located in the `/etc` directory. There is also a drift file that keeps track of the clock drift with respect to the peers. The location of the drift file can be found by inspecting `ntp.conf`.

3.1 Recent problems

The recent problems with clock synchronization (especially on `gaheris`) were due to the fact that the list of peers was out of date. In the case of `gaheris` it was using a non-existent computer to set the date with `ntpdate`! All the machines now use `susintp` to set the time with `ntpdate` and, as noted above, the NTP daemon uses `susintp` and three peers selected from the Australian NTP pool.