Computer System Design and Administration

Topic 6. Network time sync service: ISC NTP

José Ángel Herrero Velasco

Department of Computer and
Electrical Engineering

Secure information service: Puzzle

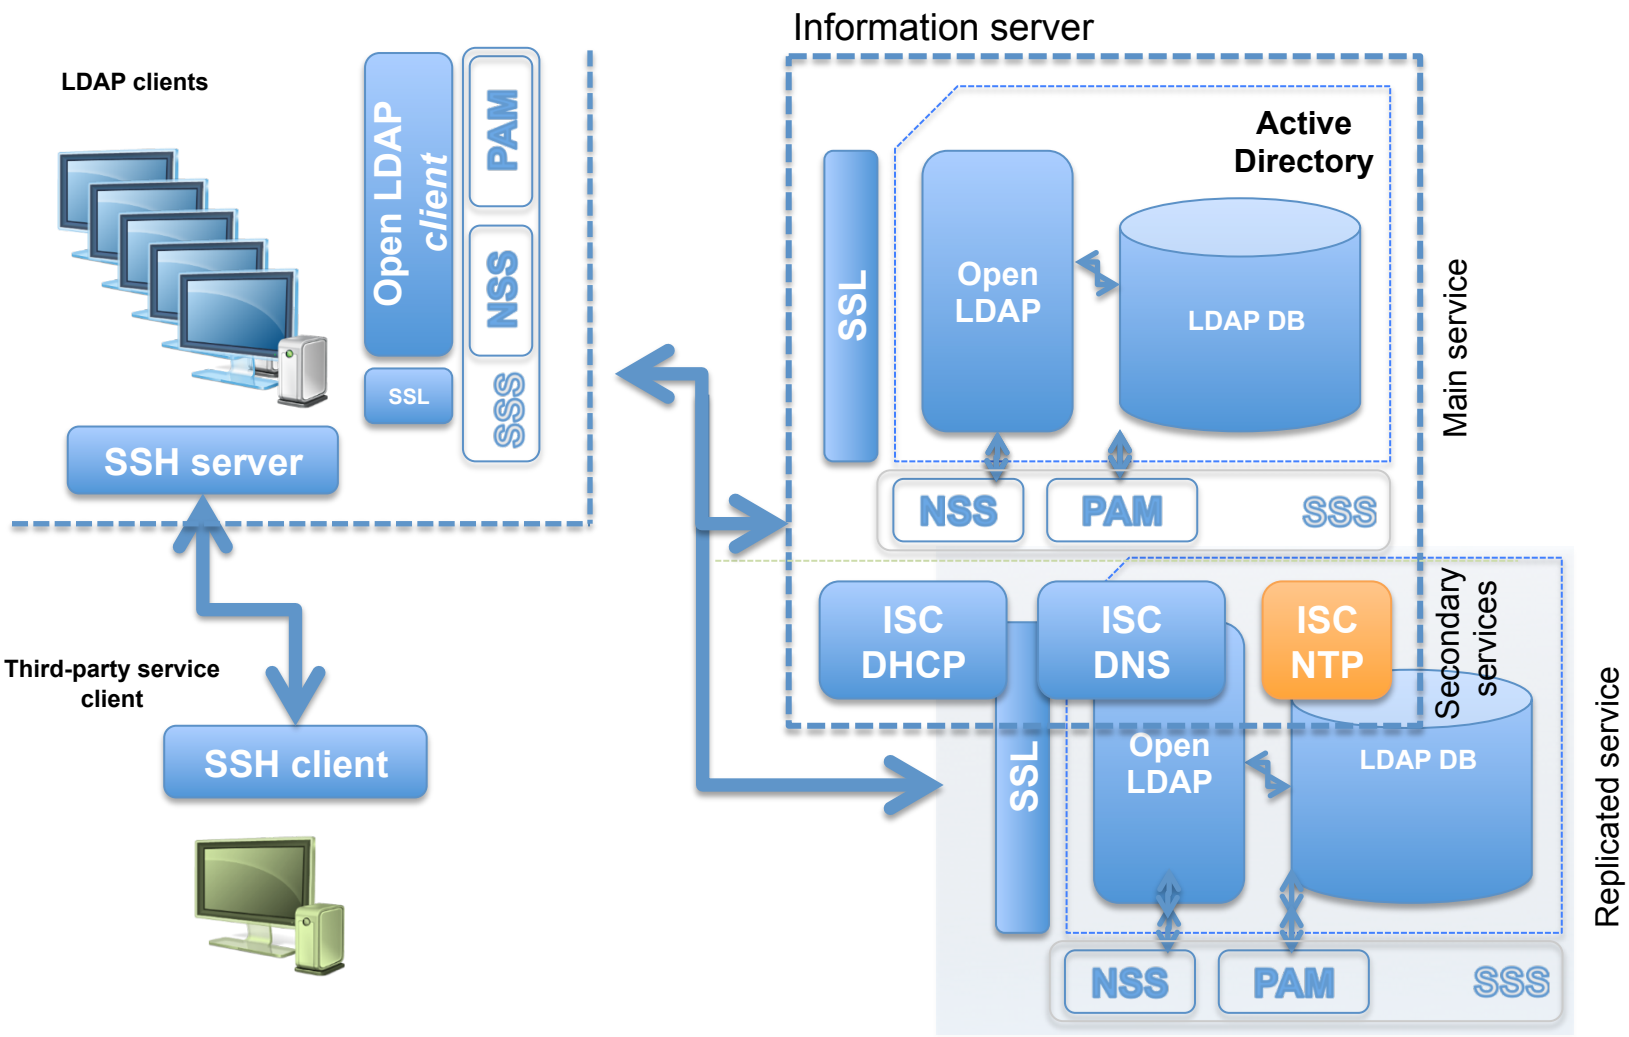

Target: …server convergence

- Installation, configuration and deployment of ***third-party*** **network services** for local *networking* management on the INTRANET:
  - Dynamic configuration service (DHCP): *ISC dhcpd.*
  - Domain name service (DNS): ISC bind9:
  - Network time service (NTP): ISC ntpd:
    - Keeping the **software time** synchronized in accordance with a common time reference.
    - EVERY network host must have the same software time:
      - **Reference time.**
      - **Regular checks (sync).**

Computer time

- In computer systems:
  - **Time** = number of seconds elapsed since a reference time (01/01/1970)[Unix]
- Every computer has **2 clocks:**
  - **Hardware clock:**
    - Integrated in *motherboard* and powered by a small battery:
      - **Computer keeps the hardware time even during shutdowns.**
      - **If you take out this battery → "Reset" (time, BIOS password ?!?!? ⋯):**
        - » Beware if battery runs out!!!!!!
    - Hardware time can be changed by OS or BIOS.
    - It is used to configure the computer **local time.**
  - **Software clock:**
    - It uses the UTC → ***Coordinated Universal Time:***
      - **Primary time standard by which the world regulates clocks and time → From 1 January 1960.**
      - **Successor to Greenwich Mean Time (GMT).**
      - **UTC (from 1970) is defined by:**
        - » **International Atomic Time** (IAT):
          - Atomic reference clocks Cesium atoms → Distributed by GPS (and radio), modems ⋯
        - » With **leap seconds** added:
          - At irregular intervals to compensate for the slowing of Earth's rotation (31s/century ΔT).
        - » UTC according to geographic zones (**Time Zones**):
          - Positive or negative offsets (24) from UTC.
    - In the past, "GMT" was used as reference → *Greenwich Mean Time:* **Mean solar time:**
      - **Astronomical base.**
      - **Stable but not constant…**

→ Both of them are independent of each other, except when OS boots:

- OS uses HW time to set up its SW time (on boot).
- Then, SW time is **synchronized** (UTC) by **NTP.**

NTP: Network time sync protocol

- *NTP: **N**etwork **T**ime **P**rotocol.*
- **Motivation:**
  - Many services and network apps need software clocks to be 100% synchronized (**timestamps**):
    - *Kerberos*, batch processing systems, distributed file systems & databases, *log systems,* developing tools *(makes),* etc…
- **Definition** (NTP):
  - NTP is a **protocol** designed to synchronize the clocks of computers in a variable-latency data network.
    → Selects the best time among **several time sources** and minimizes **cumulative delay.**
  - **Targets:**
    1. **Optimize** *local time* accuracy for *UTC.*
    2. **All *hosts*** on a LAN have their clocks synchronized (use the same software time).
- **Origins and history:**
  - One of the oldest protocols on the **Internet** (since 1979):
    - *Internet Clock Service* (RFC 778):
      - → **Internet services running over a trans-Atlantic satellite network.**
    - Accuracy of only several hundred milliseconds.
  - Versions:
    - 1985. Fuzzball and Unix implemented the **NTPv0** (RFC 958):
      - **David L. Mills (Delaware University - USA).**
    - 1988. The first complete specification: a much more complete specification in the **NTPv1** (RFC 1059).
    - 1889. Introduction of symmetric-key authentication in the **NTPv2** (RFC 1119).
    - 1992. Introduction of formal correctness principles in **NTPv3** (RFC 1305):
      - **1994: NTPv3 works for a new version of NTP: SNTP (RFC 2030).**
    - 1994-XX. Analysis of all sources of error, external pulses calibration and more new features...
    - 2010. **NTPv4** (RFCs 5905/6/7 y 8) → Continues to be a developing version:
      - **The reference implementation is currently maintained as an open source project led by Harlan Stenn.**

NTP: Basis & features

- **Fundamentals:**
  - NTP needs a **reference time** to define the true time (network time):
    - NTP system uses **UTC** as reference time, based on *International Atomic Time (IAT).*
    - This "reference time" will be assigned by the ***hierarchical system*.**
  - NTP is a **fault tolerant** protocol (Bellman-Ford shortest-path spanning tree):
    - The time data comes from **multiple sources.**
  - NTP is highly **scalable:**
    - It can increase in client numbers…
  - NTP can sync the **host time** even though network is "down":
    - Temporally… (*fudge* + `driftfile`).
- **Precision:**
  - Strongly dependent on the type of network:
    - From 5-100 ms (Internet) to 200 µs (LAN).
- **Architecture:**
  - NTP uses a *hierarchical system* of servers on the Internet (Servers → *Peers*):
    - NTP **stratum** model.
    - Each level → *stratum (ID).*
  - Many *peers* provide time **redundancy.**
- **TCP/IP protocol:**
  - **Transport layer.**
  - **NTP package** format (NTP/SNTPv4):
    - The following to IP/**UDP** headers…
    - The 64-bit timestamps:
      - **Compute the *offsets*.**

NTP architecture

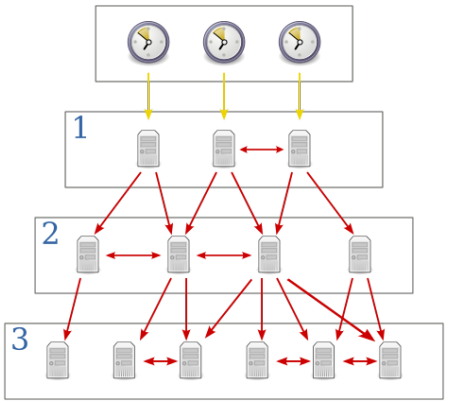

***TRANSPORT* layer (TCP/IP)**

	Octet +0								Octet +1								Octet +2								Octet +3							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
+0	LI		VN			0	0	0	*Statum*								*Poll*								*Precision*							
+4	*Synchronizing Distance*																															
+8	*Estimated Drift rate*																															
+12	*Reference clock Identifier*																															
+16	*Reference clock Timestamp*																															
+24	*Originate Timestamp*																															
+32	*Receive Timestamp*																															
+40	*Transmit Timestamp*																															

NTP: Computing the "right time"

- **NTP algorithms** for time computing:
  - The key:
    - → Selects the *best time* among many sources.
    - → Minimizes **cumulative delay** (minimizes the accumulated error).
  - **Architecture and Algorithms:**
    - **(1) Clock Filter** algorithm:
      - **Time references are calculated based on *round trip delay* and interval observations.**
      - **Then, it selects the *offset* with minimum delay.**
    - **(3) Clustering** algorithm:
      - **Selects the best suite of servers *(peers)* and *combines* their differences to determine the *offset.***
    - **(2) Intersection** algorithm (**default**):
      - **Based on Marzullo's algorithm.**
      - **A typical NTP client will regularly *pool* 3 or more servers on diverse networks:**
        - » Client must compute their time ***offset*** and ***round-trip delay*.**
        - » Among several servers, it requires that the midpoint of the interval be at the *intersection.*
    - **(4) Combinational** Algorithm:
      - **Computes the mean time *offsets.***
    - **(5) Clock Discipline** Algorithm:
      - **It is an adaptive parameter, hybrid phase/frequency-lock feedback loop → Minimize the jitter (dispersion).**

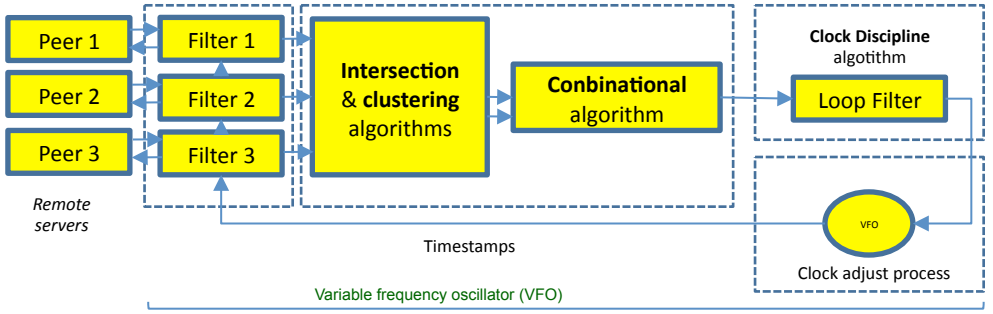

NTP: Service architecture (Topology)

- NTP uses a **hierarchical**, semi-layered system of time sources:
  - Each level of this hierarchy is termed a **"stratum".**
  - Each **stratum** is assigned a ID (0 .. N).
- The **stratum ID** represents the **distance from the reference clock (n + 1):**
  - Stratum is not always an indication of quality or reliability.
- **Clock strata:**
  - ***Stratum* 0:**
    - **High-precision timekeeping devices** → **Atomic** (cesium, rubidium) **clocks.**
  - **Stratum 1:**
    - These are computers whose system clocks are synchronized within a **few microseconds** of their attached **stratum 0 devices.**
    - They may peer with other stratum 1 servers (backups).
  - **Stratum 2, 3… to 14 (although it supports up to 256):**
    - These are computers that are synchronized over a network to a **stratum 1, 2… to 13 servers.**
    - They can themselves act as servers for stratum 3 computers, and so on.

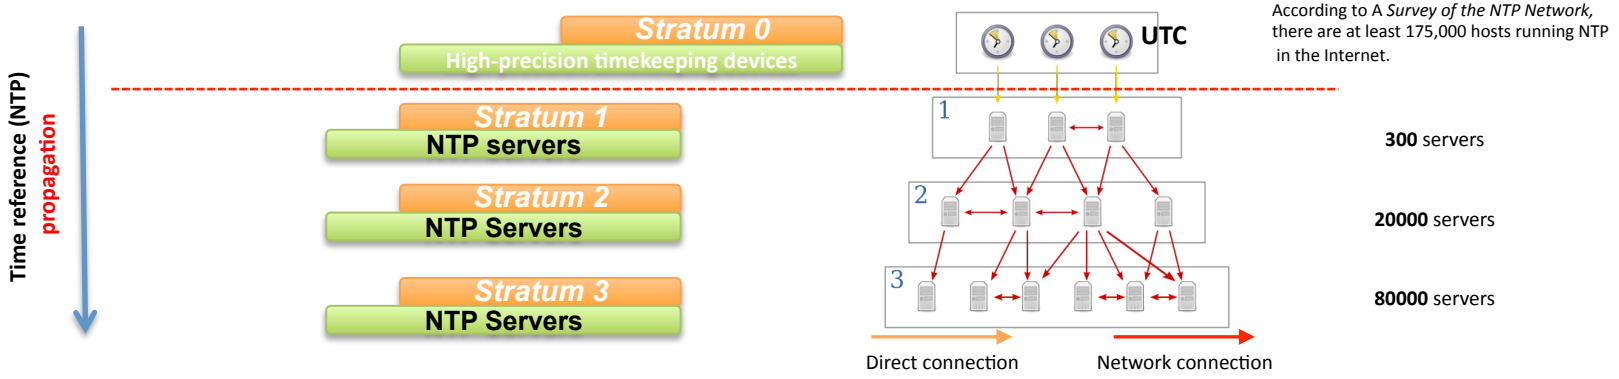

NTP: Operational basis

- When a NTP client requests a time sync (**Client/server mode**):
  - If server is a direct **time source (stratum 0):**
    - The server sends its **"local time"**, **"time zone"** and **stratum.**
  - Else:
    - The server sends a *computed time:*
      - **Using data from servers of the same or higher stratum.**
      - **Using NTP algorithms.**
- The client must recalculate the time obtained:
  - Using the **Intersection** algorithm:
    - **Time offset** and **round-trip delay.**
- Public NTP server list:
  - Public NTP Primary (stratum 1) Time Servers:
    - http://support.ntp.org/bin/view/Servers/StratumOneTimeServers.
  - Public NTP Secondary (stratum 2) Time Servers:
    - http://support.ntp.org/bin/view/Servers/StratumTwoTimeServers.
  - Public NTP Pool Time Servers:
    - http://support.ntp.org/bin/view/Servers/NTPPoolServers.

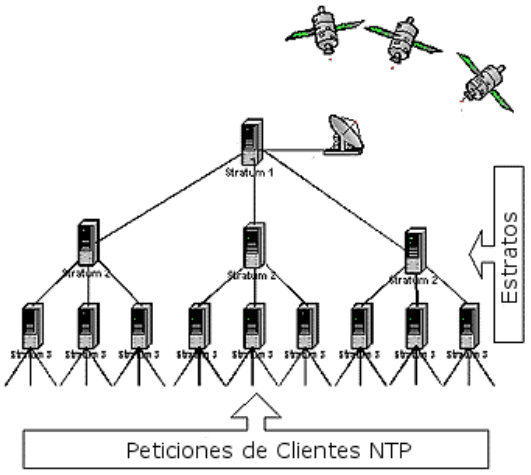

Source: https://memoria.rnp.br.

NTP: Network time sync service

- **NTP** on **Linux/UNIX:**
  - Service managed by the **`ntpd` daemon** (most of protocol is implemented in it):
    - **Operation** modes:
      - **Client/server mode:**
        - » The client requests "time sync" to a particular NTP server.
      - **Broadcast mode (client/server):**
        - » Many clients may be sync with one or more NTP servers.
        - » Operation:
          - Server sends "time" to *everybody.*
          - Clients listen only!!!
          - → It reduces network traffic (**LAN**).
      - **Multicast mode:**
        - » One or more servers periodically *multicast* the time to the servers in the network.
        - » Only in NTPv4.
      - **Symmetric mode:**
        - » It enables NTP servers to synchronize with each other to provide "time reference" copies (**Horizontal sync**):
          - To improve the accuracy of their synchronization over time.
    - NTP is defined for TCP/IP networks:
      - **UDP 123**
    - NTP security:
      - **NTP (v4) is able to guarantee the server authenticity.**
      - **NTP may use symmetric-key and public key-cryptography modes:**
        - » Public/private keys.

- **Protocol alternatives:**
  - Thera are different deployments of the same protocols (NTP)
    - **Protocol variants.**
    - **SNTP** (**S**imple **N**etwok **T**ime **P**rotocol): RFC 5905:
      - **More simple (no storage of previous connections) and less precise!!!**
      - **For embedded devices.**

NTP: Service installation (ISC NTP)

Source: www.isc.org.

- **NTP server and tools** installation (**Server**):
  - Stage 1. **Hardware clock** setting:

```
$ hwclock --set --date="10/11/2010 16:27:30"
$ hwclock --hctosys
```

  - Stage 2. **Time zone** setting (local):

```
$ dpkg-reconfigure tzdata
```

  - Stage 3. Service software **installation:**

```
$ apt-get update
$ apt-get install ntp ntp-doc
$ update-rc.d ntp defaults
```

- **Lab 2.** We should deploy a **local NTP service.**

NTP: Service configuration

- **NTPd** *service* **main** configuration file:

  `$ vi /etc/ntp.conf`

  - Main configuration entries:
    - `server <ip>`:
      - **NTP source public servers list (1/2 stratums).**
      - **It is recommended to have at least 3 servers.**
    - `restrict <ip> [opciones]`:
      - **Access control restrictions.**
      - **By default, the NTP server will be accessible from all internet hosts.**
      - **It establishes which hosts can use the NTP service and which do not.**
    - `fudge <ip> stratum <num>`:
      - **Routing control (pseudo IP) → Backup.**
      - **It is only used when NTP servers fail (unavailable):**
        - » NTP server sync itself.
    - `keys <fichero>`:
      - **Key file for queries.**
    - `driftfile <fichero>`:
      - **Drift file → The drift file is used to store the frequency offset between the system clock running at its nominal frequency and the frequency required to remain in synchronization with UTC. Default: `/var/lib/ntp/ntp.drift`.**
    - `statsdir <directorio>`:
      - ***Logs* and statistics file for NTP service.**
    - `broadcast <ip>`:
      - **Server configuration in broadcast mode.**

NTP: Daemon configuration

- **NTPd** *daemon* **main** configuration file:

  `$ vi /etc/default/ntp`

  - NTP daemon (**`ntpd`**) parameters defined as variables:
    - They are used by startup script:
      - **`/etc/init.d/ntp.`**
  - Sample:
    - `NTPD_OPTS='-g'`
  - To view the options available in the NTP service:
    - `$ man ntpd.`

- More important things about NTP service:
  - **Firewalls:**
    - It is necessary to keep port 123 open for UDP:
      - **For incoming and outgoing traffic.**

Examples: Service configuration

Sample

- /etc/ntp.conf

```
driftfile /var/lib/ntp/ntp.drift
statsdir /var/log/ntpstats/

statistics loopstats peerstats clockstats
filegen loopstats file loopstats type day enable
filegen peerstats file peerstats type day enable
filegen clockstats file clockstats type day enable


server hora.rediris.es
server 0.pool.ntp.org
server 127.127.1.0
fudge 127.127.1.0 stratum 13

restrict default kod notrap nomodify nopeer noquery
restrict 127.0.0.1 nomodify

broadcast 192.168.0.255
```

Server list for ½ stratum.

Pseudo-IP address. If any error happens, NTP syncs itself.

- /etc/default/ntp

```
NTPD_OPTS='-g'
```

NTP: Client installation

- NTP client is based on the scheduled run of **`ntpdate-debian`** command.
- NTP client **installation** (**client**):
  - The recommendations for server installation, as in the previous steps, are also valid for NTP clients.
  - Stage 1. **Hardware clock** setting:

```
$ hwclock --set --date="10/11/2010 16:27:30"
$ hwclock --hctosys
```

  - Stage 2. **Time zone** setting (local):

```
$ dpkg-reconfigure tzdata
```

  - Stage 3. Client software **installation:**

```
$ apt-get update
$ apt-get install ntpdate
```

NTP: Client configuration

- **ntpdate-debian** configuration:

```
$ vi /etc/default/ntpdate
```

  - Options:
    - `DATE_USE_NTP_CONF`:
      - **It's only used if host runs `ntpd`.**
        - `/etc/ntp.conf`.
    - `NTPSERVERS`:
      - **NTP servers list used by ntpdat-debian.**

Examples: Client configuration

Sample

- /etc/default/ntpdate

```
# The settings in this file are used by the program ntpdate-debian, but not
# by the upstream program ntpdate.

# Set to "yes" to take the server list from /etc/ntp.conf, from package ntp,
# so you only have to keep it in one place.
NTPDATE_USE_NTP_CONF=no

NTPSERVERS="192.168.0.11 0.debian.pool.ntp.org 1.debian.pool.ntp.org"

# Additional options to pass to ntpdate
NTPOPTIONS=""
```

NTP: Client configuration (regular sync)

- To maintain a client **regular time sync,** we must use the CRON service:
  - **Option 1.** Root *crontab:*

```
$ crontab -e
 */15 * * * * /usr/sbin/ntpdate-debian
$ /etc/init.d/cron reload
```

  - **Option 2.** Temporary *crontab* `/etc/cron.{daily,hourly}`

```
$ vi /etc/cron.hourly/ntpdate
 /usr/sbin/ntpdate-debian
$ chmod 755 /etc/cron.hourly/ntpdate
$ /etc/init.d/cron reload
```

NTP: Checking

- **Checking** if NTP service is *"running":*

```
$ /etc/init.d/ntp restart
$ pgrep ntpd
$ ps -elf |grep ntp
$ netstat -atunp
```

- **Checking** if a *firewall* is setting:

```
$ iptables -L
```

- **Checking** the NTP service sync according to the upper stratum:

```
$ ntpq -p (prints the current software time).
$ ntpdc -loopinfo (prints how the software time is drifted).
$ ntpdc -kerninfo (prints the current aggregated correction).
$ ntptime
```

- **Sync** the client software time:

```
$ ntpdate-debain <ntp server>
```

- **Sync** the client hardware time according to hardware time:

```
$ hwclock --systohc
```