Established technology: NPT for network time synchronization

Alexander Bähr
Open Source Automation Development Lab (OSADL) eG

Main topics

- Concepts about time and time measurement
- Short retrospective in history of time measurement and timescales
- The Network Time Protocol (NTP) and how it works
- NTP measurement results of various setups

Clocks (mechanical/natural oscillator)

- A **clock** consists of an **oscillator** and a **counter** that record the number of increments since initialized with a given value at a given **time**.
- First **clocks**: Sundials, water clocks, incense clocks, hour glasses.
- Since ~1400: Pendulum **clocks** controlled by a harmonic oscillator (the regular swing of a pendulum).

Clocks (electronic/physical oscillators)

- Crystal clock: **Electronic oscillator** regulated by a **crystal** to keep time, accuracy ~0.5 s/day
- **TXCO** (**T**emperature **C**ontrolled **C**rystal **O**scillator), accuracy ~4.3 ms/day
- **Rubidium** clock, accuracy ~1 µs/day
- **Cesium** clock, accuracy ~1 ns/day

Scientific timescales

- These are based on astronomical observations of the sun, the moon and the stars as reference. Prior to 1958: **Ephemeris Time** (ET) based on one complete revolution of the earth around the sun, the **ET second** was defined as 1/86,400 of the mean solar day.
- In 1958: **ET second** was redefined as 1/31,556,925.9747 of the tropical year 1900, the year as 365.242 days, varies +/- 50 ms, increases ~5.3 ms per year.

Time measurement

- A timescale is a continuum of monotonically increasing values that denote time in some frame of reference.
- To measure time in a comparable way, a reference is essential.
- Until ~1960: astronamical observations as reference
- Nowadays physical constants are used (Crystal oscillator, cesium, Rubidium hyperfine transition)

Time definition and standards (1)

- The official SI definition of the second is as follows:

  The **second** is the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium 133 atom.

Time definition and standards (2)

- Time is one of the seven fundamental physical quantities in both the International System of Units (SI) and International System of Quantities and the strictest defined quantity.

Base quantity:	time
Symbol for quantity:	T
Symbol for dimension:	t
SI based unit:	second
SI unit symbol:	s

Civil timescales

- The basis for civil time scale is the International Atomic Time (TAI), a weighted average of the time kept by over 400 atomic clocks worldwide.
- TAI is a continuous monotonically increasing timescale.
- Coordinated Universal Time (UTC) synchronized with TAI since 1958

Leap seconds

- In UTC irregularity of earth rotaional motion is compensated by inserting or skipping leap seconds.
- Necessary to convert the NTP/UTC timescale to TAI.
- Difference (today) 37 s → TAI = UTC + 37s.
- Inserted in the following second 23:59:59 at the last day of June or December → becomes 23:59:60 of that day.

Network Time Protocol NTP

- Designed by David L. Mills of the University of Delaware.
- 1985 NTP version 0 was implemented → RFC 958.
- To synchronize all participating computers to **UTC** within a few milliseconds.
- Usually in a client-server model, can also be used in a peer-to-peer relationship.
- Sends and receives NTP packets using UDP port 123.

Network Time Protocol NTP

- In 2010, RFC 5905 was published containing a proposed specification for NTPv4.
- The reference implementation is currently maintained as an open source project.

NTP terminology

- Resolution → the degree to which a clock reading can be distinguished from another → reciprocal of clock oscillator frequency: 2 GHz → 0.50 ns.
- Precision → the degree to which an application can distinguish one clock reading from another →the latency of a system.
- Accuracy → the degree to which a clock reading differs from the real time.

NTP Strata

Stratum

0 → Reference Clock

1 → NTP Server

2 → NTP Server

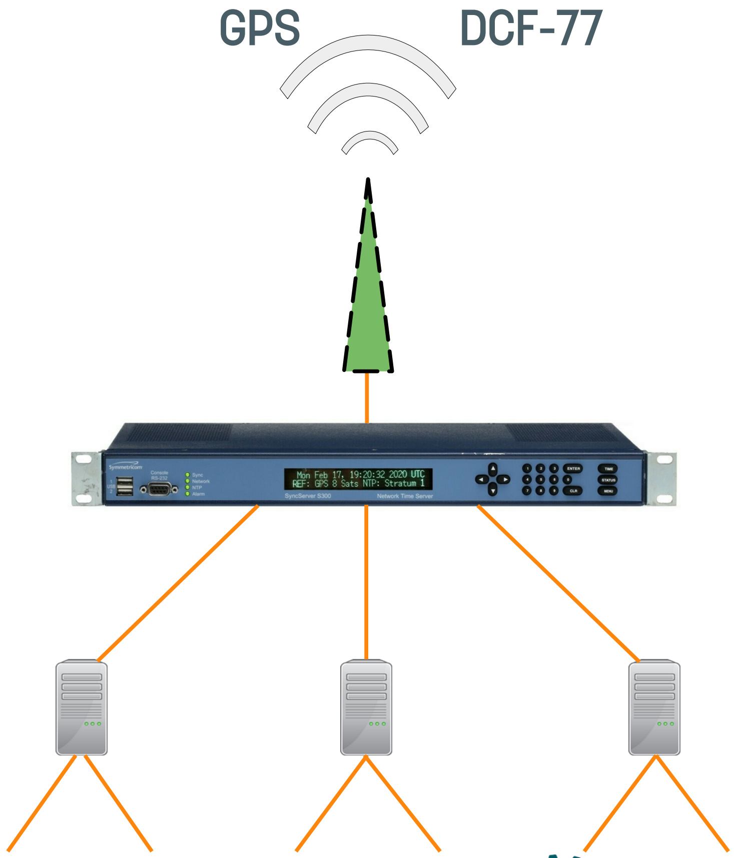

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode     |Stratum                         |Poll intervall             |Precision                 |
Root delay                                                                                                      |
Root dispersion                                                                                                 |
Reference clock identifier                                                                                      |
Reference timestamp: Seconds (32)                                                                               |
Reference timestamp: Fraction (32)                                                                              |
Originate timestamp: Seconds (32)                                                                               |
Originate timestamp: Fraction (32)                                                                              |
Receive timestamp: Seconds (32)                                                                                 |
Receive timestamp: Fraction (32)                                                                                |
Transmit timestamp: seconds (32)                                                                                |
Transmit timestamp: Fraction (32)                                                                               |
Key Identifier (optional)                                                                                       |
Message Digest (optional)                                                                                       |
```

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode    |Stratum                      |Poll intervall           |Precision          |
Root delay                                                                                          |
Root dispersion                                                                                     |
Reference clock identifier                                                                          |
Reference timestamp: Seconds (32)                                                                   |
Reference timestamp: Fraction (32)                                                                  |
Originate timestamp: Seconds (32)                                                                   |
Originate timestamp: Fraction (32)                                                                  |
Receive timestamp: Seconds (32)                                                                     |
Receive timestamp: Fraction (32)                                                                    |
Transmit timestamp: seconds (32)                                                                    |
Transmit timestamp: Fraction (32)                                                                   |
Key Identifier (optional)                                                                           |
Message Digest (optional)                                                                           |
```

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode    |Stratum                    |Poll intervall          |Precision             |
Root delay                                                                                          |
Root dispersion                                                                                     |
Reference clock identifier                                                                          |
Reference timestamp: Seconds (32)                                                                   |
Reference timestamp: Fraction (32)                                                                  |
Originate timestamp: Seconds (32)                                                                   |
Originate timestamp: Fraction (32)                                                                  |
Receive timestamp: Seconds (32)                                                                     |
Receive timestamp: Fraction (32)                                                                    |
Transmit timestamp: Seconds (32)                                                                    |
Transmit timestamp: Fraction (32)                                                                   |
Key Identifier (optional)                                                                           |
Message Digest (optional)                                                                           |
```

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode    |Stratum                        |Poll interval          |Precision             |
Root delay                                                                                            |
Root dispersion                                                                                       |
Reference clock identifier                                                                            |
Reference timestamp: Seconds (32)                                                                     |
Reference timestamp: Fraction (32)                                                                    |
Originate timestamp: Seconds (32)                                                                     |
Originate timestamp: Fraction (32)                                                                    |
Receive timestamp: Seconds (32)                                                                       |
Receive timestamp: Fraction (32)                                                                      |
Transmit timestamp: Seconds (32)                                                                      |
Transmit timestamp: Fraction (32)                                                                     |
Key Identifier (optional)                                                                             |
Message Digest (optional)                                                                             |
```

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode     |Stratum                        |Poll intervall          |Precision         |
Root delay                                                                                          |
Root dispersion                                                                                     |
Reference clock identifier                                                                          |
Reference timestamp: Seconds (32)                                                                   |
Reference timestamp: Fraction (32)                                                                  |
Originate timestamp: Seconds (32)                                                                   |
Originate timestamp: Fraction (32)                                                                  |
Receive timestamp: Seconds (32)                                                                     |
Receive timestamp: Fraction (32)                                                                    |
Transmit timestamp: Seconds (32)                                                                    |
Transmit timestamp: Fraction (32)                                                                   |
Key Identifier (optional)                                                                           |
Message Digest (optional)                                                                           |
```

The NTP packet

```
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32|
LI    |Version |Mode    |Stratum                   |Poll interval          |Precision          |
Root delay                                                                                      |
Root dispersion                                                                                 |
Reference clock identifier                                                                      |
Reference timestamp: Seconds (32)                                                               |
Reference timestamp: Fraction (32)                                                              |
Originate timestamp: Seconds (32)                                                               |
Originate timestamp: Fraction (32)                                                              |
Receive timestamp: Seconds (32)                                                                 |
Receive timestamp: Fraction (32)                                                                |
Transmit timestamp: Seconds (32)                                                                |
Transmit timestamp: Fraction (32)                                                               |
Key Identifier (optional)                                                                       |
Message Digest (optional)                                                                       |
```

The 64 bit NTP timestamp

Bit: 0 31	32 64
Seconds since 1900	Fraction of second

Begin: 1 January 1900, 00:00:00
Ends: 19 January 2038, 03:14:08

Theoretical resolution: 232 ps = 0.000,000,000,232s

Year 2038 problem → 32 bit int overflow

The 64 bit NTP timestamp

Bit: 0 31	32 64
Seconds since 1900	Fraction of second

Begin: 1 January 1900, 00:00:00
Ends: 19 January 2038, 03:14:08

Theoretical resolution: 232 ps = 0.000,000,000,232s

Year 2038 problem → 32 bit int overflow

The 64 bit NTP timestamp

Bit: 0 31	32 64
Seconds since 1900	Fraction of second

Begin: 1 January 1900, 00:00:00
Ends: 19 January 2038, 03:14:08

Year 2038 problem → 32 bit int overflow

Theoretical resolution: 232 ps = 0.000,000,000,232s

Network Time Protocol NTP

A (client) reads current time T_1

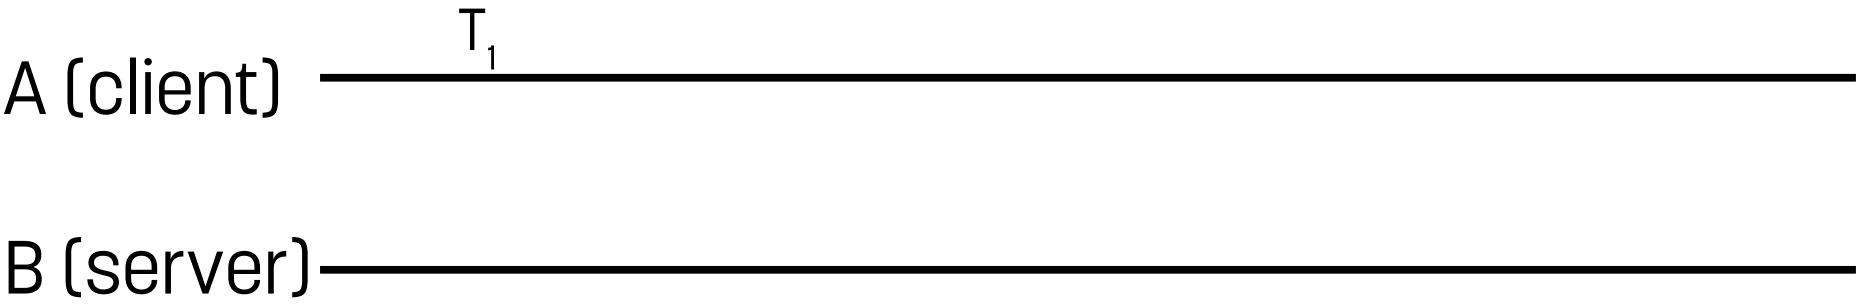

Network Time Protocol NTP

A (client) reads current time T_1 and sends it to B (server)

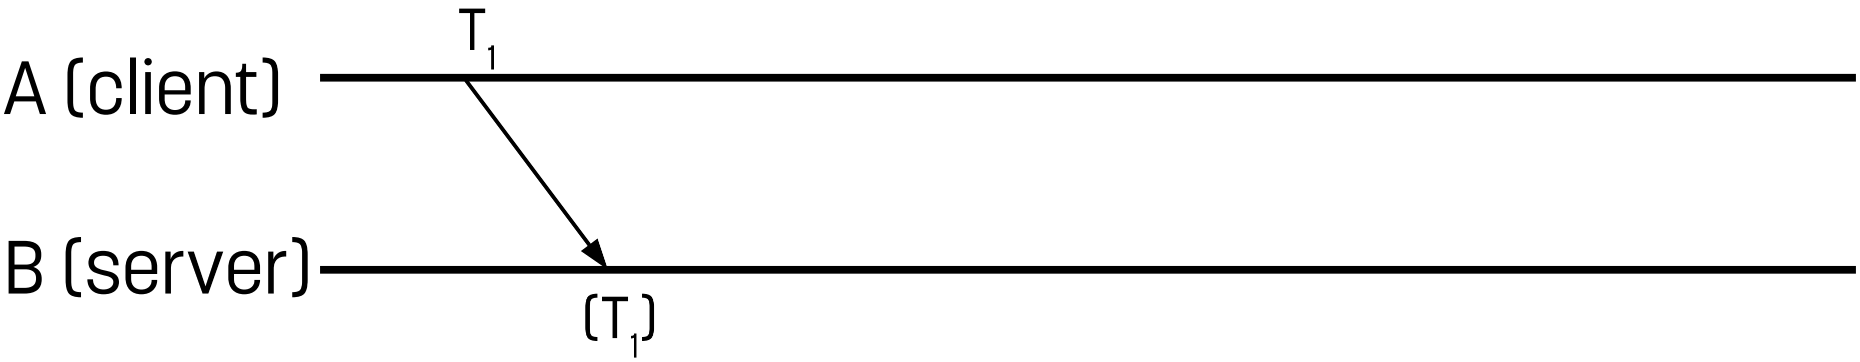

Network Time Protocol NTP

- B (server) reads current time T_2 and saves T_1 and T_2

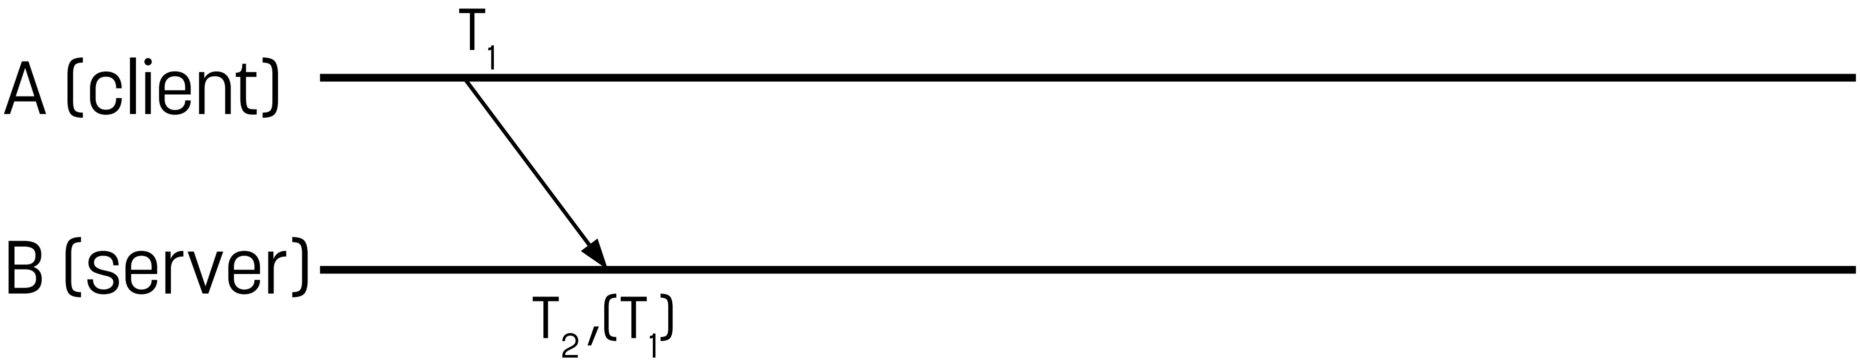

Network Time Protocol NTP

- B (server) reads current time T_3

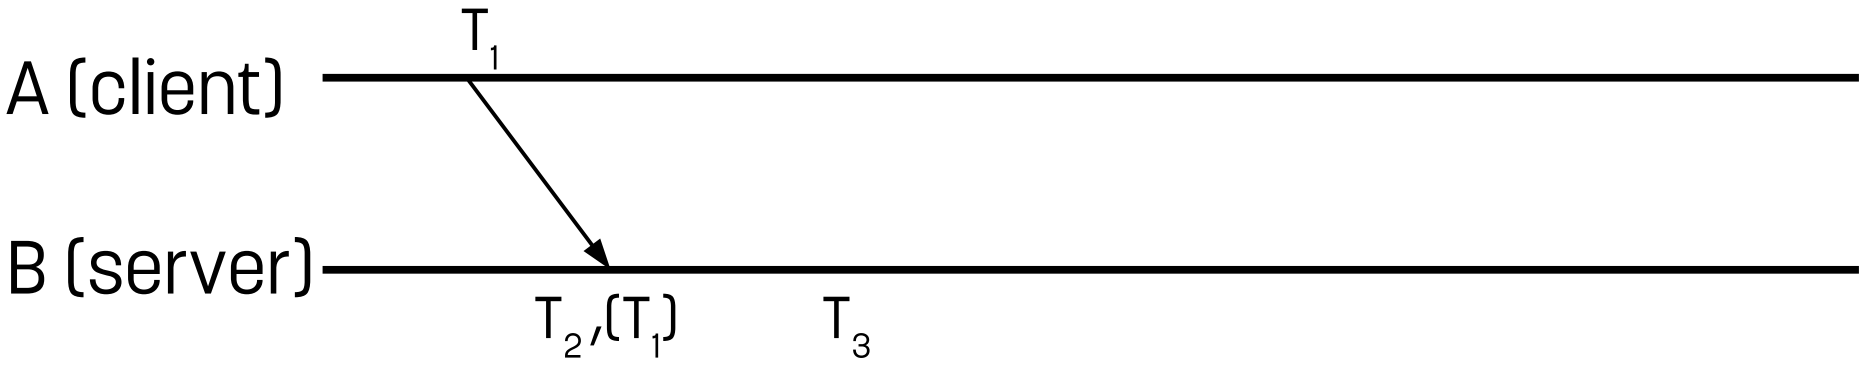

Network Time Protocol NTP

- B (server) reads current time T_3 and sends it along with the saved T_1 and T_2 to A (client)

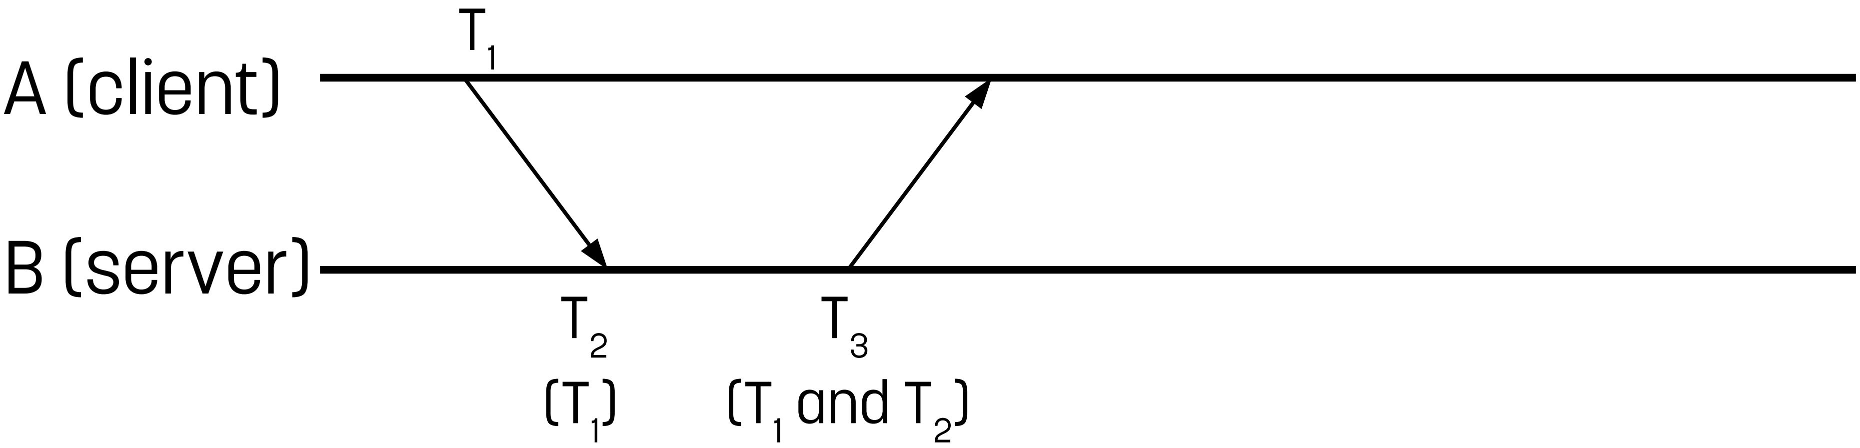

Network Time Protocol NTP

- A (client) reads current time T_4 on arrival

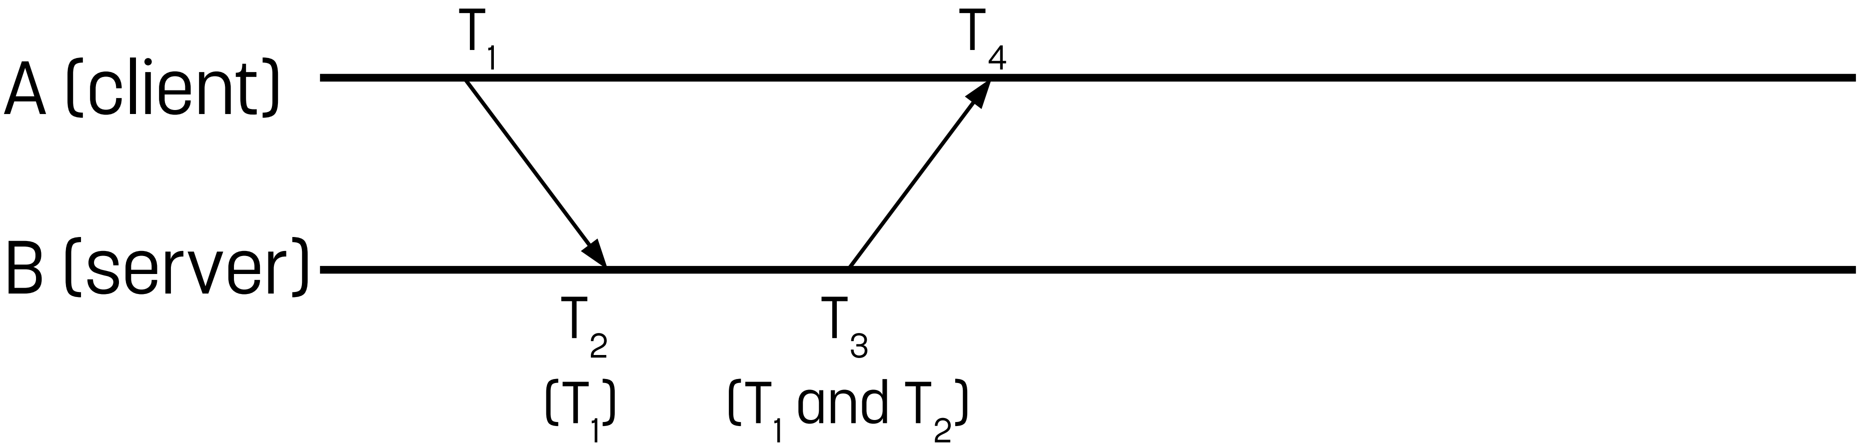

Network Time Protocol NTP

- A (client) reads current time T_4 on arrival and computes offset and round-trip delay relative to B out of $T_1 \dots T_4$

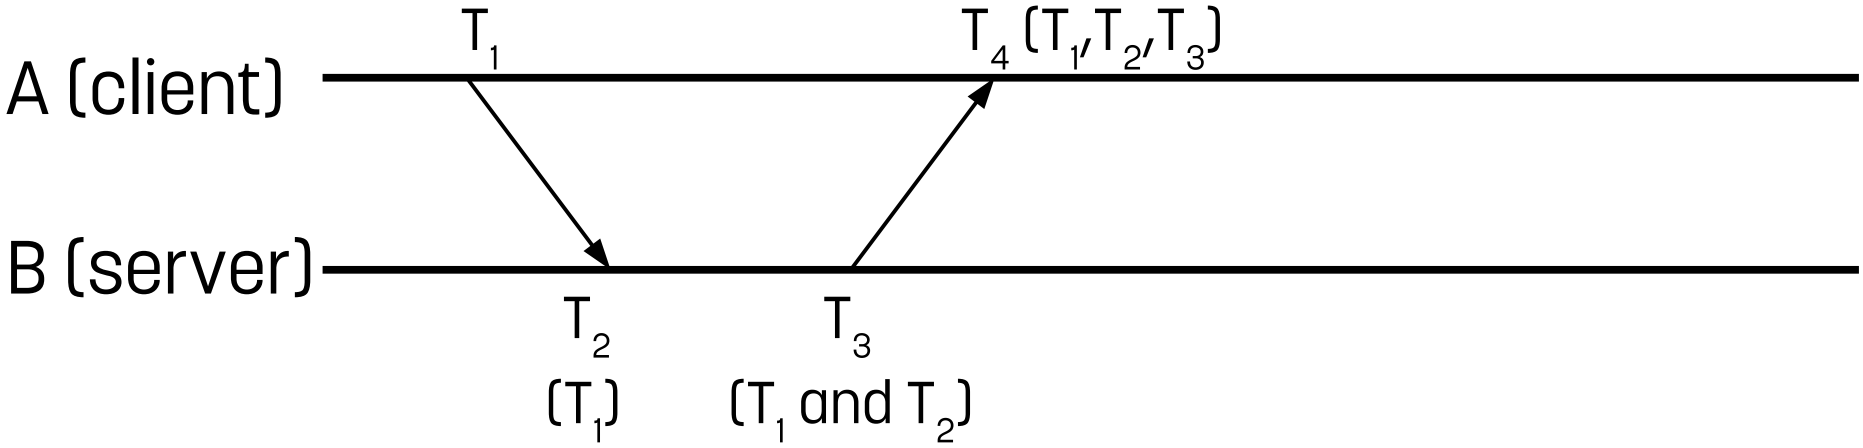

Time sources in a computer

- RTC → Real Time Clock (crystal stabilized), battery powered
- Software clock (or system clock, kernel clock), only running, when the system is up
- Network time (via NTP), read the time periodically from a network time server, and continuously adjust the rate of the system clock so that the time data always match.

How NTP disciplines system clock

- **Small** time offsets (**step threshold**, per default 128 ms)
  - Adjust time smoothly, max 500 µs/s (1.8s/hour)

- **Large** time offsets
  - Adjusts time at NTP server start
  - If disciplined system detects timer offset that exceeds **step threshold** → stepout interval (300s)
  - If time offset > **panic threshold** (1000s) the NTP server terminates → set clock manually

NTP measurement

QA-Farm setup:

1. Standard distribution NTP setup
2. Local NTP server
3. Local NTP server peer to peer, optimized

NTP measurement: equipment

System Board: AAEON/UP-WHL01

- CPU x86 Intel Core i3-8145UE @2200 MHz
- Integrated GPU UHD 620 @300 MHz
- Kernel: 5.10.27-rt36 #4 SMP PREEMPT_RT
- Architecture: Whiskey Lake

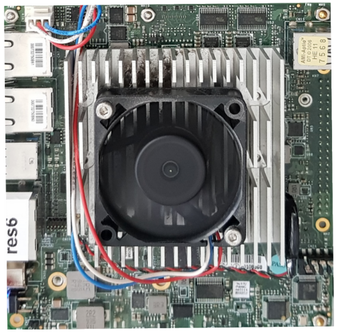

NTP measurement: setup 1

- Connect to the standard local NTP server

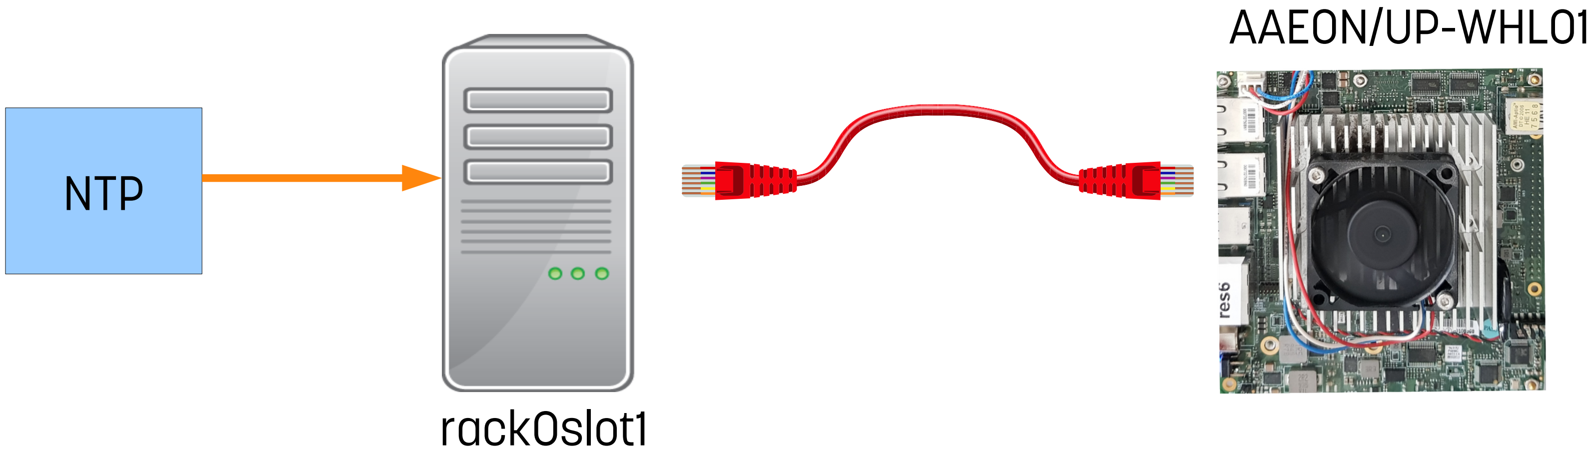

NTP measurement: setup 1

- Connect to the standard local NTP server

```
ntpq> pe
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*funky.f5s.de    131.188.3.221    2 u   74 1024  377    7.174    0.025   3.639
+185.168.228.59  131.188.3.222    2 u 1377 1024  372    9.573    0.055   1.207
+electrode.felix 85.10.240.253    3 u  657 1024  377    6.882   -2.468   1.593
-time.cloudflare 10.71.1.91       3 u  344 1024  377    2.905   -3.134   1.161
```

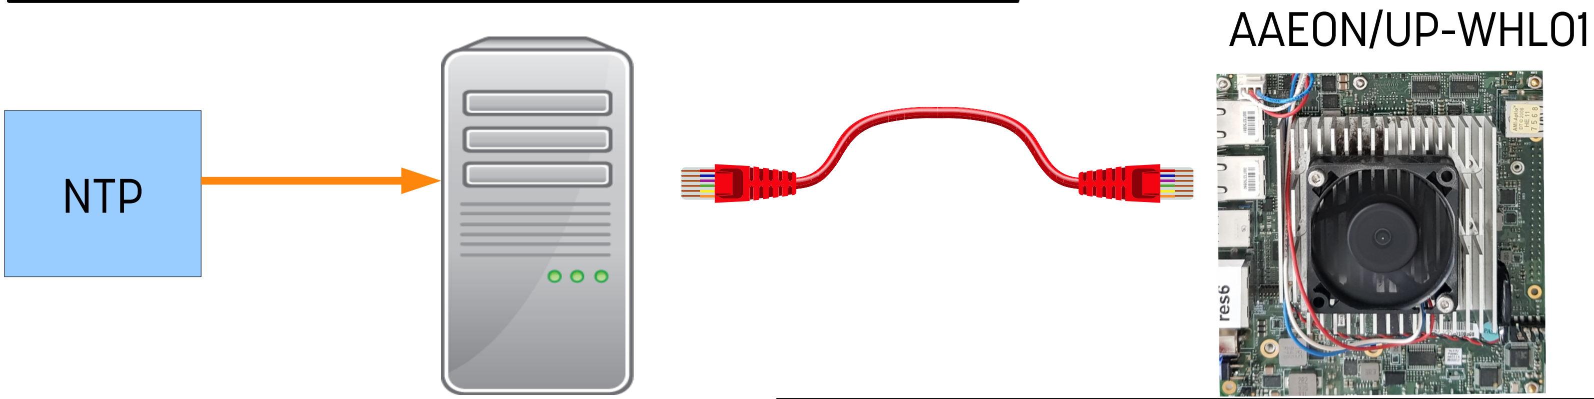

rack0slot1

```
ntpq> pe
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*rack0slot1.osad 144.76.59.37     3 u 1005 1024  377    0.422   -1.100   0.688
```

NTP measurement results: setup 1

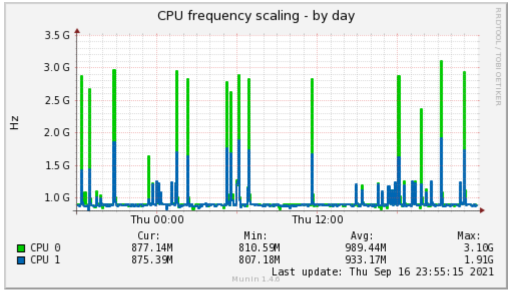

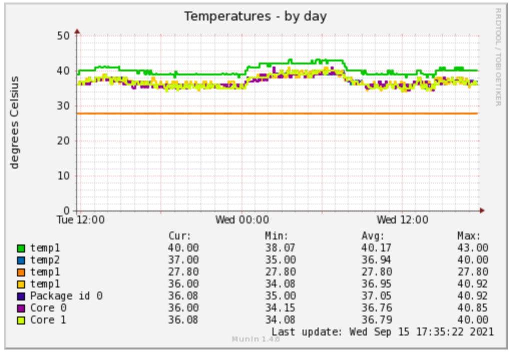

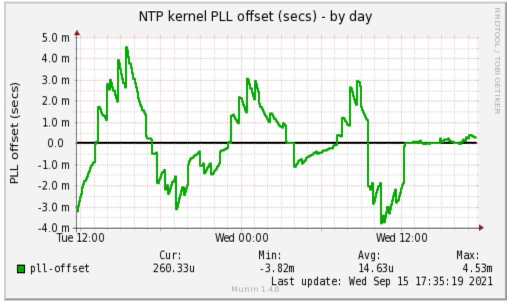

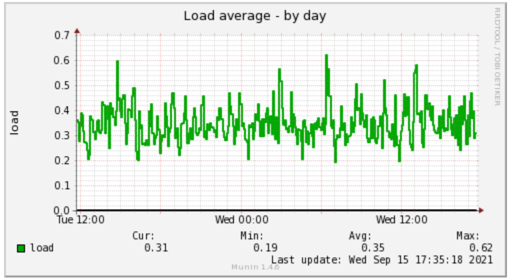

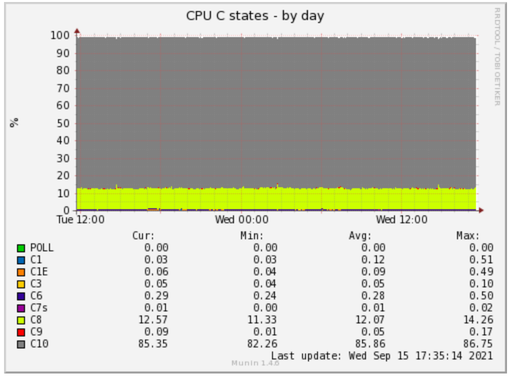

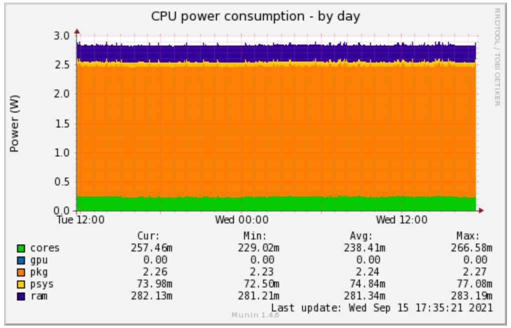

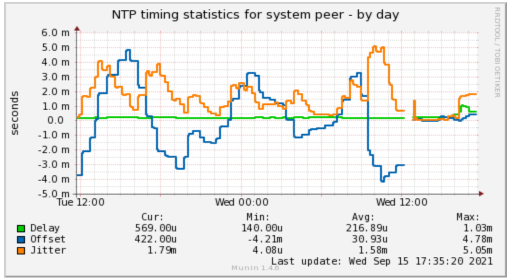

NTP measurement: setup 2

- Connect to local NTP server

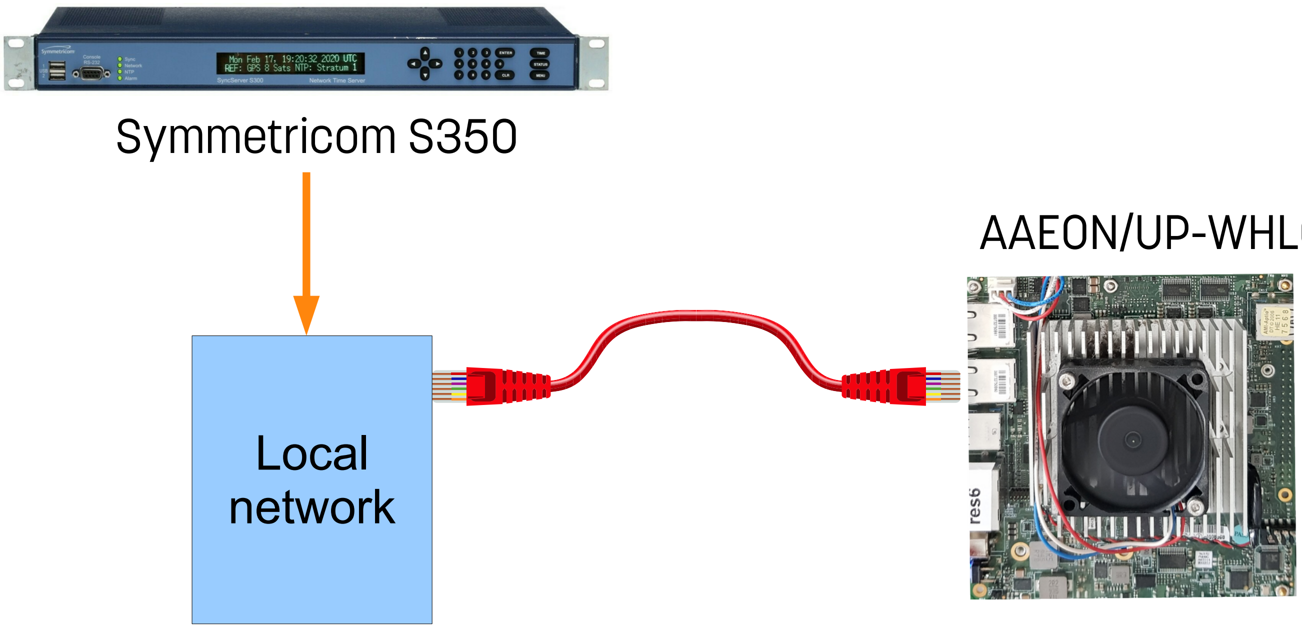

NTP measurement: setup 2

- Connect to local NTP server

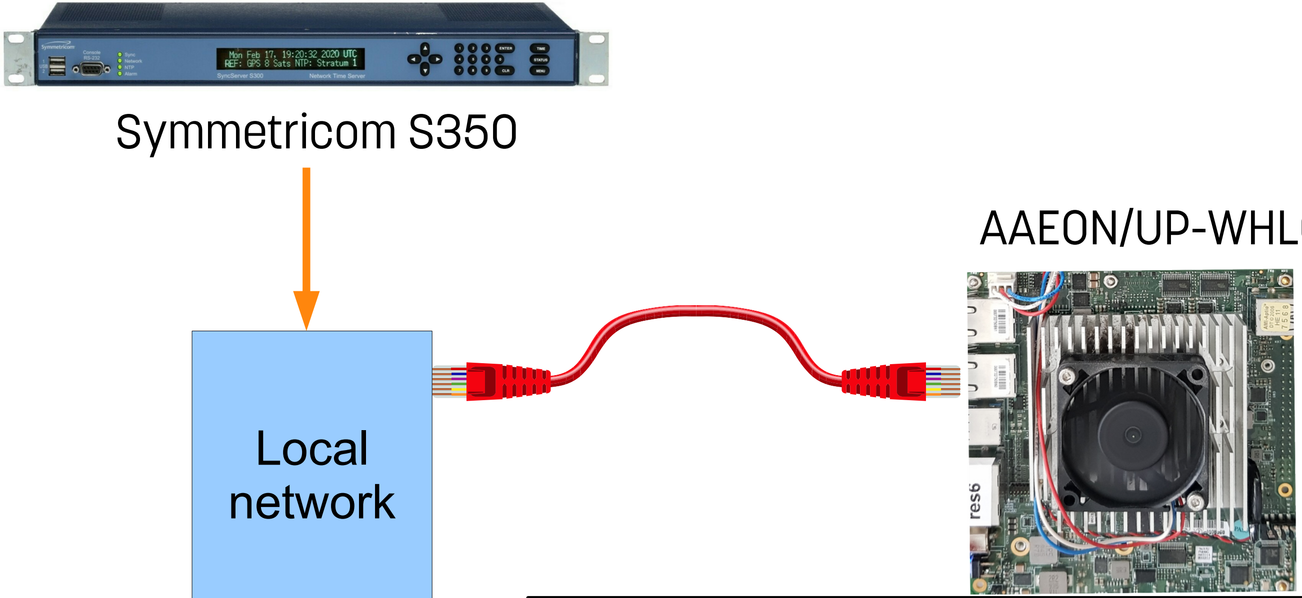


```
ntpq> pe
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*ntp.osadl.org   .PPS.            1 u  689 1024  377    1.001   -1.550   0.793
```

NTP measurement results: setup 2

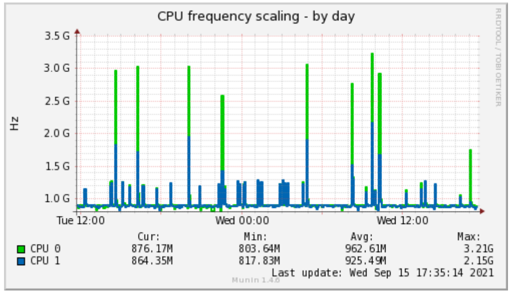

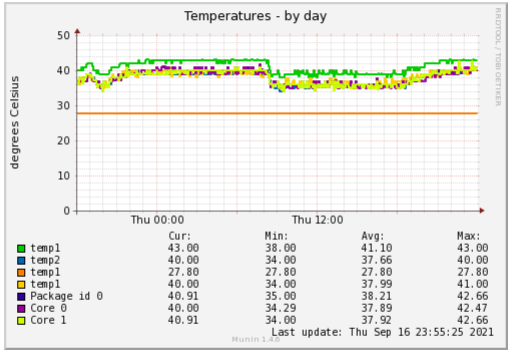

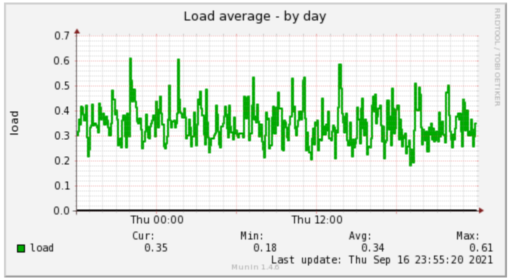

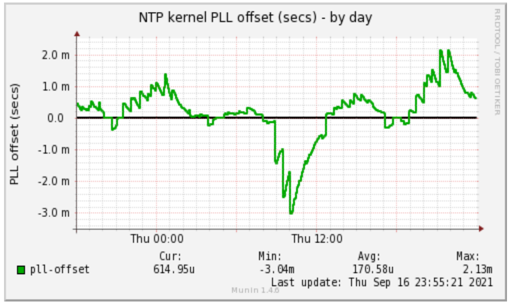

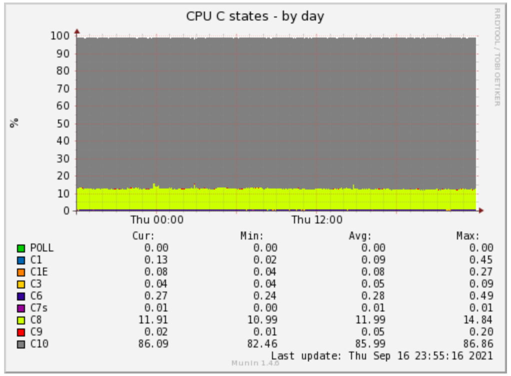

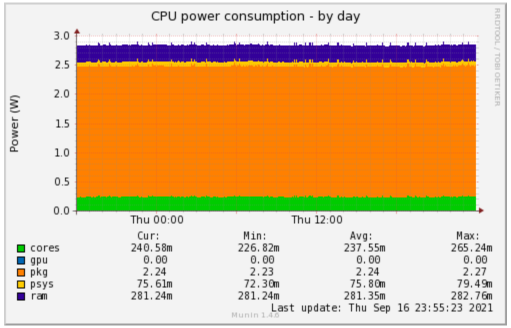

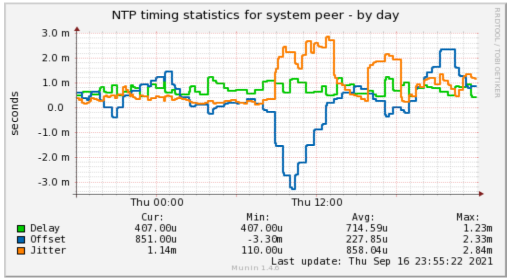

NTP measurement: setup 3

- Connect to local NTP server P2P
  - dedicated interface
  - network irq priority 90
  - network irq affinity 1
  - taskset -c 0 chrt -f 91 /usr/sbin/ntpd

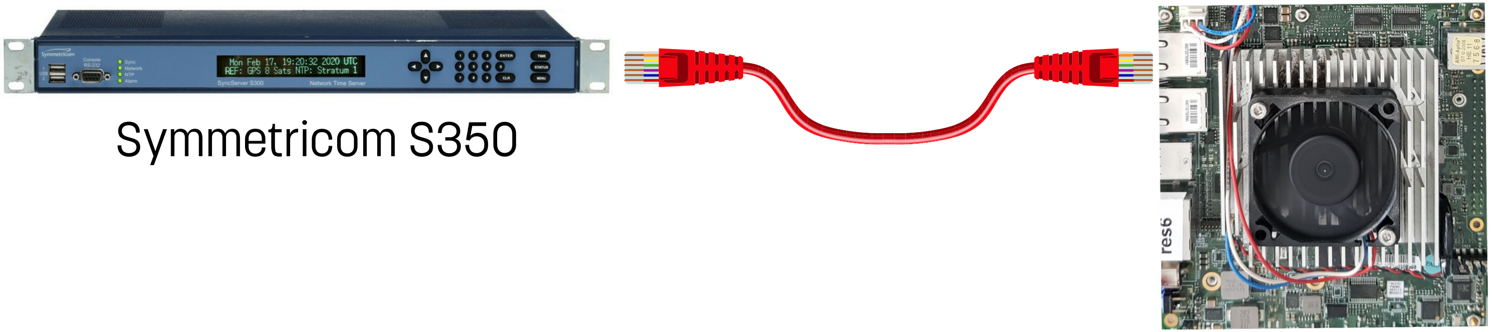

NTP measurement: setup 3

- Connect to local NTP server P2P
  - dedicated interface
  - network irq priority 90
  - network irq affinity 1
  - taskset -c 0 chrt -f 91 /usr/sbin/ntpd

Symmetricom S350

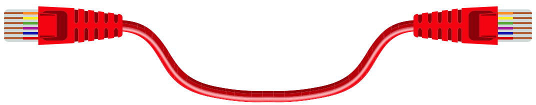

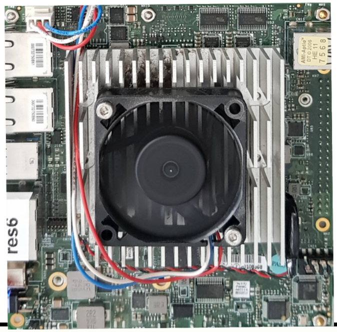

AAEON/UP-WHL01

```
ntpq> pe
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*ntp.osadl.org   .PPS.            1 u  689 1024  377    1.001   -1.550   0.793
```

NTP measurement results: setup 3

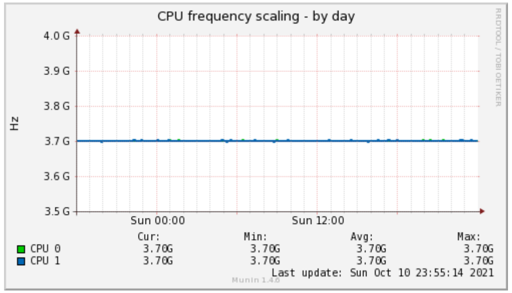

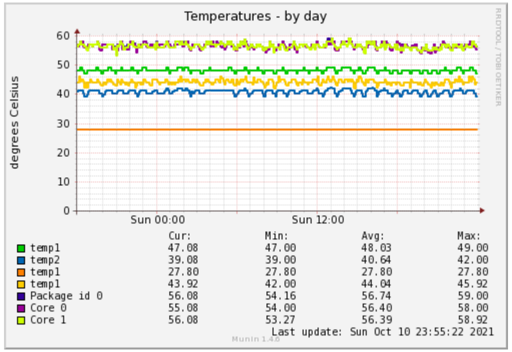

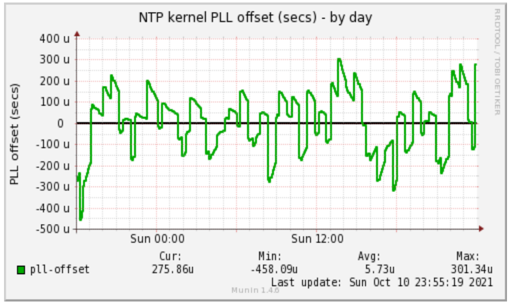

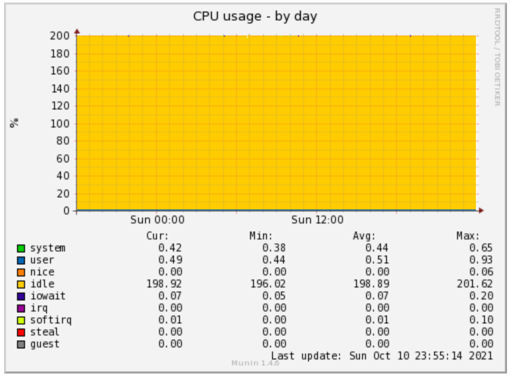

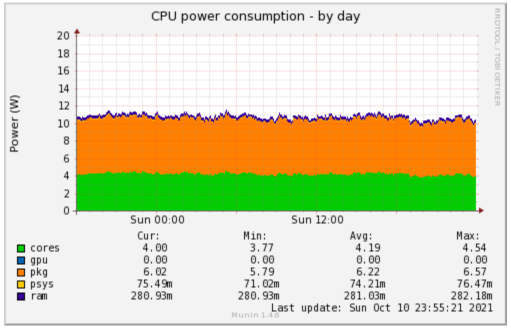

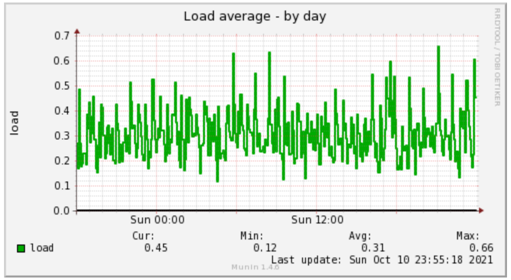

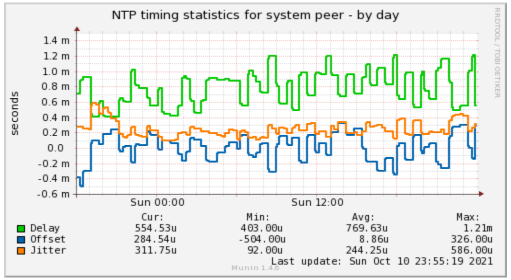

Summary of measurement results

	Setup_1	Setup_2	Setup_3
NTP_Source	external NTP-Server	local NTP-Server	local NTP-Server
Optimization	none	none	rt/p2p-connection
NTP delay	-0.1/1.0/0.2	0.4/1.2/0.7	0.4/1.2/0.77
NTP offset	-4.2/4.7/0.03	-3.3/2.3/0.2	-0.5/0.3/0.009
NTP jitter	-0.004/5/1.6	0.1/2.8/0.86	0.09/0.5/0.24
Pll offset	-3.8/4.5/0.015	-3.0/2.1/0.17	-0.46/0.3/0.006

all values in ms

Summary of measurement results

	Setup_1	Setup_2	Setup_3
NTP_Source	external NTP-Server	local NTP-Server	local NTP-Server
Optimization	none	none	rt/p2p-connection
NTP delay	-0.1/1.0/0.2	0.4/1.2/0.7	0.4/1.2/0.77
NTP offset	-4.2/4.7/0.03	-3.3/2.3/0.2	-0.5/0.3/0.009
NTP jitter	-0.004/5/1.6	0.1/2.8/0.86	0.09/0.5/0.24
Pll offset	**-3.8/4.5/0.015**	**-3.0/2.1/0.17**	**-0.46/0.3/0.006**

all values in ms

Conclusion

- NTP is the established technology to synchronize networks and devices
- Global availability
- Optimization can be achieved through
  - the use of a NTP server in LAN
  - Peer to peer connection
  - Linux-RT optimizations (taskset, irq-prio/-affinity