Open
Source
Automation
Development
Lab eG

OSADL eG · Im Neuenheimer Feld 583 · 69120 Heidelberg · www.osadl.org · info@osadl.org

The OSADL QA Farm

Rationale

Embedded systems must fulfill a number of indispensable requirements for the use in industrial products. In addition to general reliability and stability, case-specific functionality must be provided without exception. Therefore, the OSADL QA Farm was established, continuously monitoring approximately 180 different industrial embedded and other computer systems. The resulting data are primarily observed by companies who have provided the systems, and the data are used to optimize and fine-tune the system hardware and software. Furthermore, a large part of the data is made available to the public. Last but not least, these data, which are collected and documented under controlled circumstances, may also serve to create testing scenarios to answer scientific questions on the behavior and suitability of hardware and software components of embedded systems used in industry.

Realization

The OSADL QA Farm consists of a number of open test racks each of which provides eight slots for test tablets. The racks can be located anywhere in the world; the only requirements are power supply and Internet access. At this time, OSADL test racks are located at two different test centers in Germany. Every test rack is equipped with a number of centralized control and communication units. In particular, these are

- 10/100/1000 Mb/s network switch with port mirroring
- remote control of power distribution units with individual power metering
- serial-to-network adapter
- serial-to-USB adapter
- USB hub
- KVM adapter with network access
- server for cross development and as peer for generating network load

The systems under test are equipped with DIN rails – one per each of the eight slots. Each tablet provides a 230 V AC wall plug providing electrical power as well as two RJ45 sockets used for network access and an RS232 system console interface. They are connected to the respective counterparts of the rack control systems. An optional VGA graphics connector as well as a USB or PS/2 keyboard and mouse connector may be added if needed. If all lines are connected, a particular test system is ready for operation within the rack.

All test racks are connected to a central processing server via VPN channels for maintaining, storing and visualizing the collected data (see Figure 1). This server is equipped with appropriate software to allow inspection of the data from all over the world using a standard web browser. Alarm thresholds are defined for many variables and categorized into “warning” and “critical” level. In case an alarm threshold is exceeded, a previously assigned contact person will be notified by an escalation system. Email, text messaging, fax and voice messaging may be selected as means of communication. In addition, the particular variables that are in “warning” or “critical” state are highlighted in the web interface in yellow or red color, respectively.

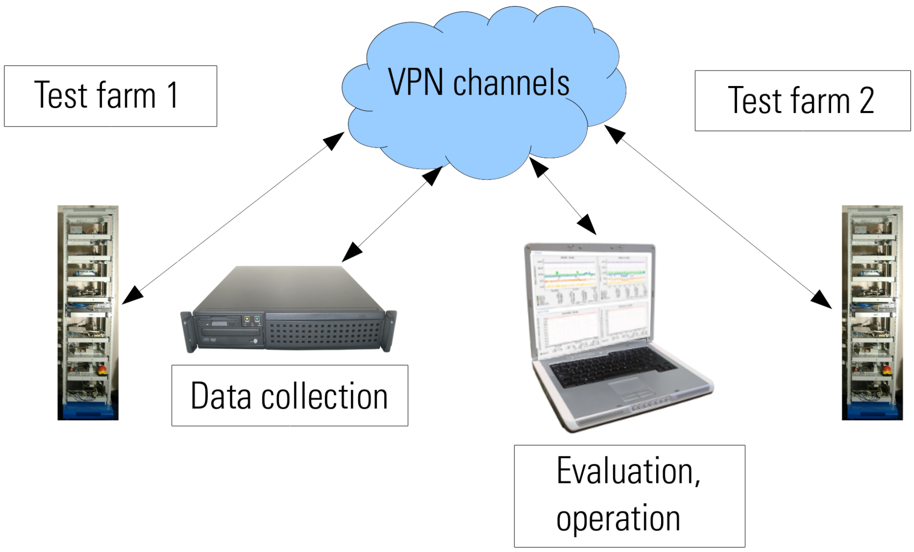

Figure 1: Communication among the various OSADL QA Farm components

Hardware testing

The embedded systems under test either originate from OSADL member companies that are interested in the stability and reliability data of their systems, or are provided by OSADL for the development of the Linux kernel. The latter are updated continuously to the latest kernel version, while the member-provided systems, depending on the arrangement, may either also be updated regularly or run a stable kernel. The former approach assures the companies that an update to a new kernel version is possible anytime without requiring substantial development and testing.

According to the usual deployment in industrial embedded systems, the systems are equipped with processors of the ARM, MIPS, PowerPC and x86 family. Some of the processors have just been launched on the market, while some others may have a 20-year old design. The latter is important because industrial systems could– and particularly medical and military system – could have such a long life cycle, and it must be assured that it is possible to install the latest Linux kernel on them should this be required.

Among the various processors tested in the OSADL QA Farm particular emphasis is put on the availability of a wide range of clock frequency, memory size and processor topology whenever possible. For instance, the clock frequencies range from 133 to 4,000 MHz, and the memory size ranges from as little as 26 MByte up to 128 GByte. Beside single-core processors of former generations, multi-core processors with up to 32 cores as well as mixed multi-core systems with several sockets and nodes are used. In addition, when selecting chip sets and peripheral devices care was taken that a wide variety of different controllers of different manufacturers are under test. An overview of all systems installed in the QA Farm can be found at `https://www.osadl.org/?id=879`.

Monitored variables

The variables measured at the OSADL QA Farm can be divided into the groups benchmark, disk, network, NFS, processes, real-time system, email, sensors, time synchronization, system and virtualization. Table 1 shows an overview of the measured variables. The current plots of the measured variables for all QA Farm systems can be found at `http://munin:munin@munin.osadl.org/munin/`.

Benchmarks
GL benchmark gltestperf
UnixBench (multi-core)
UnixBench (single-core)
UnixBench 2D graphics performance

Disk
Disk IOs per device
Disk latency per device
Disk throughput per device
Disk usage in percent
Disk utilization per device
File system mount-scheduled checks
File system time-scheduled checks
Filesystem usage (in bytes)
Inode usage in percent
IO Service time
IOstat
S.M.A.R.T values of every drive

Network
eth0 errors
eth0 traffic
Firewall Throughput
HTTP loadtime of a page
Netstat

NFS
NFS Client
NFSv4 Client

Processes
Fork rate
Number of threads
Processes
Processes priority
Vmstat

Real-time system
5-min max. timer and wakeup latency
5-min max. timer offsets
5-min max. wakeup latency
RT Features

Email
Sendmail email traffic
Sendmail email volumes
Sendmail queued mails

Sensors
Fans
HDD temperature
Power consumption
Temperatures

Time synchronization
NTP kernel PLL estimated error (secs)
NTP kernel PLL frequency (ppm + 0)
NTP kernel PLL offset (secs)
NTP states
NTP timing statistics for system peer

System
Available entropy
C states
CPU frequency
CPU usage
File table usage
Individual interrupts
Inode table usage
Interrupts and context switches
Kernel version
Load average
Logged in users
Memory usage
Split memory usage
Application memory usage
Swap in/out
Uptime

Virtualization
Virtual domain block device I/O
Virtual domain CPU time
Virtual domain memory usage
Virtual domain network I/O

Table 1: Monitored variables at the OSADL QA Farm (selection)

Variables with special importance for the use of embedded systems in industry

Variables that are especially important for the use of industrial embedded systems are those that are related to a system's response to asynchronous external and internal events (real-time capability) as well as to the temperature profile and power consumption for different load scenarios. Additionally, the particular performance characteristics of CPU, FPU and GPU must be determined and registered for comparison purposes. Last but not least, it is necessary to record the version and release numbers of the Linux kernel in order to be able to verify whether a kernel upgrade introduced a regression of one or several of the measured variables.

OSADL QA Farm Schedule for the evaluation of real-time capabilities

The PREEMPT_RT patch turns Linux from a general purpose operating system into a real-time operating system with a POSIX programming interface. The real-time capability of an operating system can be measured both by internal and external means. All systems installed in the OSADL QA Farm undergo a 12-hour test cycle in which they are measured repeatedly under idle conditions and under load. In the first two hours of the six-hour load period, timer interrupts are triggered at a frequency of 5 kHz as the only load. During the following four hours, the systems are additionally exposed to a specific load with memory allocations, network accesses and file system calls. The OSADL QA Farm schedule is shown in Figure 2.

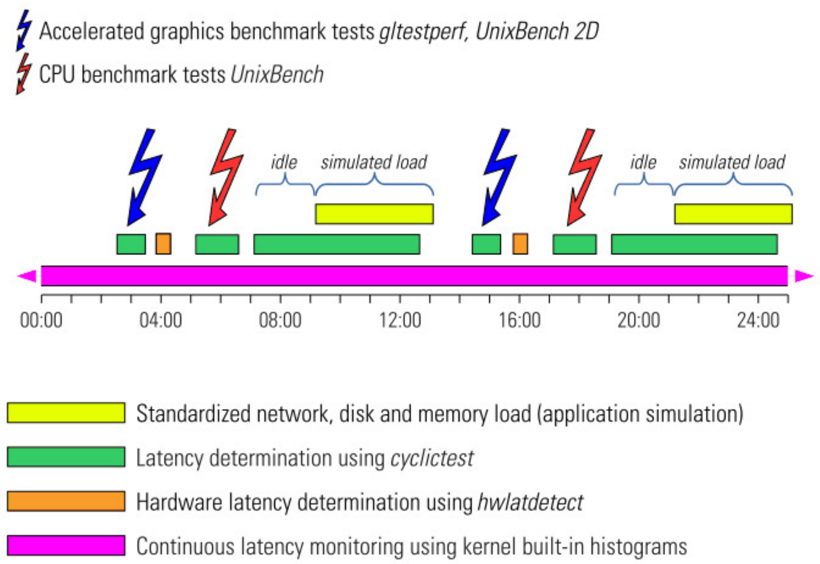

Figure 2: Schedule of the OSADL QA Farm test procedures

Continuous internal real-time latency monitoring is based on built-in kernel histograms that were part of the Linux kernel trace subsystem, but are now maintained by OSADL. The patches required to implement this additional functionality into the Linux kernel can be downloaded from the OSADL website at `https://www.osadl.org/?id=2943`. With this method the Linux kernel records the current time at the beginning of interrupt processing and immediately before the user space regains control. The difference between these two time stamps is recorded in a histogram of which the highest value within a five-minute interval is read out using the *debugfs* virtual file system. This total system latency consists of three partial latencies that are also recorded individually (Figure 3): the offset between the intended and the actual expiration of the timer (*missed_timer_offset*), the scheduler time until enqueuing the user space program (*wakeup*) and the time of the context switch until returning from system call (*switchtime*).

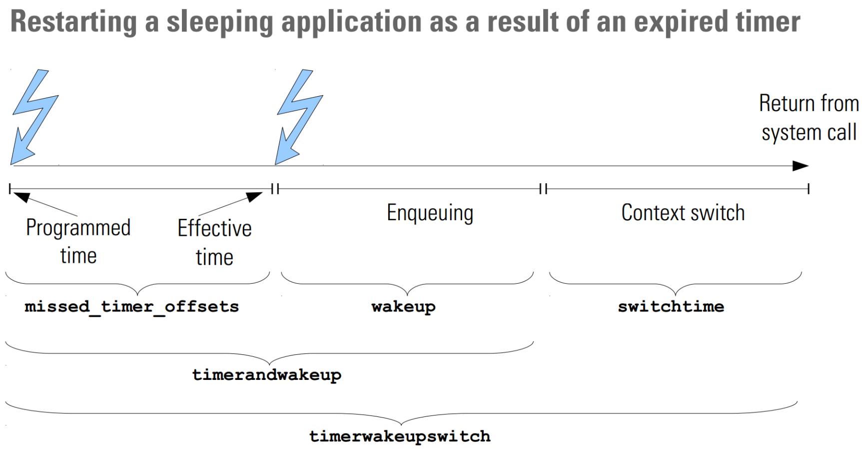

Figure 3: Latencies during a wakeup procedure

In addition, sums of partial latencies are calculated and saved, such as the total sum of timer latency and scheduler latency. This sum corresponds – to a very large extent – to the so-called preemption latency and represents an important measure of the real-time capability of a given system. All systems in the OSADL QA Farm are continuously monitored this way by recording the latency maximum of consecutive five-minute measurement periods. Last but not least, characteristic data (name, ID and priority) of the process with the maximum latency of a five-minute interval ("victim") are determined In addition, such data are also saved of the process from which the switch was made ("culprit")). This culprit-victim analysis is done to obtain information about why a longer latency occurred in an individual case so that these issues may be resolved. For the external latency measurement the test program *cyclictest* is used to determine a system's worst case latency per measurement period under standard, graphics and CPU load. It uses cyclic timer interrupts with a frequency of 5 kHz to simulate external events and measure the delay between the programmed and the effective wake-up time directly in user space. The results are presented in a so-called latency plot: a histogram with a linear x axis and a logarithmic y axis to visualize even very low sample sizes. This is important, because the longest ever measured wake-up latency may occurs only a few times or even only once. However, this value is the most important result of the measurement as it classifies the real-time capability of the system. Latency plots for all systems installed in the OSADL QA Farm are published at `https://www.osadl.org/?id=895`. Figure 4 shows the latency plot for an Intel Celeron G1620 CPU with 2.7 GHz running a Linux kernel version 5.4.17-rt9 as an example (the current latency plot can be found at `https://www.osadl.org/?id=1967`). This histogram shows the result of one five and a half hour measurement under normal load with 100 million cycles per CPU core.

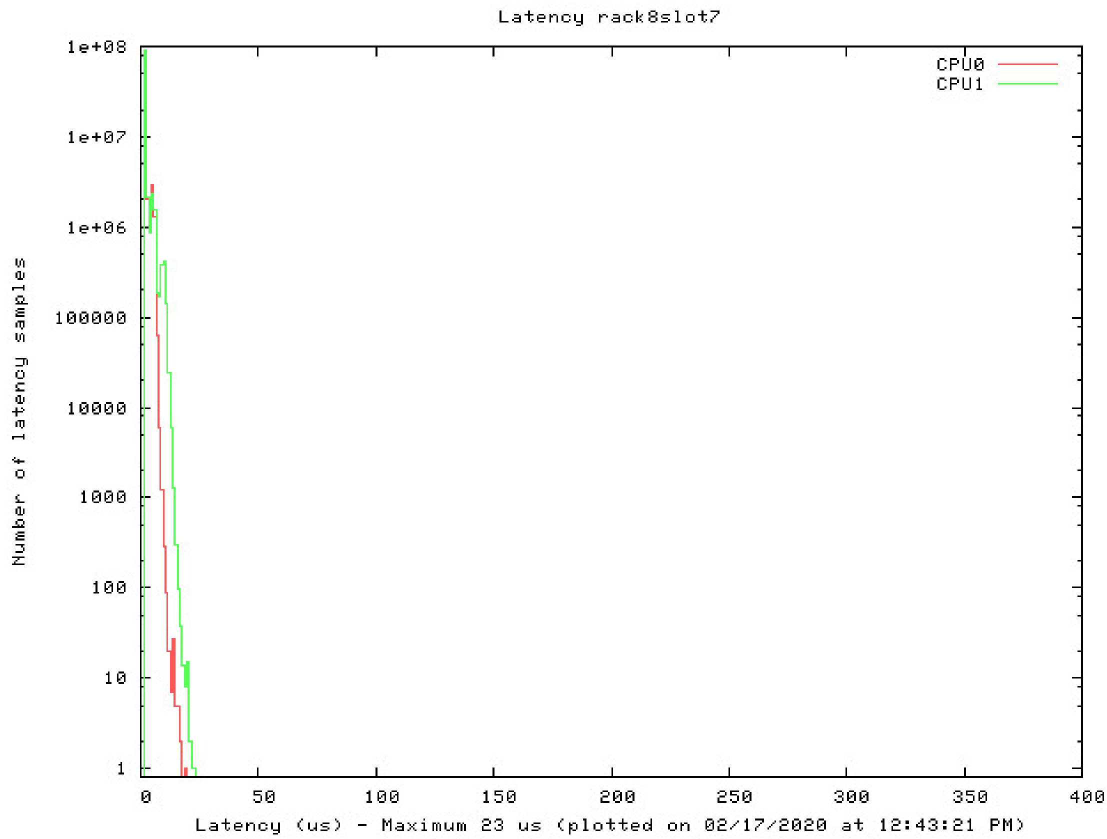

Figure 4: Latency plot of an Intel Celeron G1620 CPU @2.7 GHz, Linux kernel version 5.4.17-rt9

In order to achieve high statistical confidence, the individual histograms of a large number of recordings are combined in a joint evaluation. For the selected CPU such a long-term latency plot is shown in Figure 5 (the current plot is available at https://www.osadl.org/?id=1968). The time axis("Repetitions") of these long-term plots runs backwards, i.e. they are to be read from back to front with the newest measurements located at repetition 0). As can be seen in Figure 5, sporadic higher latencies were recorded at the beginning. After removing the causes for these latencies, the system has not been showing any further high latency spikes for over three years indicating that this system may rightfully be called deterministic. In addition, these long-term measurements prove to be important in the analysis of a system's real-time behavior with regard to changes in the run-time conditions which might lead to improved or worsened latencies. This is exemplified in Figure 5 between repetitions 1400 and 1200. During this time, the maximal latency was recognizably higher. As the version of the Linux kernel (major, patch level, sub-level and RT release) is recorded as well, the improvement of the latency at repetition 1200 could have been linked to a kernel version update.

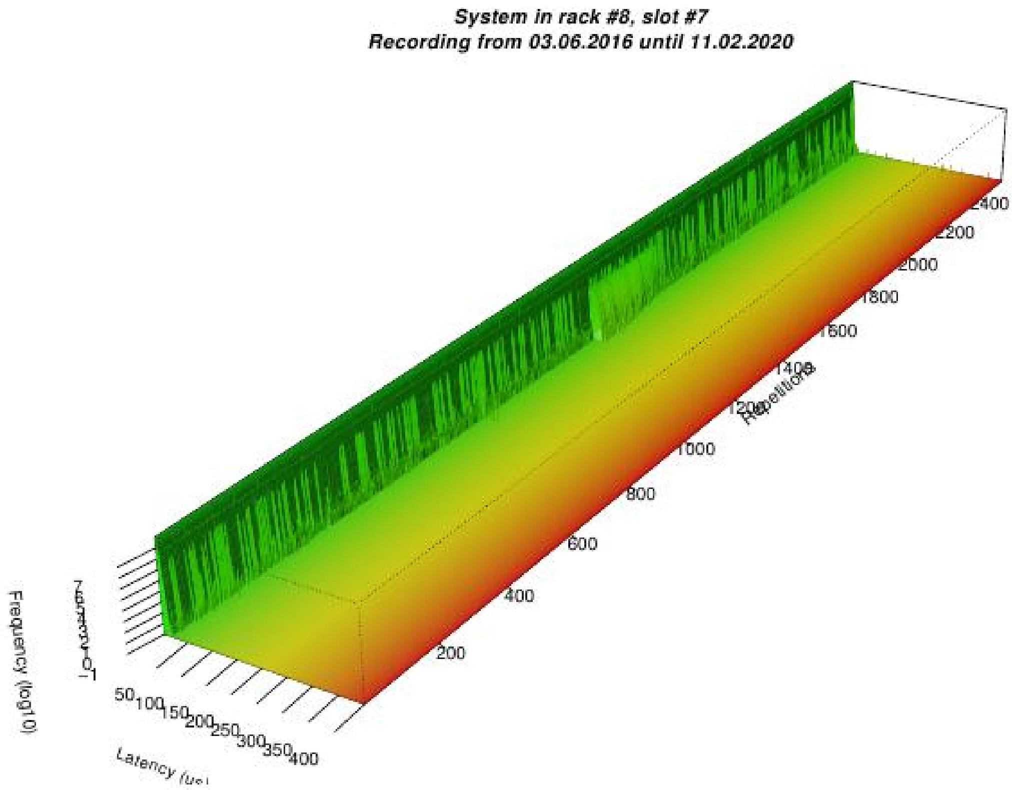

Figure 5: Long-term latency plot of an Intel Celeron G1620 CPU @2.7 GHz, Linux kernel version 5.4.17-rt9 (to be read from back to front, i.e. the newest measurements are at repetition 0)

Conclusion

OSADL's QA Farm collects and provides data on stability, reliability and real-time behavior of a large and varied number of processors and computer systems. They are used to

- provide selection criteria for automation hardware,
- generate availability and stability data of individual systems,
- study a system's change of behavior due to new kernel versions,
- generate reliable data for certification purpose (e.g. real-time, safety),
- study real-time behavior of special test setups (e.g. peer-to-peer UDP Duplex-link, PTP/TSN, network load, virtualization).

References

Emde, C.: Quality assessment and assurance of embedded systems.
`https://www.osadl.org/uploads/media/OSADL-QA-Farm-Description-ENG.pdf`

Emde, C.; Gleixner, T.: Quality assessment of real-time Linux boards & solutions + ECE, Sept 2011
`https://www.osadl.org/uploads/media/ECE-2011-09.pdf`

OSADL Im Neuenheimer Feld 583 • 69120 Heidelberg, Germany
Phone: +49 (0) 6221 98504 0 • Fax: +49 (0) 6221 98504 80 • info@osadl.org • http://www.osadl.org

As of February 2020