Real-Time OS Overview

Ing. Michal Sojka, Ph.D.

ČVUT v Praze,

michal.sojka (at) cvut.cz

October 30, 2024

Real-Time OS Overview

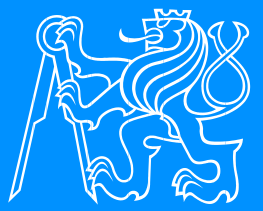

- There exist hundreds of RT operating systems...
  - more in the past, less today
- VxWorks (WindRiver)
  - Commercial OS
  - Multiprocessor support, optional memory protection, POSIX API, WindAPI, Eclipse-based development environment + tools
  - Flew to Mars
- RTEMS
  - Open-source
  - Multiprocessor support, POSIX API, RTEMS API
  - Active development community
  - Good documentation

Real-Time OS Overview

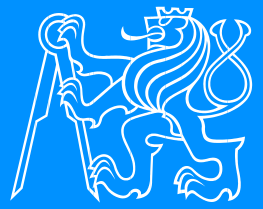

- Zephyr (https://zephyrproject.org/)
  - Modern, open-source RTOS, managed by Linux foundation
    - supported by many hardware vendors
  - Introduced in 2016
  - No memory protection by default (ala VxWorks DKM), user mode with protection (ala VxWorks RTP) available for some hardware.
  - One goal is to provide safety certification as a commercial service
- Apache NuttX
  - Open-source

Real-Time OS Overview

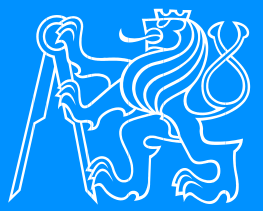

- QNX – commercial, microkernel-based
- Dual-kernel OS (real-time addons to generic OS)
  - RTAI – open-source, Linux runs as a task with lowest priority
  - ~~RT Linux – ditto, patent problem (now Wind River/Intel)~~
  - Ardence/RTX – real-time addon to Windows
- PikeOS (SysGo) – commercial hypervisor, safety critical applications
- FreeRTOS – for small micro-controllers, single address space
- Windows CE – Real-Time OS from Microsoft, Win32 API
- Azure RTOS (ThreadX) – Small RTOS from Microsoft, MS cloud integration
- ~~eCos – professional and open-source version. Development is not too open. Interesting HAL. Offered as an alternative OS for Siemens's PLCs.~~

Linux and real-time

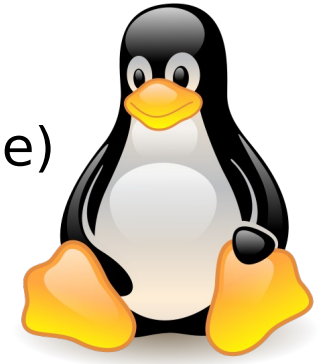

- Standard kernel did not have RT properties
  - Kernels 2.4.x were not preemptive (2.6+ is preemptive)
  - Many companies (MontaVista, TimeSys) tried to turn Linux into RT OS
  - They didn't work with the community – they usually offered old versions
  - Later they hired Ingo Molnar a Thomas Gleixner, to make Linux RT capable and push the patches to the mainline version.
- Around 2004 – birth of PREEMPT_RT patch
  - In the beginning it was distributed as a single huge (unmaintainable) patch to prevent the companies selling untested early development versions.

PREEMPT_RT patch

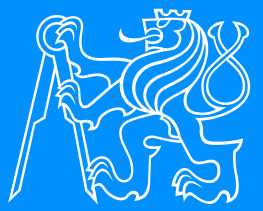

- **Problem**:
  - Even in Linux 2.6.x it was not possible to preempt large amount of kernel code
    - Interrupt handlers, SMP critical sections (spinlock protected), …
    - Those disable interrupts – response time to external events is unnecessary long
- **Solution: rt-preempt patch**
  - Most spinlocks are replaced by mutexes (critical sections become preemptive)
  - Implements priority inheritance (prevents priority inversion)
  - IRQ handlers and softirqs are converted to threads (become preemptive)
  - Timers were reworked to provide high resolution

PREEMPT_RT current status

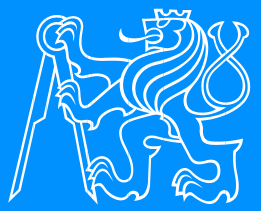

- **PREEMPT_RT was merged** into Linux 6.12 (release in a few weeks)!
- In 2015, the preempt-rt project was funded by Google (via Linux Foundation)
- Writing RT applications for Linux is simple:
  - https://wiki.linuxfoundation.org/realtime/documentation/howto/applications/application_base
  - Configure scheduler (SCHED_FIFO + inheritance)
  - Use mlockall(), and “page-in” your stack.
- rt-tests package contains useful tools:
  - cyclictest – allows measuring worst-case scheduling latencies (similar to our lab task “scheduler latency bechmark”).

Real-Time preempt configuration

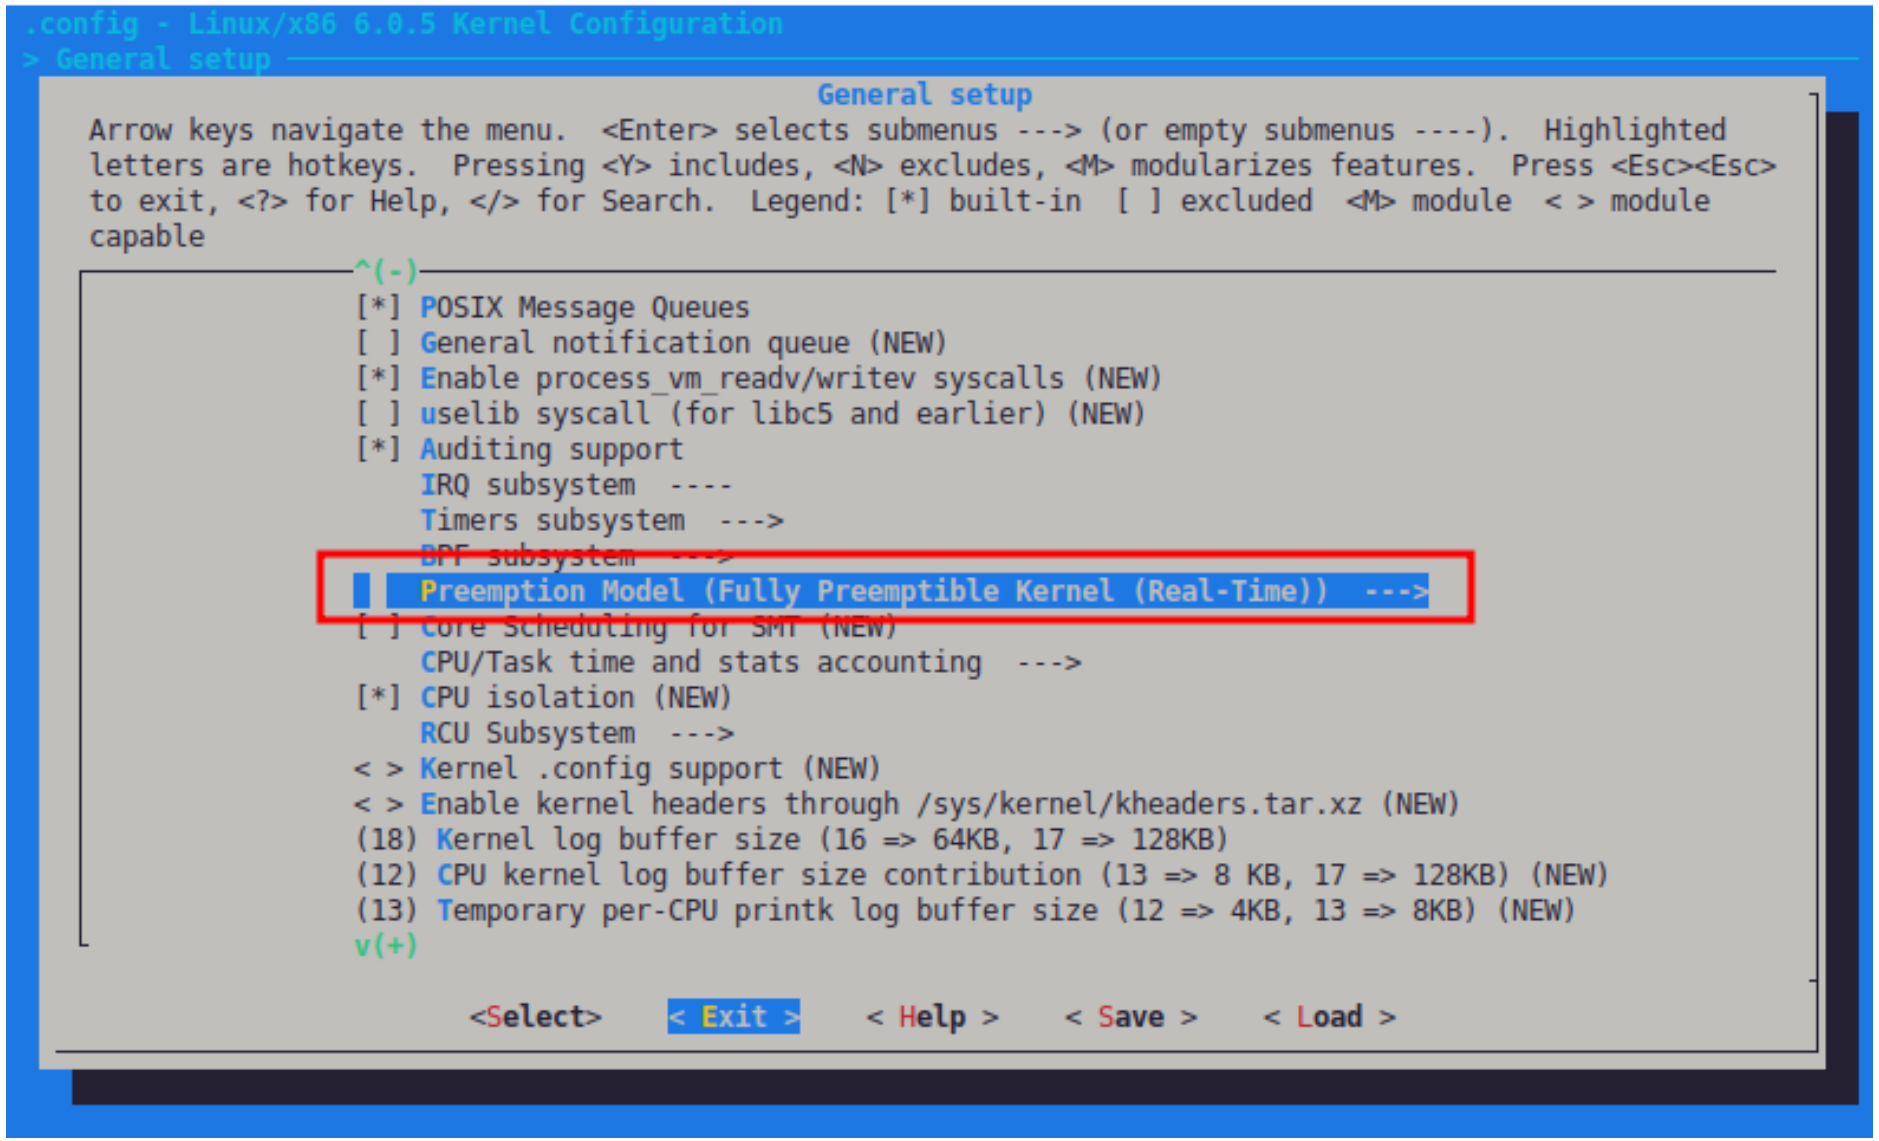

Other Linux-related solutions

- **SCHED_DEADLINE** – Linux EDF scheduler (in mainline)
  - Implements Constant Bandwidth Server (CBS) that provides temporal task isolation (protects WCET from overruns)
  - https://github.com/jlelli/sched-deadline
  - http://www.evidence.eu.com/sched_deadline.html
- **JailHouse** – partitioning hypervisor for Linux
  - Small hypervisor for real-time safety critical applications
  - https://github.com/siemens/jailhouse
- **Xenomai** – hard real-time from Linux user space
  - Adeos – IRQ and syscall virtualization

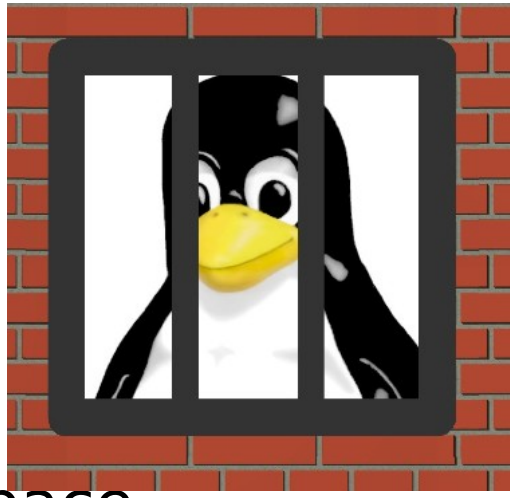

Linux rt-preempt – links

- Real-Time Linux Wiki:
  https://wiki.linuxfoundation.org/realtime/start
  (older site: http://rt.wiki.kernel.org)
- Mailing list: linux-rt-users@vger.kernel.org
- OSADL: http://www.osadl.org
- „Latest stable“ Real-Time Linux
  http://www.osadl.org/Latest-Stable.latest-stable-realtime-linux.0.html
- Linux Weekly Newsletter: http://lwn.net/

Microkernels (TU Dresden)

- L4Re (https://l4re.org/)
  - L4 microkernel + run-time environment
  - Currently used in VW cars etc.
- NOVA Microhypervisor (http://hypervisor.org)
  - Small Trusted Computing Base
- Genode OS framework
  - Unified user space for microkernels
  - Cool!
  - UNIX emulation...