



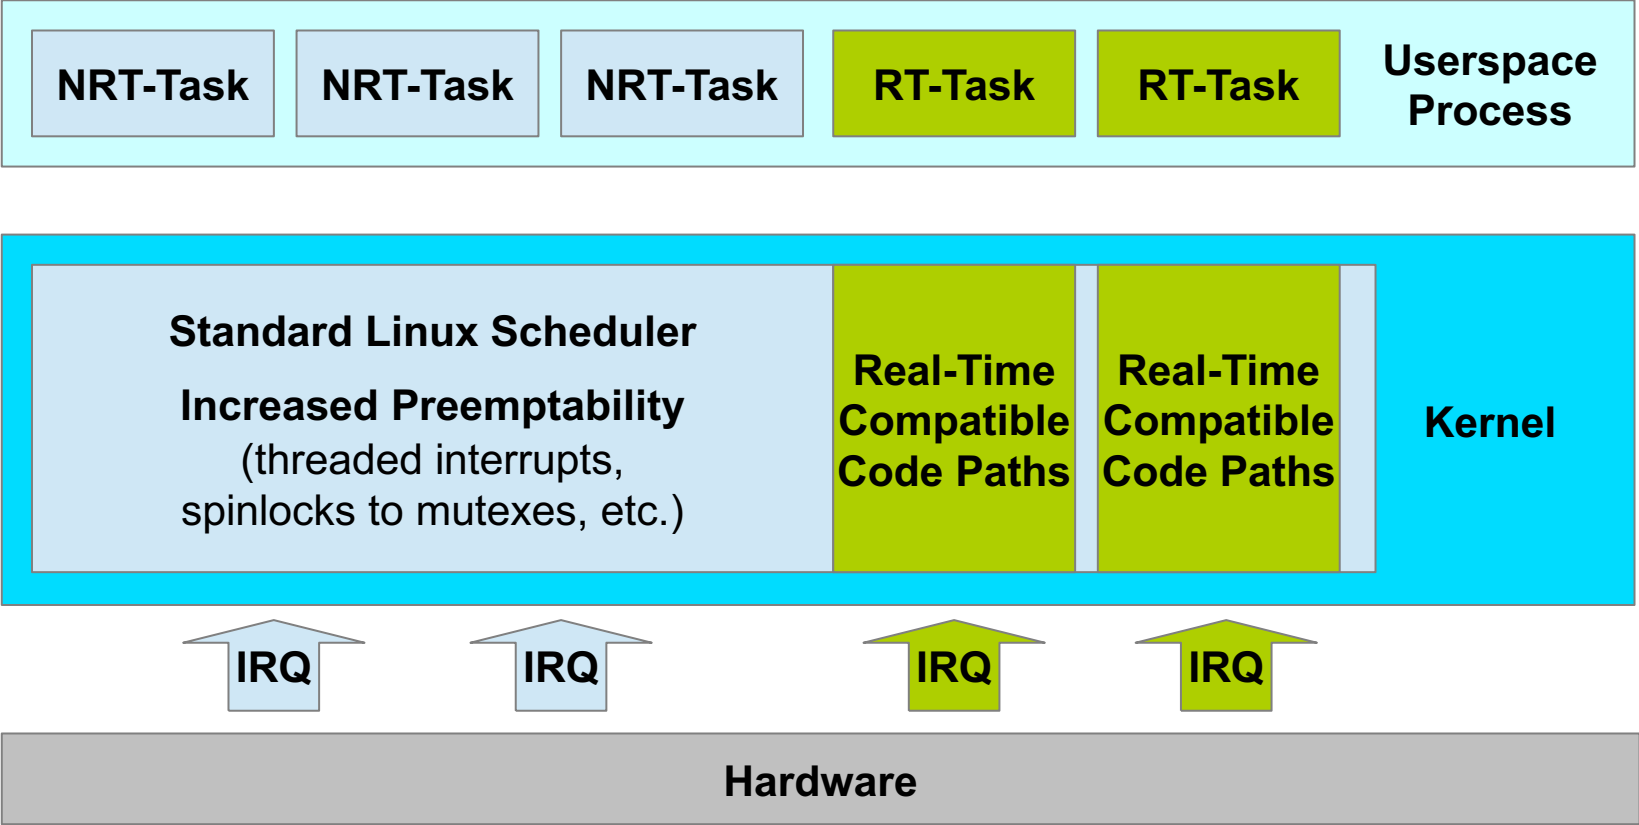
Xenomai & Real-Time Linux

Driving OSS Projects for SIEMENS

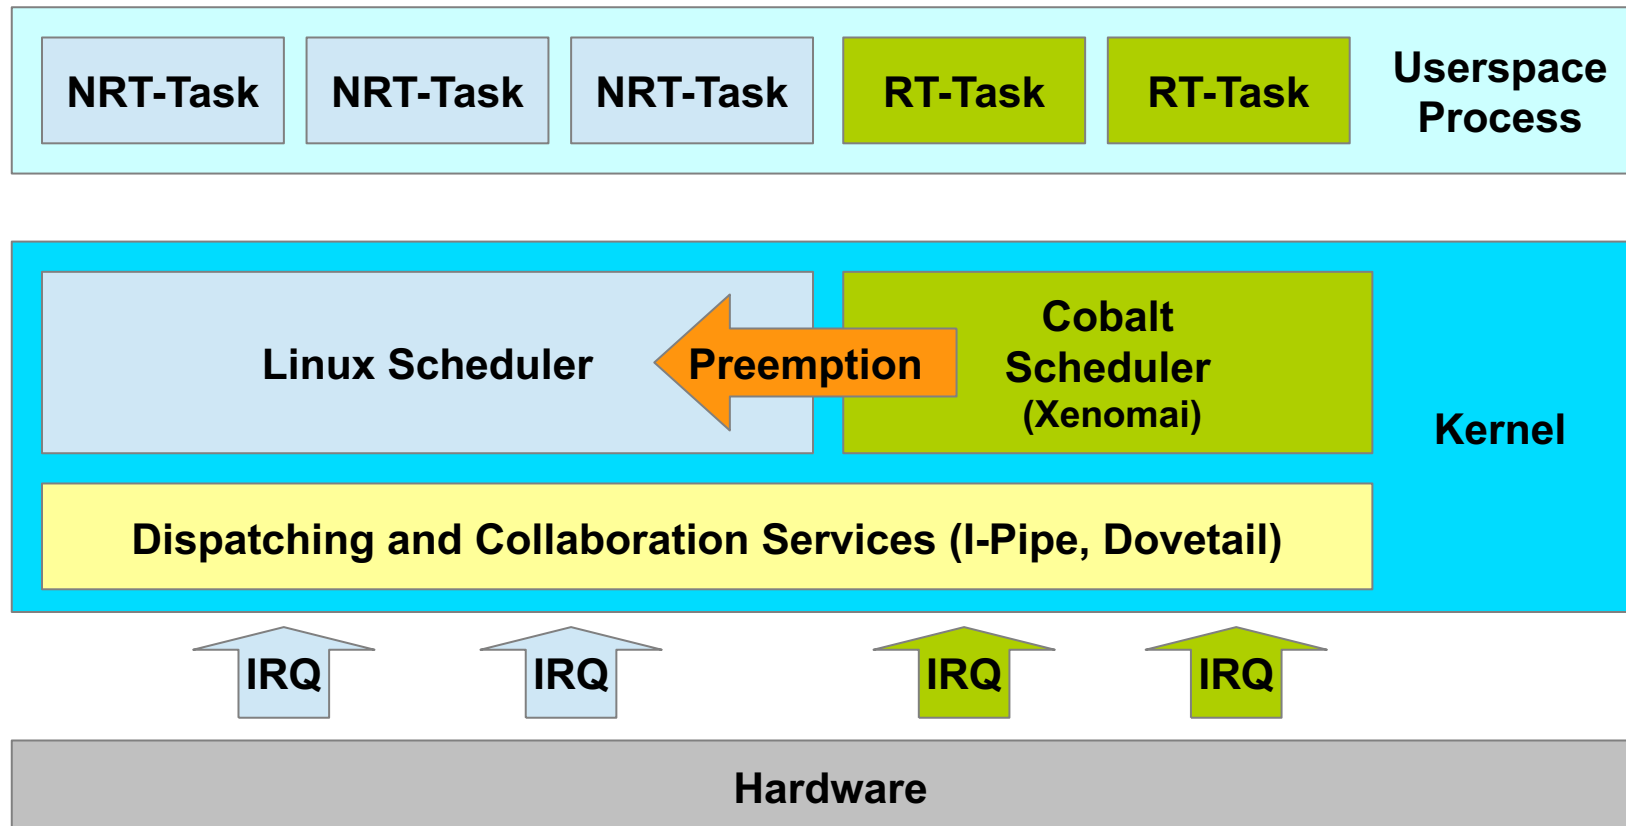
Agenda

- Real-Time Linux: Co-Kernel vs. Preempt RT
- The Role of Xenomai for Siemens
- Strengthening the Xenomai Community
 - Attracting Contributions
 - Improving Infrastructure
 - Building the Future
- The Role of Preempt RT for Siemens
- Siemens Activities around Preempt RT

Native Real-time for Linux – PREEMPT-RT



The Co-Kernel Concept



Why still using a co-kernel? (With PREEMPT-RT becoming mainline)

Allows to port complex RT application stacks to Linux

- Different scheduling mechanisms
- Scheduler tunings
- Additional APIs

Architectural separation of RT and non-RT

- Application gets early feedback when leaving RT path
- Even libraries cannot disturb RT silently

Can provide better latencies

- Lower overhead on low-end platforms
- Can decouple from scalability issues (~ 8+ cores)

The Xenomai Project

Xenomai 1.0

- Started in 2001 as RTOS portability framework for Linux
- Became part of RTAI

Xenomai 2.0

- Stand-alone again since 2005 (different design goals)
- I-pipe patch to hook into the kernel

Xenomai 3.0

- Released in 2015 with renovated co-kernel core („cobalt“)
- Support for RTOS layers over standard Linux / PREEMPT-RT

Xenomai 3.2

- Released in 2021
- New kernel patch (dovetail)

On the horizon: Xenomai 4



Xenomai Applications at Siemens



- Hard real-time requirements with minimal latencies
- API-extensions to serve needs of existing (large) code base
- Long-term support



- Maintain separation of RT and non-RT architecturally
- High through-put demands aside RT workload

Siemens Taking the Lead

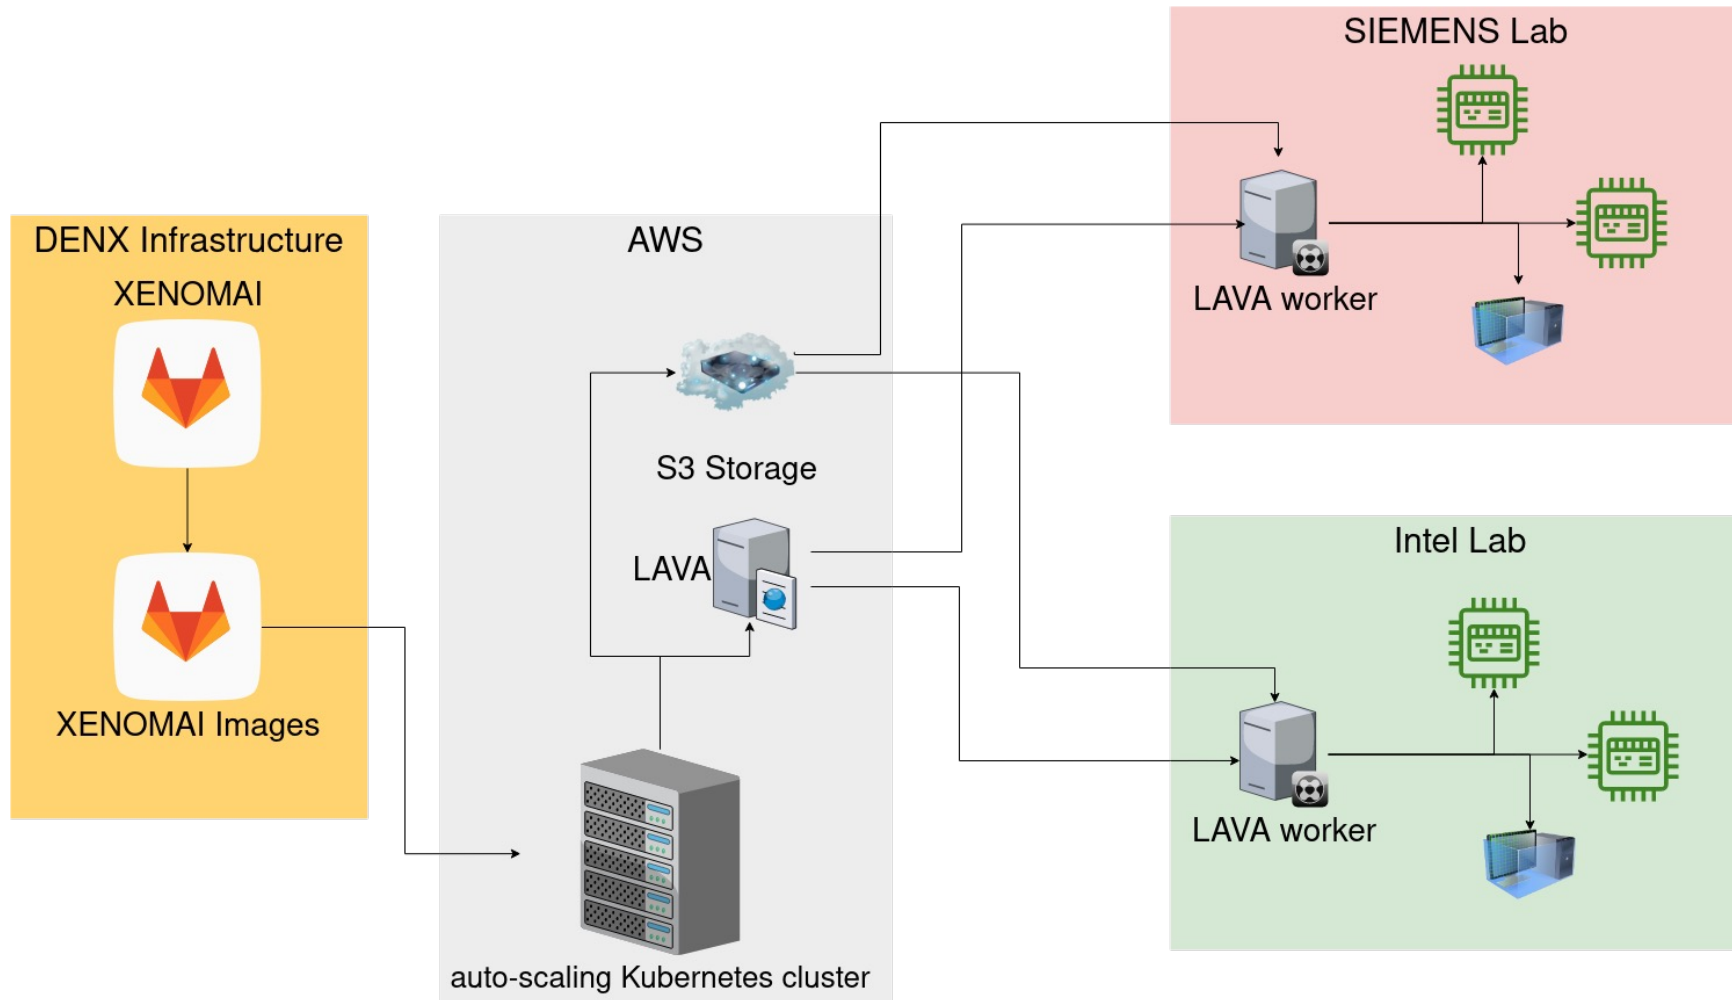
- Xenomai's active community was never particularly large
- 2016: Key contributor passed away
- 2018: Maintainer announced to step back from daily business
 - Wanted to focus on next generation („EVL“ -> Xenomai 4)
 - As long-term contributor, Jan Kiszka and, thus, Siemens was asked to take over
- What to do?
 - Migrate products to PREEMPT-RT?
 - Take the lead – and invest further?
- Affected business units decided to enable the latter
 - Jan becomes Xenomai maintainer in late 2018
 - 2-3 further Siemens Technology engineers working on upstream topics (part-time)
 - Siemens invests infrastructure and tries to grow the community
- 2020: Intel becomes new and important Xenomai contributor

Xenomai Development Activities

- Addressing the Y2038 problem in Xenomai
 - Y2038 problem in a nutshell: `time_t` will overflow at Tuesday, 19 of January 2038, 3:14:07 UTC
 - Only compat applications (32 bit applications) are affected
 - Xenomai project still cares about 32 bit applications
- Linux: The Y2038 problem is addressed in Linux ≥ 5.10 , glibc ≥ 2.34 required
- Current state
 - Kernel: Nearly all y2038 affected syscalls have a y2038 safe implementation now, test suite has been extended to test them
 - Kernel: Some more work around struct `sched_param` necessary
 - User space: Tests with glibc ≥ 2.34 currently ongoing

```
struct timespec {  
    time_t tv_sec;  
    long tv_nsec;  
};
```

Xenomai CI Infrastructure



<https://lava.xenomai.org>

<https://source.denx.de/Xenomai/xenomai-images>

Xenomai Test Matrix for 3.2 stable

- Architectures
 - x86_64
 - ARM
 - ARM64
- Kernel versions
 - 4.19 (ipipe)
 - 5.4 (ipipe)
 - 5.10 (dovetail)
- Boards (SIEMENS Lab)
 - Beagle bone black (arm)
 - Hikey (arm64)
 - IPC 427E (x86_64)
 - Virtual targets (x86_64, arm, arm64)

The screenshot displays the Xenomai Test Matrix interface for the 'stable/v3.2.x' branch. The interface is organized into several sections:

- Navigation:** A top-left sidebar contains buttons for 'next', 'stable/v3.0.x', 'stable/v3.1.x', and 'stable/v3.2.x' (which is selected and highlighted).
- Branch Selection:** A central box shows 'stable/v3.2.x' with the commit hash '#11812' and a 'Child' button.
- Build Results:** A 'Build' column lists 18 build configurations, each with a green checkmark indicating success. The configurations include combinations of kernel versions (4.19 and 5.4) and architectures (beagle-bone-black, hikey, qemu-amd64, qemu-arm64, qemu-armhf, and x86-64-efi).
- Test Results:** A 'Test' column lists 18 test configurations, each with a green checkmark indicating success. The configurations mirror the build configurations, using 'lava-test' as the prefix.

Building the Future – Xenomai 4

- Use cases for co-kernels are expected to stay
- Xenomai's core is ageing
 - Limited SMP scalability
 - Several unmaintained RT drivers
 - Likely no longer relevant features
- Renovate again: EVL (“Even Less”)
 - Fresh restart, no baggage
 - Focus on co-kernel
 - Maintain closer to kernel versions
- No new island without bridges
 - Dovetail patch already shared between Xenomai 3 and EVL/Xenomai 4
 - Common Xenomai Platform define shared APIs, tools etc.
 - Smooth migration from 3 to 4 is the goal

| PREEMPT_RT @ SIEMENS

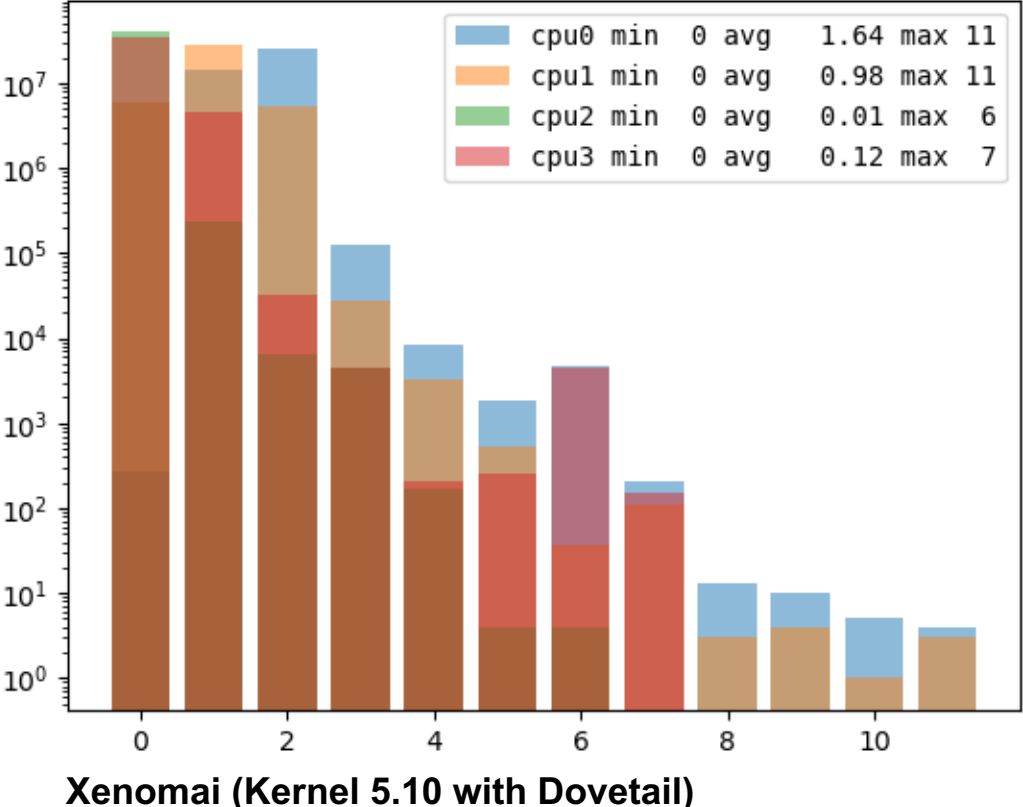
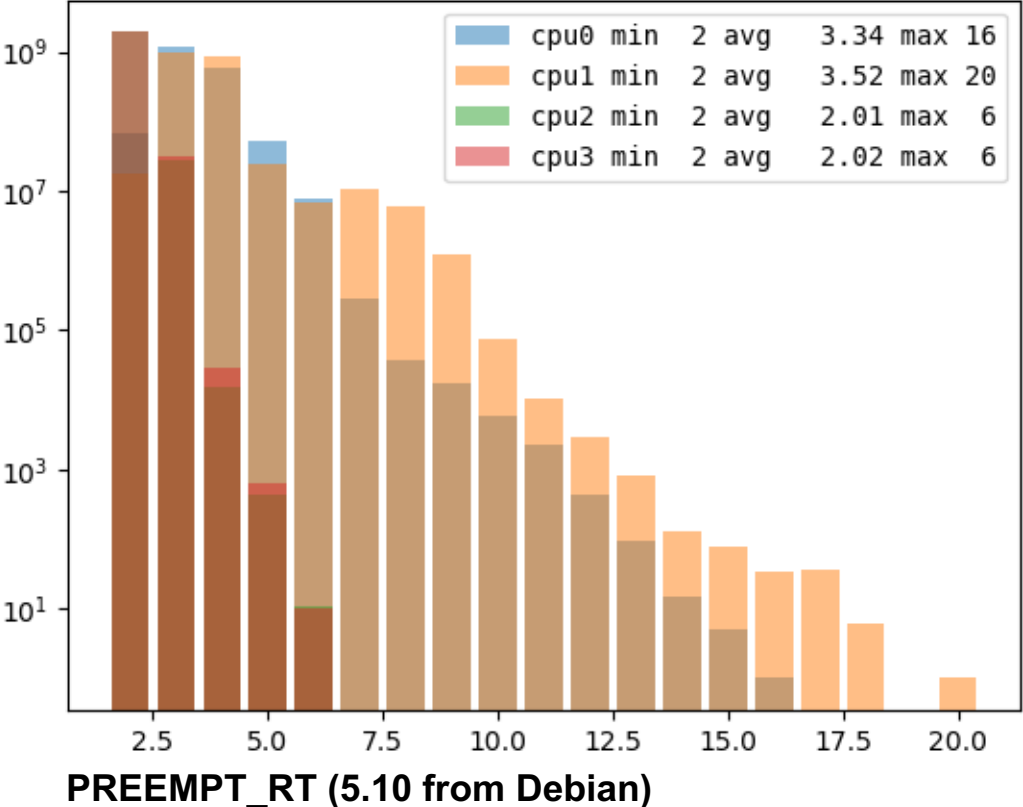
Ongoing PREEMPT_RT related activities

PREEMT_RT – Preparing the next generation RT platform

- Product line driving a new vision: Realtime enabled Linux platform
- Moving away from dedicated “RT only” systems
- Opening the system for third party integration
- RT applications shall also be delivered by external partners

- Which platform to use?

PREEMT_RT vs. Xenomai – A low level comparison



Values in μ s
 CPU cores 2+3 were isolated, serving RT tasks only
 Hardware: SIEMENS IPC 427E



PREEMT_RT – Challenging Topics

- Resource management becomes a major topic
 - cgroups v2 for flexible resource management at runtime
 - Tuned profiles for optimizing RT capabilities (frequency scaling, cstates, pstates, ...)
- Combine RT with the further Linux features
 - eXpress data path (XDP), Time Sensitive Networking (TSN)
 - Container runtimes (docker, podman, OpenStack, ...)
- System tuning
 - Validation of system configuration for best RT performance
 - Regression testing, following upstream releases as soon as possible
- Provide all that in an easy-to-re-use manner
 - ISAR based build process
 - Easy to use reference images

PREEMT_RT – Public Contributions

Public contributions to address our needs as early as possible

- PREEMPT_RT
 - Regression reporting
 - Debugging
 - tuned improvements
 - Financial support (Premier level)
- Civil Infrastructure Platform (CIP)
 - SLTS (Super Long Term Support) kernel maintenance
 - Financial Support (Platinum level)



<https://wiki.linuxfoundation.org/realtime/start>

<https://www.cip-project.org/>

Summary

No “one size fits it all”

- PREEMPT-RT for common RT use cases
- Xenomai for harder separation and special needs

Siemens is committed to both

- Maintaining and evolving Xenomai
- Sponsoring PREEMPT-RT
- Contributing requirements and code
- Share know-how, tools, tests between both



| Contact

Jan Kiszka

T CED
Otto-Hahn-Ring 6
81739 München
Germany

E-mail jan.kiszka@siemens.com

Florian Bezdeka

T CED SES-DE
Günther-Scharowsky-Str. 1
91058 Erlangen
Germany

E-mail florian.bezdeka@siemens.com