

Implementierung gemischt kritischer Systeme durch statische Hardwarepartitionierung

Ralf Ramsauer*, Jan Kiszka†, Daniel Lohmann‡, Wolfgang Mauerer*†

* Ostbayerische Technische Hochschule Regensburg

† Siemens AG, Corporate Research and Technology, München

‡ Gottfried Wilhelm Leibniz Universität Hannover

Frühjahrstreffen der Fachgruppen Betriebssysteme, Hannover

2. März 2018



OSTBAYERISCHE
TECHNISCHE HOCHSCHULE
REGENSBURG

SIEMENS

11
102
1004

Leibniz
Universität
Hannover

Motivation: Mixed Criticality

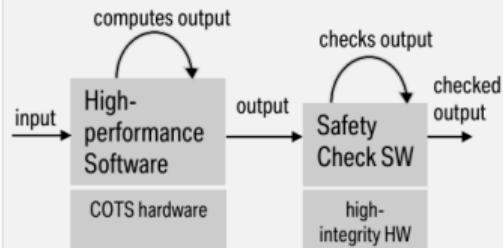
- ▶ Driven by consolidation of physical hardware units [1]
- ▶ SMP is everywhere
- ▶ Reduction of physical control units
- ▶ **System of systems** on a chip
 - ▶ Increasing complexity
 - ▶ Scalability
 - ▶ Maintainability
- ▶ Consolidation of multiple software stacks requires safe isolation



Separate Automotive Control Units

Image © CVEL

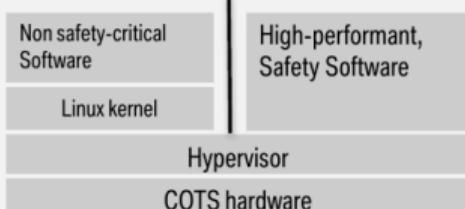
System Architecture A



Assumptions/Drawbacks:

- Needs decomposition of safety of complex function to a simple checking on lower performance high-integrity HW
- Checking must detect subtle errors from HW and OS of high-performance computation.

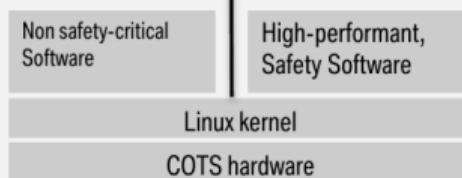
System Architecture B



Assumptions/Drawbacks:

- Hypervisor & HW guarantees isolation
- Safety SW without underlying OS
- If needed, scheduling, multi threading & file system is implemented in safety SW

System Architecture C



Assumptions/Drawbacks:

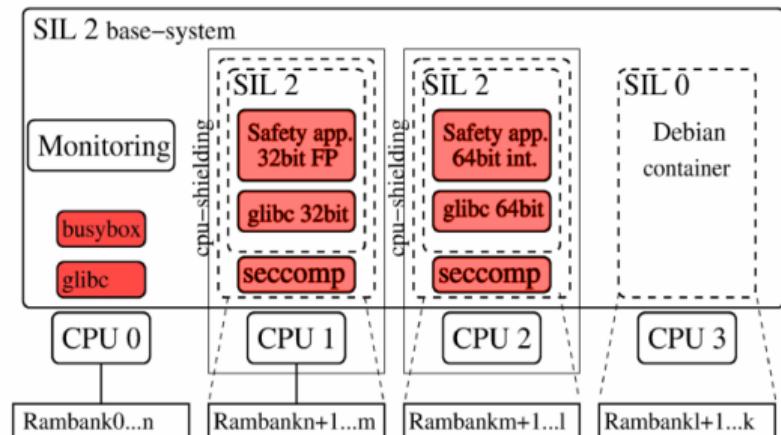
- Linux kernel provides sufficient isolation
- Safety SW uses Linux scheduling, multi threading, file system etc.
- Parts of Linux kernel and HW functionality must be qualified.

Is Linux Kernel Development Good Enough to Make Your Life Depend on It?

Image © Lukas Bulwahn

Main Architecture Decisions:

- Safety-critical and non-safety critical applications run on the same kernel
- Isolation is achieved with:
 - **CPU shielding**
 - use of dedicated cores and memory regions
- Unintended behavior of safety-critical applications is limited with **seccomp**
- System and applications are monitored with an application on a dedicated core
- Safety-critical applications use **glibc**



SIL2LinuxMP High-Level Architecture

Image © Nicholas McGuire

Hypervisor-assisted solutions

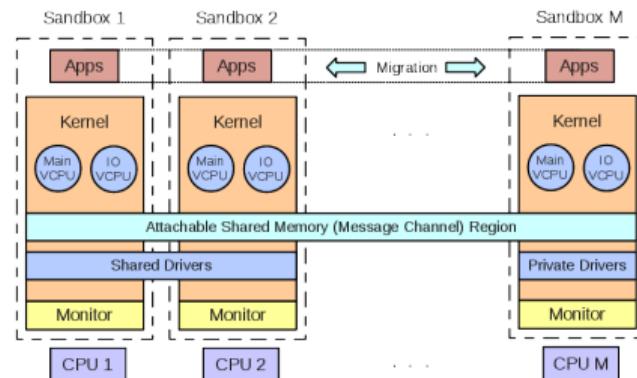
► **Quest-V** [2]

- ▶ Allows direct I/O access
- ▶ Rich set of device drivers (OS + VMM)
- ▶ Virtualisation *only* for isolation
- ▶ Communication: Shared memory + IPI
- ▶ Only trap on violations
- ▶ Traditional boot sequence

► **PikeOS** [3]

► **XtratuM** [4], **LTZvisor** [5], ...

QUEST



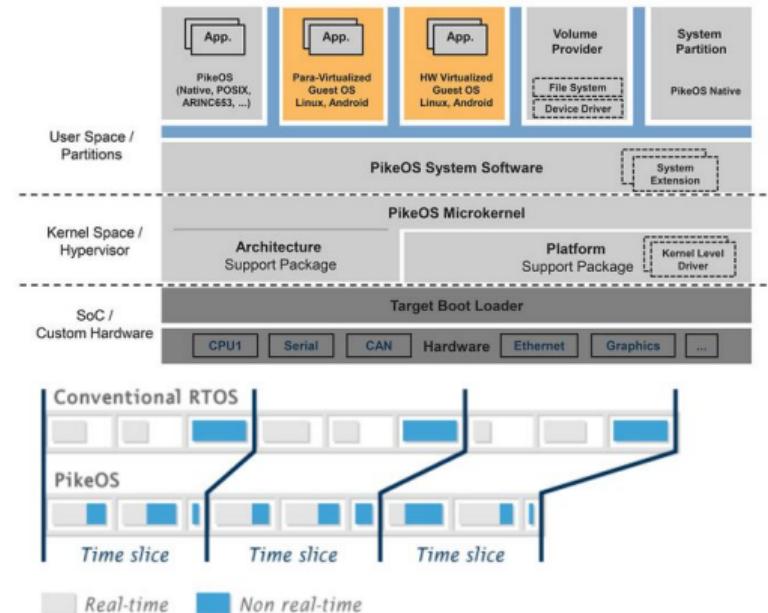
Quest-V overview

Images © Rich West

Hardware Partitioning

Hypervisor-assisted solutions

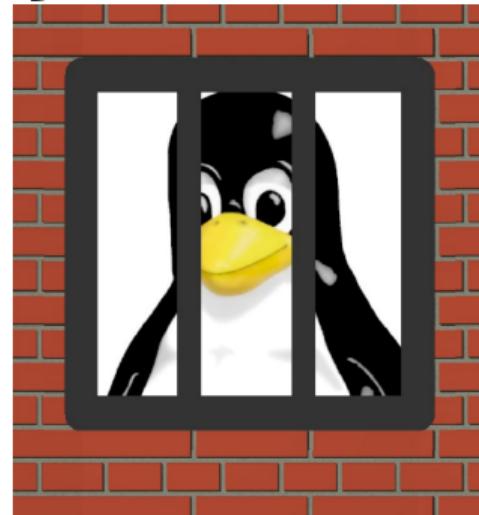
- ▶ **Quest-V** [2]
- ▶ **PikeOS** [3]
 - ▶ Allows direct I/O access
 - ▶ Paravirtualisation, hardware-assisted virtualisation
 - ▶ Time or Event triggered scheduling
- ▶ **XtratuM** [4], **LTZvisor** [5], ...



PikeOS architecture

Images © SYSGO AG

Jailhouse



Yet another hypervisor?

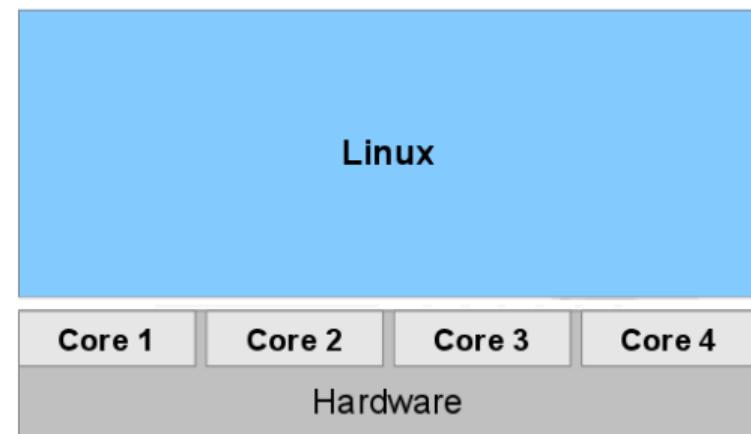
Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

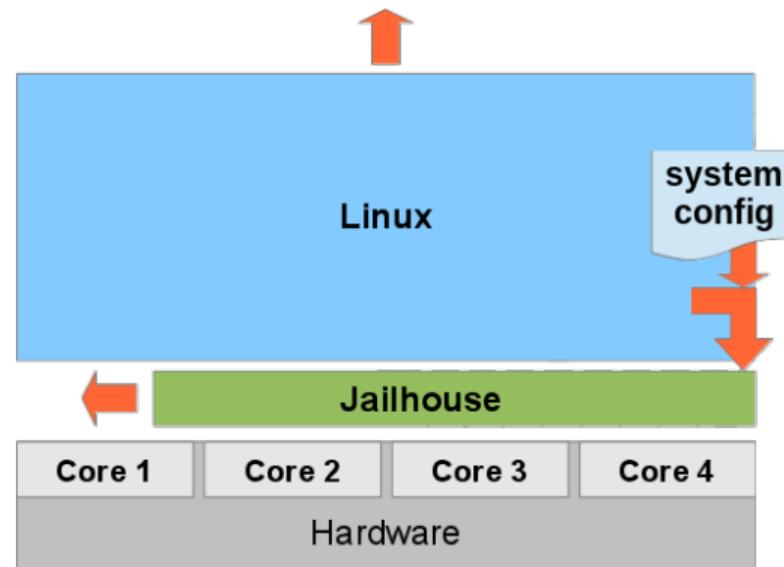


1. Boot phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

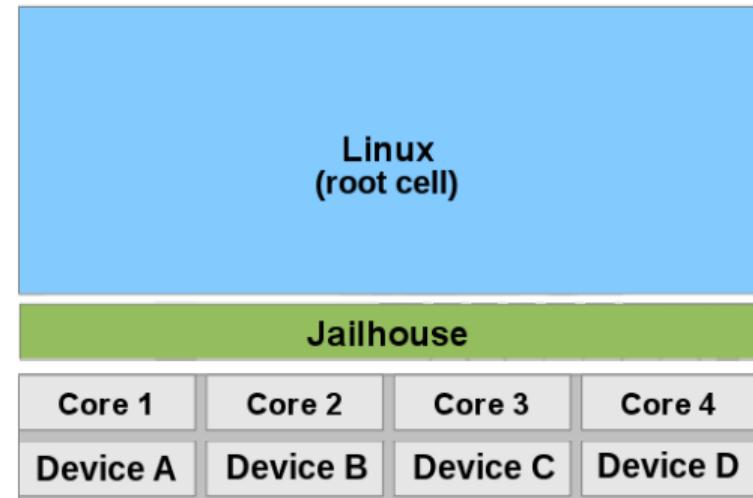


1. Boot phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

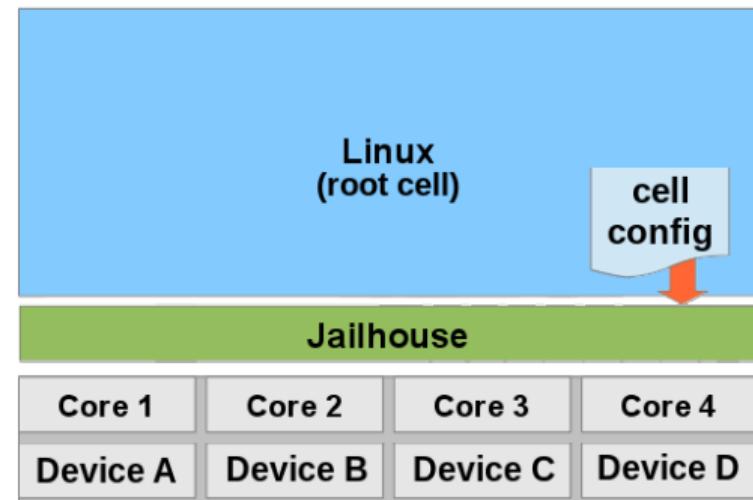


1. Boot phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

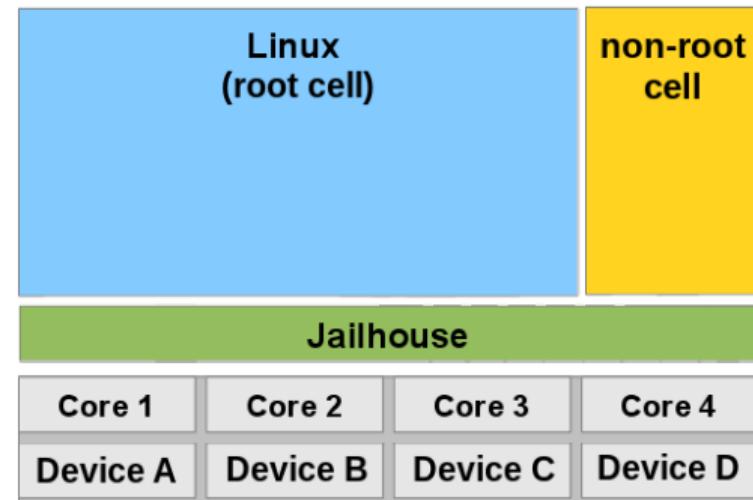


2. Partitioning phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

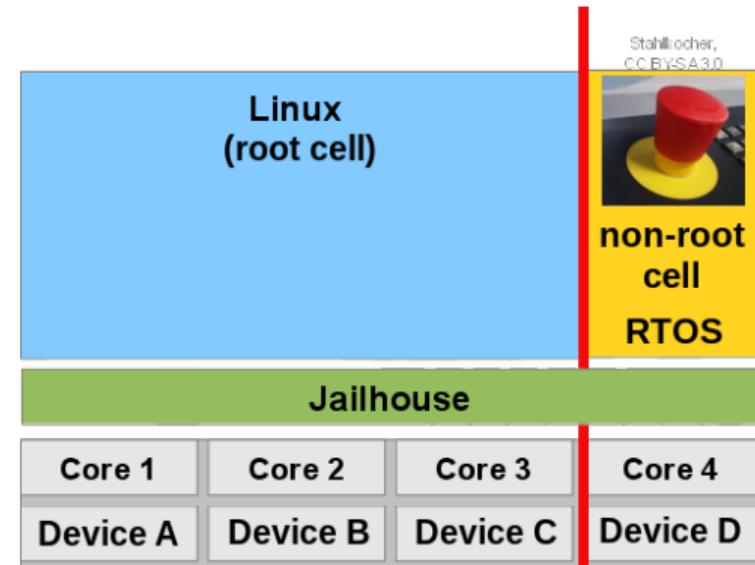


2. Partitioning phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

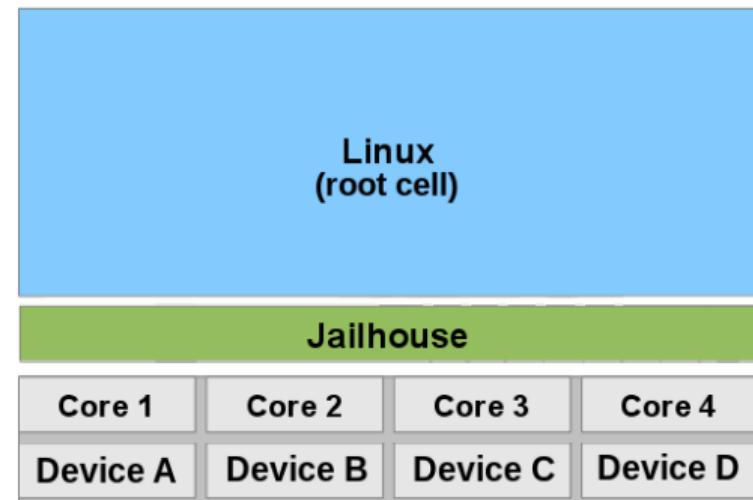


3. Operational phase

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

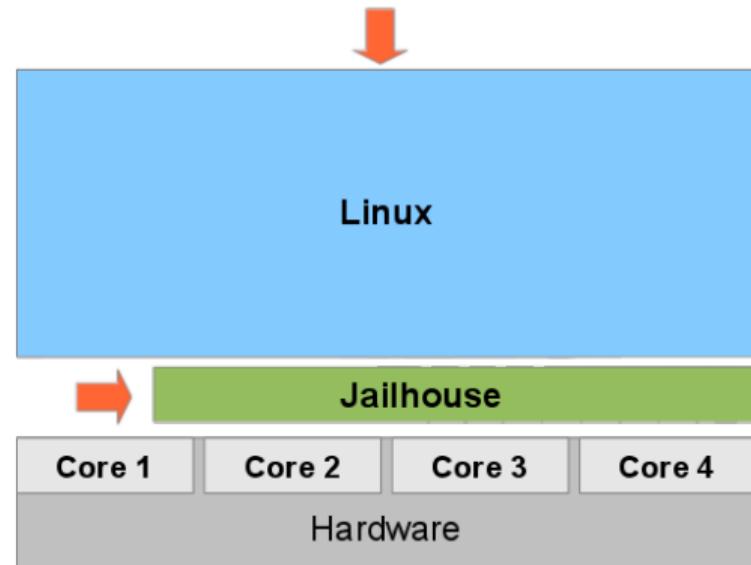


(Dev: Destroy cells)

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation

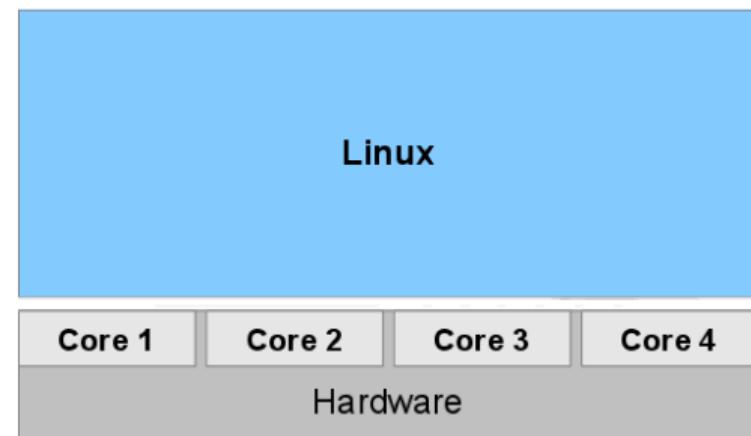


(Dev: Disable hypervisor)

What makes Jailhouse different?

Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation



1. Boot phase

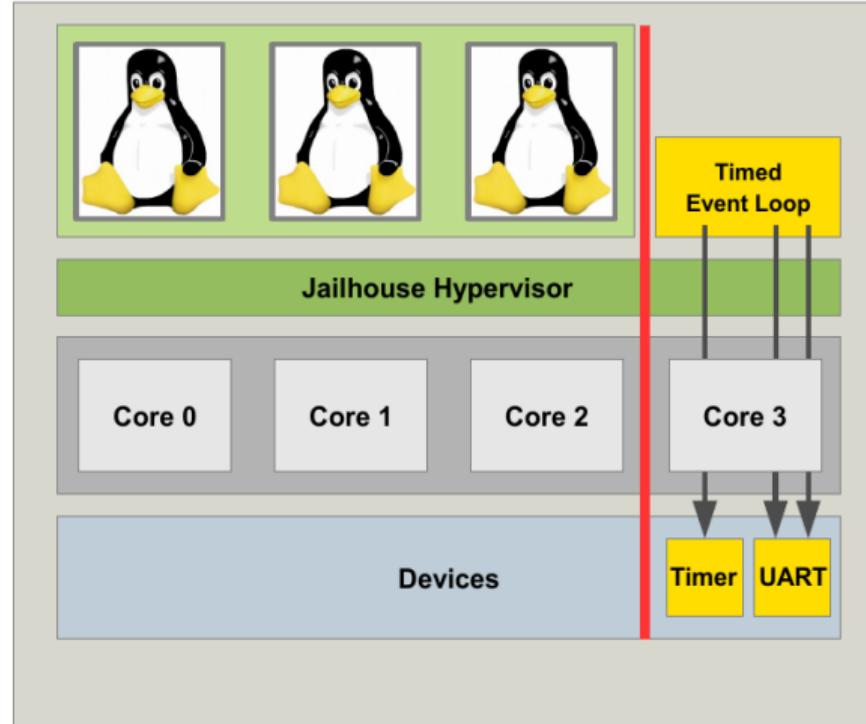
Jailhouse [6], an *Exohypervisor*

- ▶ Minimalist hypervisor skeleton
(cf. Exokernel approach [7])
- ▶ **Offload uncritical work to Linux**
 - ▶ System boot and initialisation
 - ▶ Partition »cell« management
 - ▶ Control and monitoring
 - ▶ Deferred VMM activation
- ▶ **Prefer simplicity over features**
 - ▶ Partition booted system
instead of booting Linux
 - ▶ Resource access control
instead of resource virtualisation
 - ▶ 1:1 static resource assignment
instead of scheduling

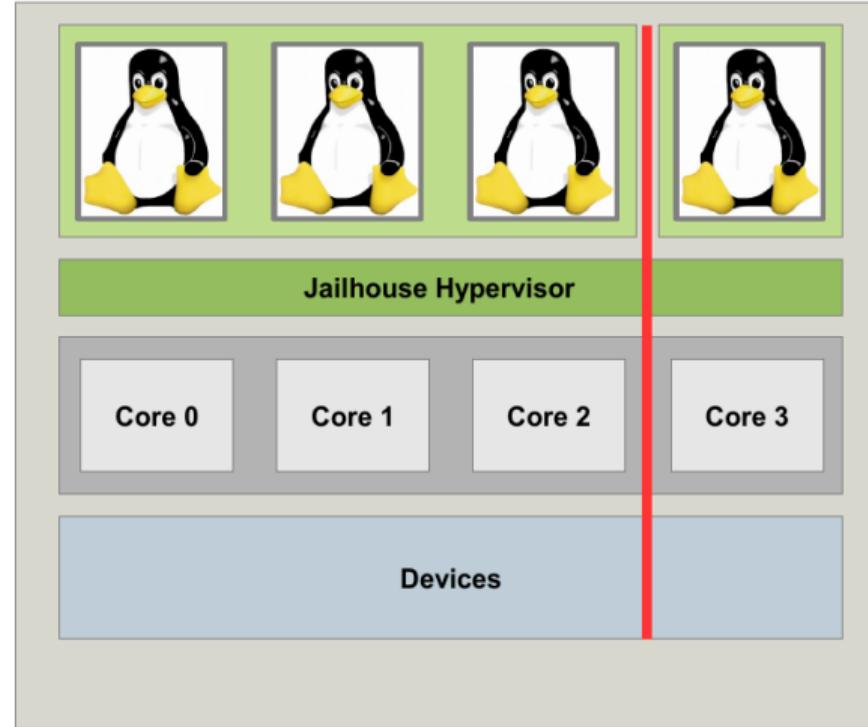
Small code-base and tiny impact

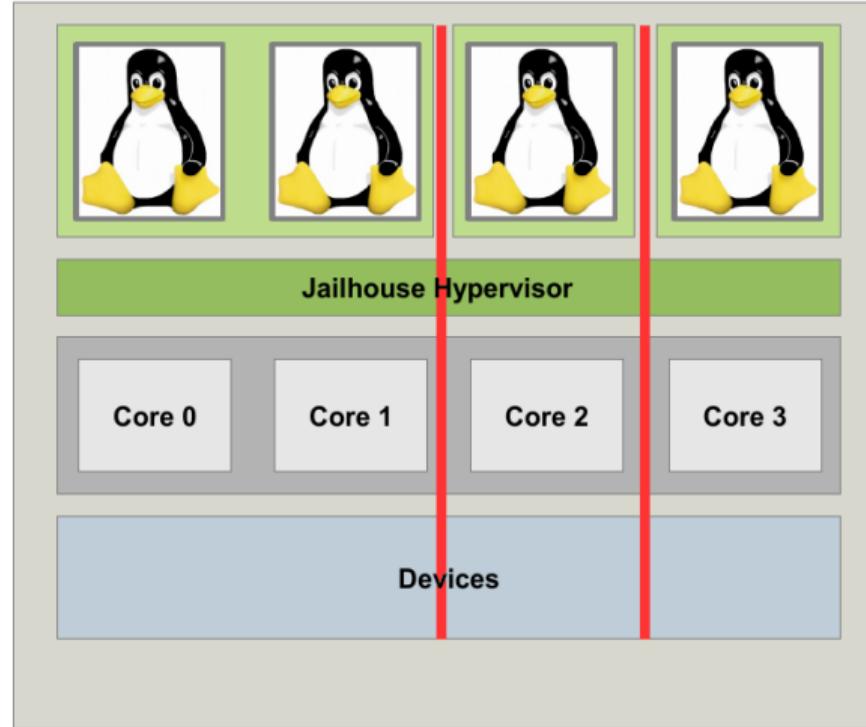
- ▶ $\approx 7\text{kLoC}$ on armv7
 - ▶ Simplifies certification efforts
 - ▶ Suitable basis for formal verification
- ▶ Try to hide (reduce traps), but don't hide existence
- ▶ $\#\text{Partitions} \leq \#\text{CPUs}$, sufficient for many real-world use cases
- ▶ \Rightarrow Maintain real-time capabilities *by design*

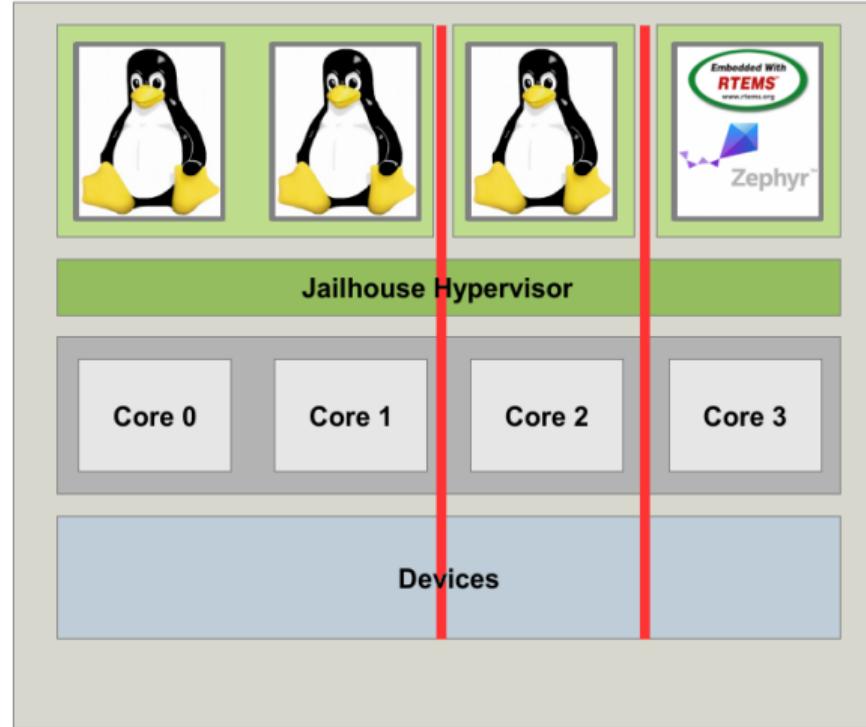
Example use cases

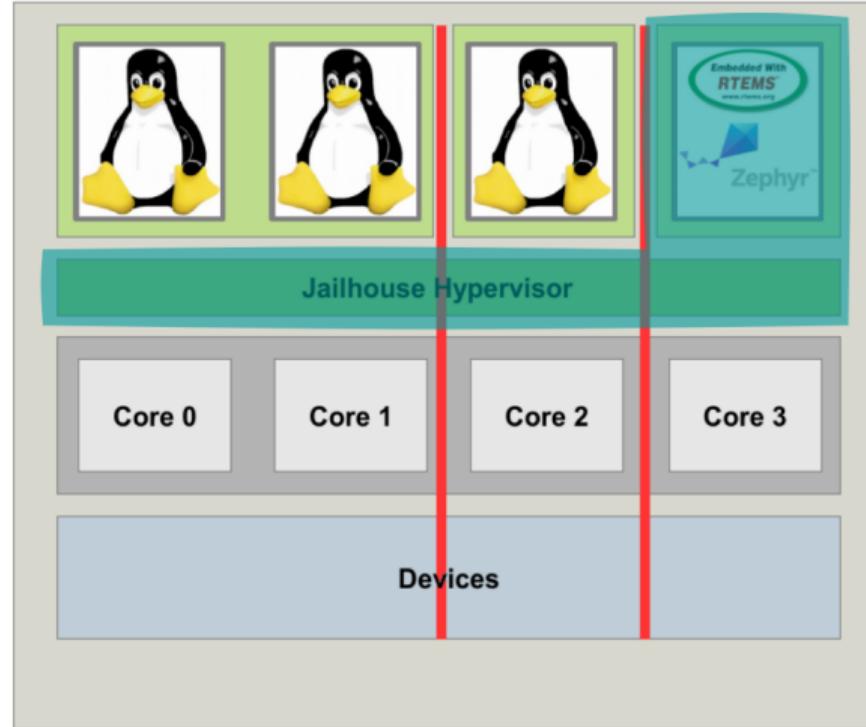


No VM Exits! (x86)









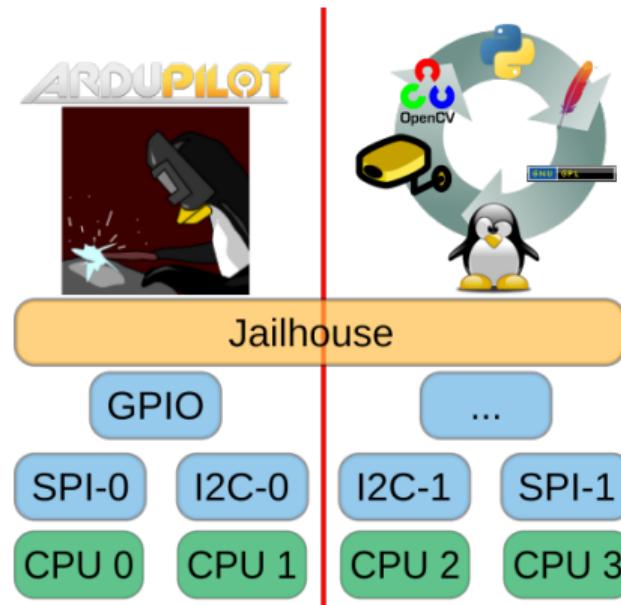
Burn-in test

- ▶ Typical mixed-criticality scenario
- ▶ Linux/RTOS as common use case
 - ▶ uncritical: computer vision task, video streaming
 - ▶ critical: flight control (hardware and software)
- ▶ Legacy software stack
- ▶ Jailhouse support out of the box
- ▶ Only **board support**



Burn-in test

- ▶ Typical mixed-criticality scenario
- ▶ Linux/RTOS as common use case
 - ▶ uncritical: computer vision task, video streaming
 - ▶ critical: flight control (hardware and software)
- ▶ Legacy software stack
- ▶ Jailhouse support out of the box
- ▶ Only **board support**



Conclusion

- ▶ Solid testament for implementing real-time safety critical systems with Jailhouse
- ▶ Jailhouse as platform or playground for other academic approaches
- ▶ Hardware/Software codesign towards zero traps

Conclusion

- ▶ Solid testament for implementing real-time safety critical systems with Jailhouse
- ▶ Jailhouse as platform or playground for other academic approaches
- ▶ Hardware/Software codesign towards zero traps

Future Work

- ▶ Linux mainline integration (Ongoing!)
- ▶ Safety certification (Ongoing!)
- ▶ Make Jailhouse Linux agnostic (WIP by TI)
- ▶ Sound quantification of hypervisor influence (there are certain traps)
- ▶ Consider extending jailhouse for heterogeneous architectures
 - ▶ +GPU
 - ▶ +FPGA
 - ▶ +PRU

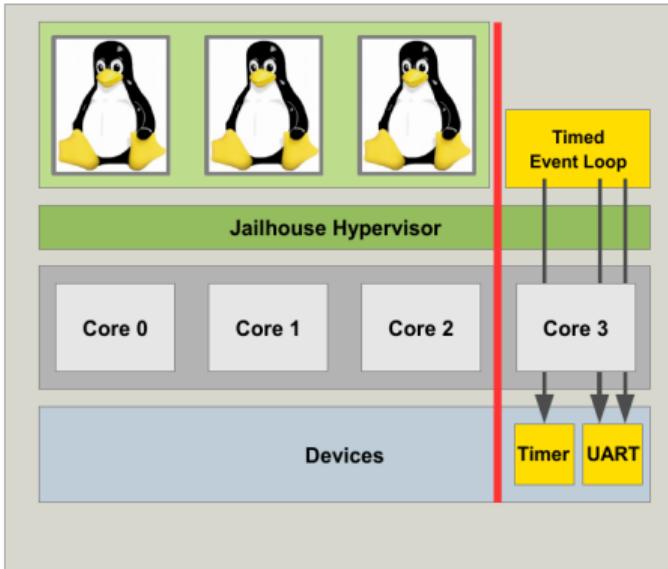
Thank you!



<https://github.com/siemens/jailhouse>
[<jailhouse-dev@googlegroups.com>](mailto:jailhouse-dev@googlegroups.com)

[<ralf.ramsauer@othr.de>](mailto:ralf.ramsauer@othr.de), [<jan.kiszka@siemens.com>](mailto:jan.kiszka@siemens.com)

- [1] Manfred Broy. "Challenges in Automotive Software Engineering". In: *Proceedings of the 28th International Conference on Software Engineering (ICSE)*. (Shanghai, China). New York, NY, USA: ACM Press, 2006, pp. 33-42.
- [2] Ye Li, Richard West, and Eric Missimer. "A Virtualized Separation Kernel for Mixed Criticality Systems". In: *Proceedings of the 10th USENIX International Conf. on Virtual Execution Environments (VEE)*. 2014.
- [3] Robert Kaiser and Stephan Wagner. "Evolution of the PikeOS microkernel". In: *First International Workshop on Microkernels for Embedded Systems*. 2007, p. 50.
- [4] Alfons Crespo, Ismael Ripoll, and Miguel Masmano. "Partitioned Embedded Architecture based on Hypervisor: The XtratuM approach". In: *Proceedings of the 8th European Dependable Computing Conference (EDCC)*. IEEE. 2010.
- [5] Sandro Pinto, Jorge Pereira, Tiago Gomes, Adriano Tavares, and Jorge Cabral. "LTZVisor: TrustZone is the Key". In: *LIPics-Leibniz International Proceedings in Informatics*. Vol. 76. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2017.
- [6] Ralf Ramsauer, Jan Kiszka, Daniel Lohmann, and Wolfgang Mauerer. "Look Mum, no VM Exits! (Almost)". In: *Proceedings of the 13th Workshop on Operating System Platforms for Embedded Real-Time Applications (OSPERT)*. 2017.
- [7] Dawson R. Engler, M. Frans Kaashoek, and James O'Toole. "Exokernel: An Operating System Architecture for Application-Level Resource Management". In: *Proceedings of the 15th ACM Symposium on Operating Systems Principles (SOSP)*. 1995, pp. 251-266.
- [8] Udo Steinberg and Bernhard Kauer. "NOVA: A Microhypervisor-based Secure Virtualization

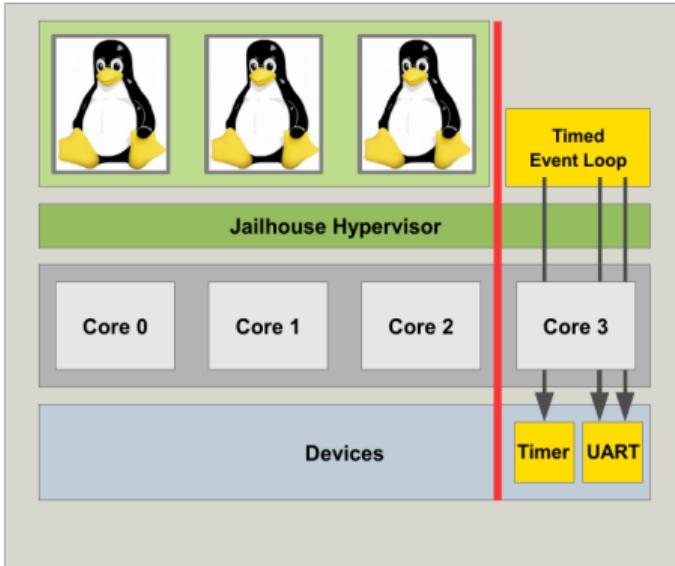


raw inmate: timed event loop (GIC demo)

cyclic timer interrupt, measure jitter

- ▶ \$ modprobe jailhouse
- ▶ \$ jailhouse enable tk1.cell
- ▶ \$ jailhouse cell create tk1-demo.cell
- ▶ \$ jailhouse cell load tk1-demo gic-demo.bin
- ▶ \$ jailhouse cell start tk1-demo

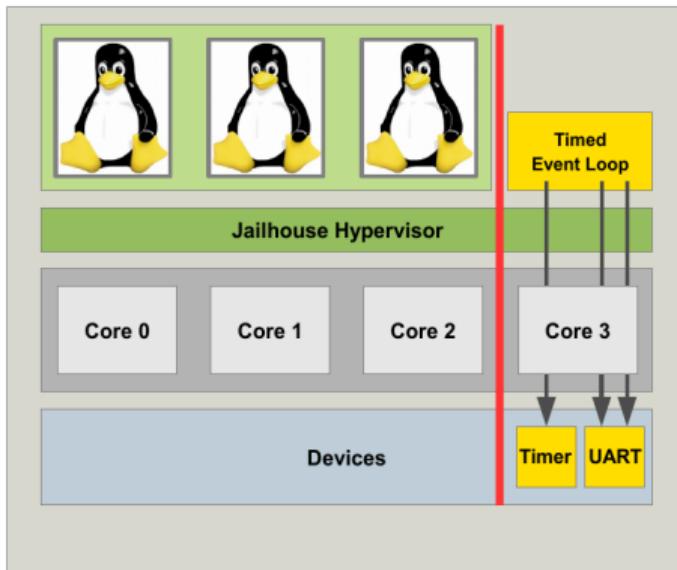
raw inmate: timed event loop (GIC demo)



```

Initializing Jailhouse hypervisor v0.7 (26-g918bec06) on CPU 1
Code location: 0xf0000040
Initializing processors:
CPU 1... OK
CPU 2... OK
CPU 0... OK
CPU 3... OK
Activating hypervisor
Created cell "jetson-tk1-demo"
Page pool usage after cell creation: mem 82/16107, remap 69/131072
Cell "jetson-tk1-demo" can be loaded
Started cell "jetson-tk1-demo"
Initializing the GIC...
Initializing the timer...
Timer fired, jitter: 3083 ns, min: 3083 ns, max: 3083 ns
Timer fired, jitter: 2333 ns, min: 2333 ns, max: 3083 ns
Timer fired, jitter: 2416 ns, min: 2333 ns, max: 3083 ns
Timer fired, jitter: 3916 ns, min: 2333 ns, max: 3916 ns
Timer fired, jitter: 3749 ns, min: 2333 ns, max: 3916 ns
Timer fired, jitter: 3499 ns, min: 2333 ns, max: 3916 ns
[...]

```



raw inmate: timed event loop (GIC demo)

- ▶ \$ jailhouse cell destroy tk1-demo
- ▶ \$ jailhouse disable

[...]

```
Timer fired, jitter: 3416 ns, min: 2166 ns, max: 3916
Timer fired, jitter: 3499 ns, min: 2166 ns, max: 3916
Timer fired, jitter: 3499 ns, min: 2166 ns, max: 3916
Closing cell "jetson-tk1-demo"
Shutting down hypervisor
Releasing CPU 2
Releasing CPU 0
Releasing CPU 1
Releasing CPU 3
```