

# Jailhouse: An Open-Source Hypervisor for connected (and many other) things

Jan Altenberg

Linutronix GmbH

## Overview

- 1 Motivation
- 2 Jailhouse
- 3 Challenges
- 4 Conclusion

## Why using a hypervisor?

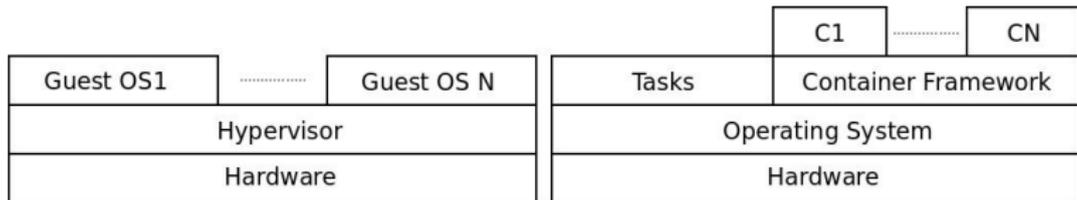
- ❑ Consolidation
- ❑ Isolation
- ❑ Mixed-Critical applications

## What do we need?

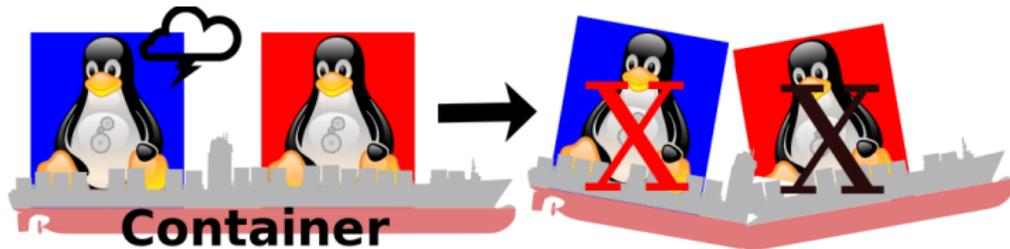
- ❑ A hypervisor that scales pretty well
- ❑ A small code base (verification and certification)
- ❑ Best possible isolation

## So, why not containers?

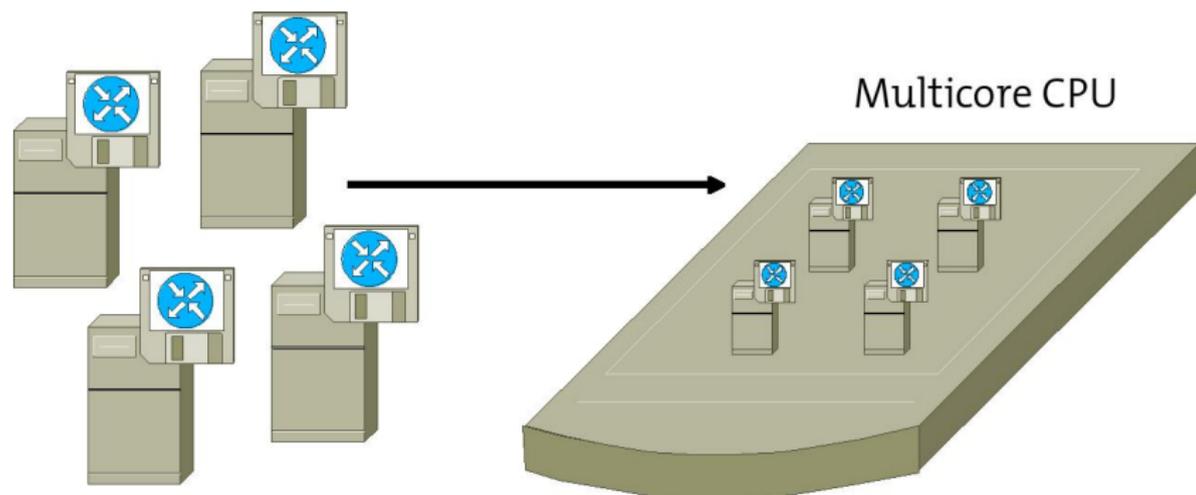
- ❏ Think about isolation!
- ❏ Different operating systems



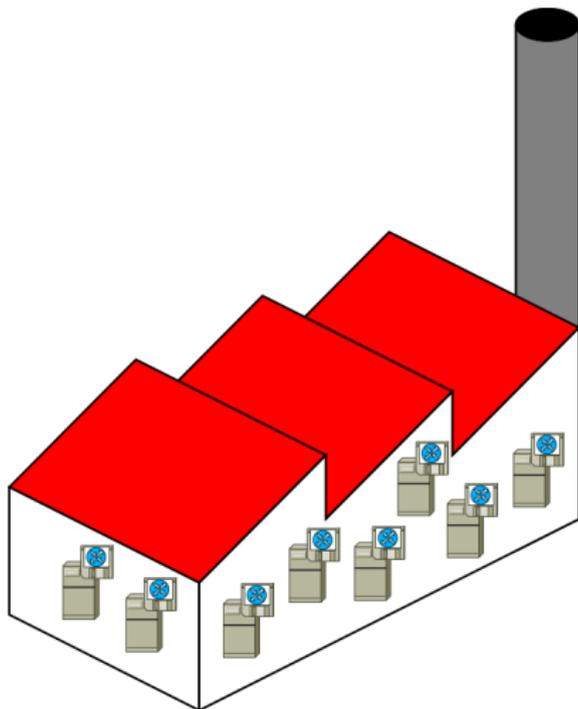
## Impact of an Operating System fault when using containers



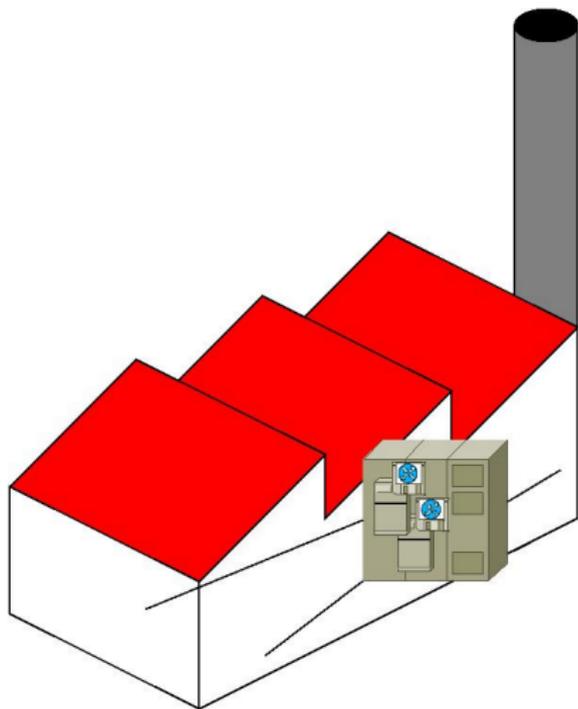
## Example: Use multicore CPUs efficiently



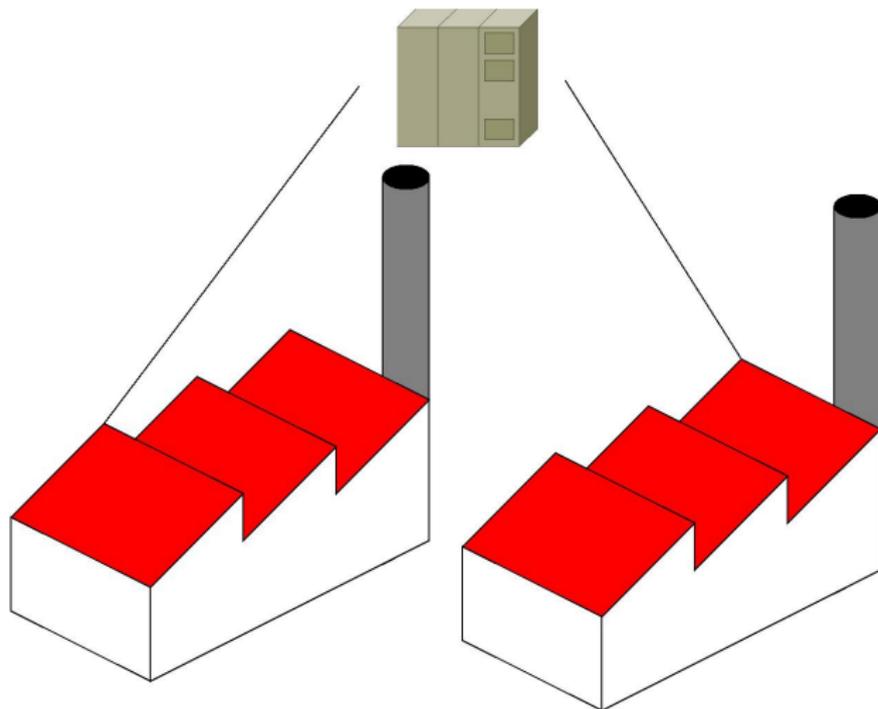
## Example: Consolidation



## Example: Consolidation ctd.



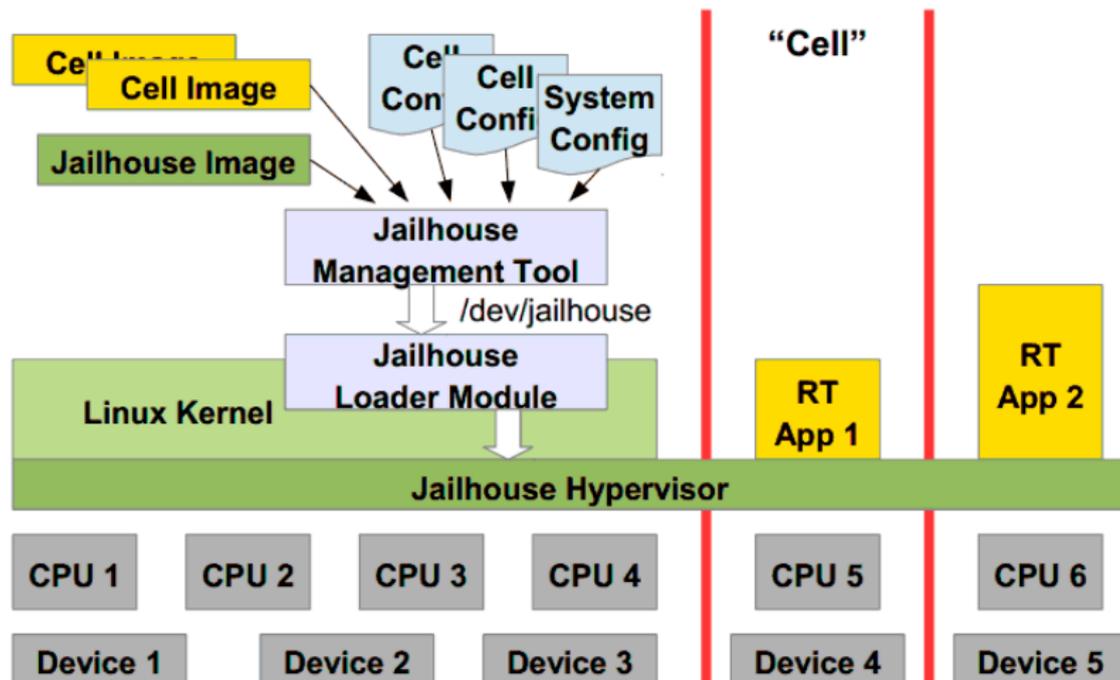
## Example: Consolidation ctd.



## What is Jailhouse?

- ❑ Hypervisor with static partitioning
- ❑ Open-Source (GPLv2)
- ❑ Developed and maintained by Siemens
- ❑ Management Interface based on Linux
- ❑ Isolates so called "cells"
- ❑ These cells can run bare-metal applications or any OS
- ❑ Currently x86 and ARM are supported

## What is Jailhouse? (Source: Siemens)



## Jailhouse Hypervisor

### Static Partitioning (NO emulation)

- ❏ Resources must be assigned to a cell
- ❏ Optimal isolation of each cell
- ❏ Full HW control
- ❏ The so called "root cell" is re-used as a management interface. This leads to:
- ❏ Small and simple code base (round about 8k LoC per architecture)

## Cell configuration

```

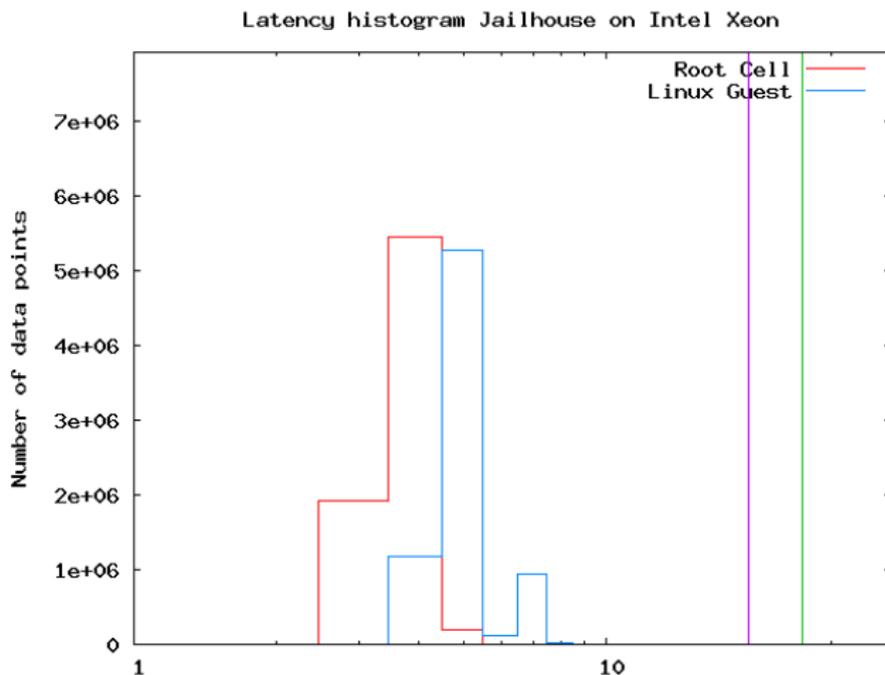
.cpus = {
    0xf,
},
.mem_regions = {
    /* RAM */ {
        .phys_start = 0x0,
        .virt_start = 0x0,
        .size = 0x3b000000,
        .flags = JAILHOUSE_MEM_READ | JAILHOUSE_MEM_WRITE |
                JAILHOUSE_MEM_EXECUTE | JAILHOUSE_MEM_DMA,
    },
[...]
.pci_devices = {
    { /* VGA */
        .type = JAILHOUSE_PCI_TYPE_DEVICE,
        .domain = 0x0000,
        .bdf = 0x0008,
    },
[...]

```

## Jailhouse in the pICASSO project

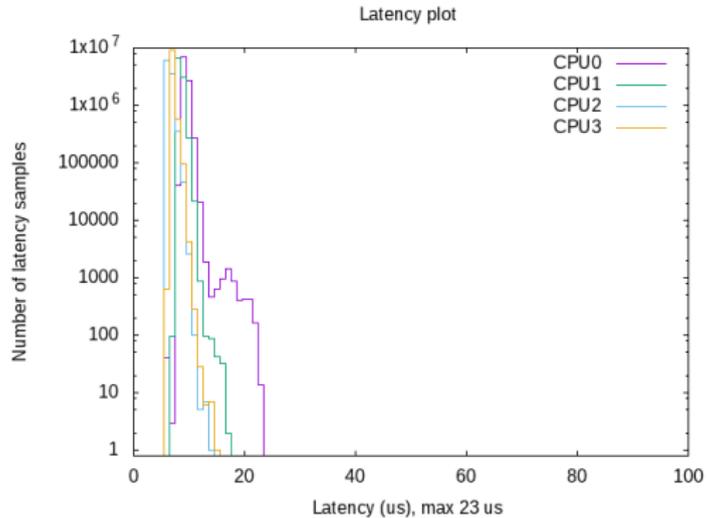
- ❏ Demonstrator at the ISW
- ❏ Connected to well known components from the "IT world" (libvirt)
- ❏ Participation in the Jailhouse project

## Results: Timertask (Linux Guest with PREEMPT\_RT)



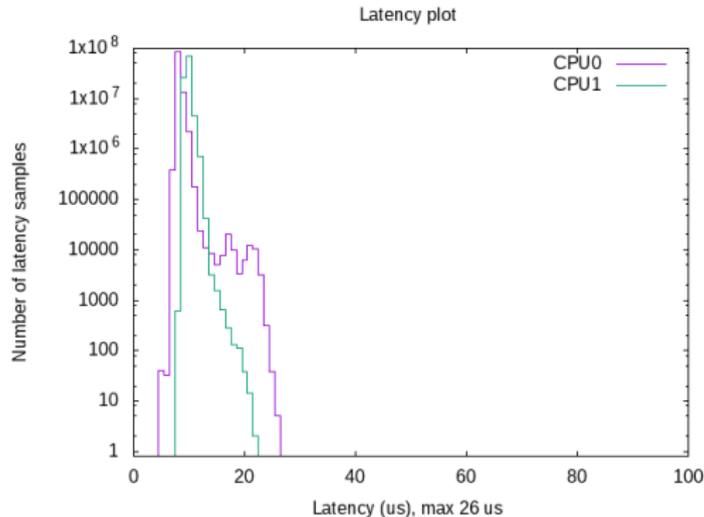
## ARM: Timertask (Zynq Ultrascale)

Native Linux system:



## ARM: Timertask (Zynq Ultrascale)

Within Jailhouse Guest (2 Cores used by the guest cell):



## Attention!!

- ❏ Full partitioning is not always possible (e.g. shared caches amongst CPUs)
- ❏ Using a hypervisor for isolation DOES NOT replace a security concept
- ❏ Just remember "Meltdown" and "Spectre" (exploiting the host from a guest is possible)

## Conclusion

- ❑ Hypervisor technologies can be used for: Consolidation and separation
- ❑ Jailhouse is completely Open-Source
- ❑ Jailhouse scales from a simple ARM CPU (with VT extensions) up to big server machines
- ❑ "mixed critical" scenarios are possible (concept discussed with TÜV Rheinland)
- ❑ Be careful: The isolation of different components doesn't replace a security concept

**Vielen Dank für Ihre  
Aufmerksamkeit**

**Linutronix GmbH**

**Bahnhofstraße 3  
88690 Uhltingen-Mühlhofen**

