XDF XILINX DEVELOPER FORUM

Heterogeneous Real-Time SoC Software Architecture

Presented By

Stefano Stabellini
Principal System Software Engineer

XILINX

Introduction

- **Stefano Stabellini**
  - Xen Project:
    - Founder of the Xen on Arm effort in late 2011
    - Xen on ARM Maintainer and Committer, Linux Maintainer
    - Develops Xen Project features on Zynq UltraScale+ MPSoC
  - Xilinx:
    - System Software Architect focusing on heterogeneous systems
    - Upstreaming Xilinx support to Xen and OpenAMP projects

Virtualization Basics

Virtualization – The Concept

- **“Virtualization”**
  - *The act of creating a virtual version of something, including virtual computer hardware platforms, storage devices, and computer network resources.*
  - Allows the deployment of multiple operating systems and independent workloads on one or more processors
- **“Hypervisor”**
  - *A hypervisor or virtual machine monitor is computer software, firmware or hardware that creates and runs virtual machines.*

- **Why Virtualize?**
  - OS/Workload consolidation
  - Lower system cost
  - Lower power consumption
  - Improved resource utilization
    - Mixed Criticality Systems
  - Fault tolerance
  - Multi-tenancy
  - Portability

Why Virtualize?

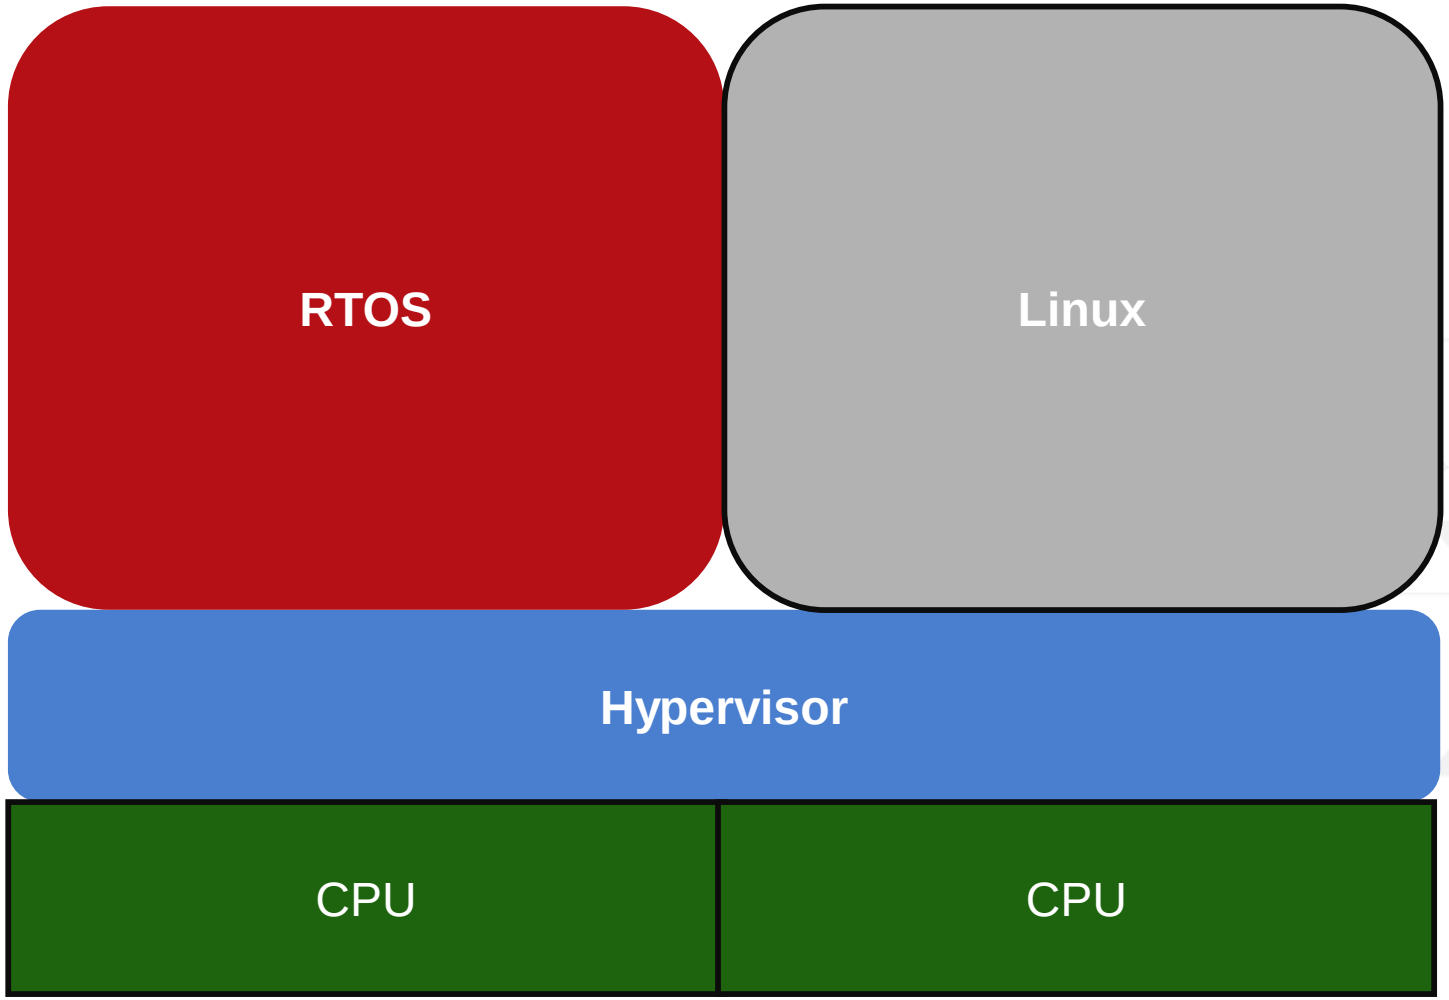

Why Virtualize?

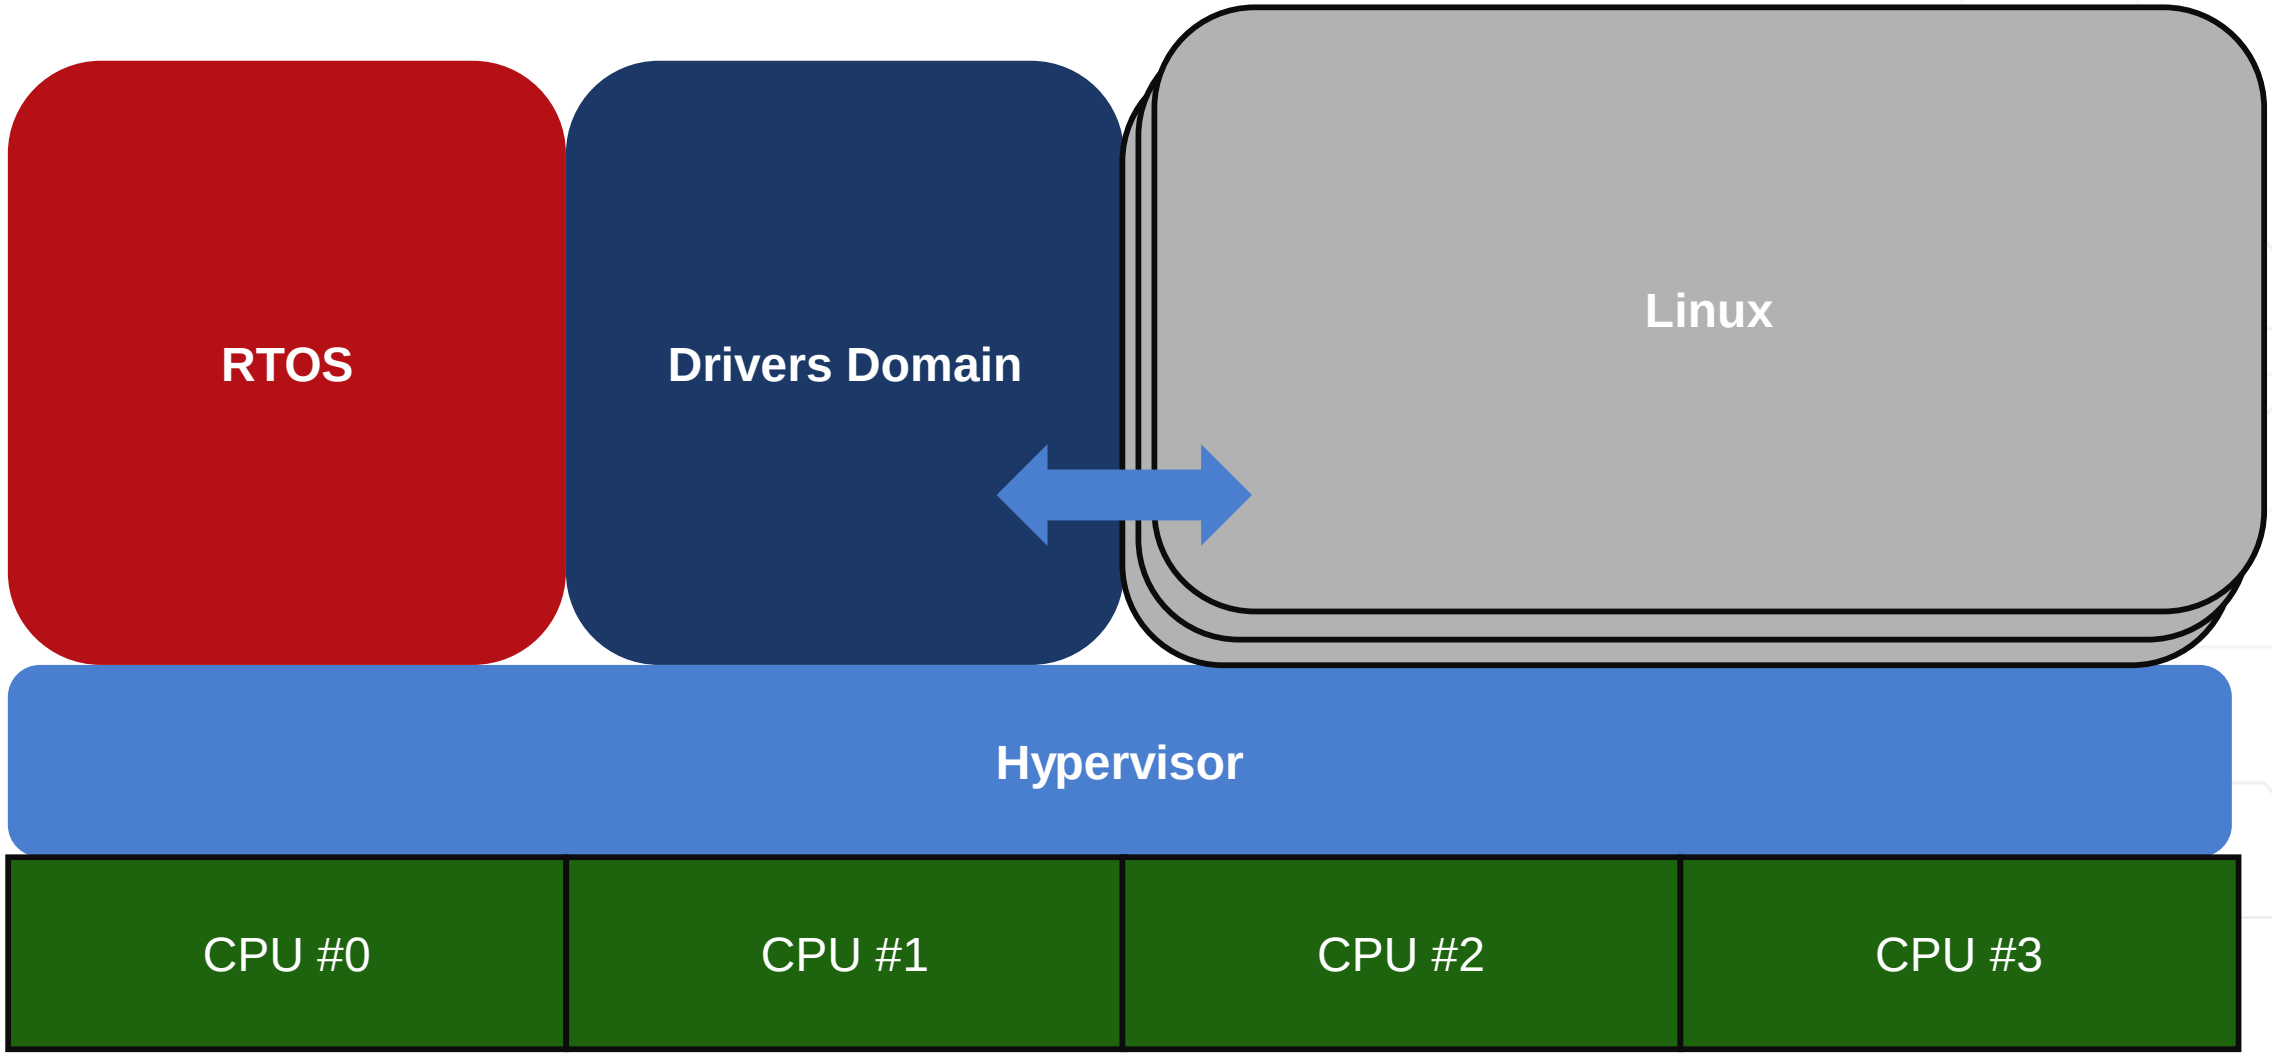

Embedded Hypervisor Requirements

- **Short Boot Times**
- **Real time**
  - Low, deterministic IRQ latency
  - Real time schedulers
  - Static CPU partitioning
- **Device Virtualization**
  - Device Assignment
  - Device Sharing
  - Driver Domains
  - VM to VM communication
- **Security, Isolation and Partitioning**
  - Memory
  - Devices
  - CPU
  - SLCRs
- **Operating System Support**
  - Linux, bare-metal, other RTOS support
- **Certifications**
  - Small code base
  - Type-1

Xen Project

Xen Project

- **Xen Project**
  - Open source hypervisor
  - Small code base implementing a micro-kernel design
  - Xen Project hosted by the Linux Foundation

- **Broad, Customizable Feature Set**
  - From servers to embedded
  - Out of box “real time” schedulers and enhancements
  - Advanced device management, partitioning, assignment
  - Independent user, control, and driver domains

- **Linux, BSDs or other OSes used for bootstrap (dom0)**
  - Linux is the most widely used but other OSes are possible

Example Xen Architecture

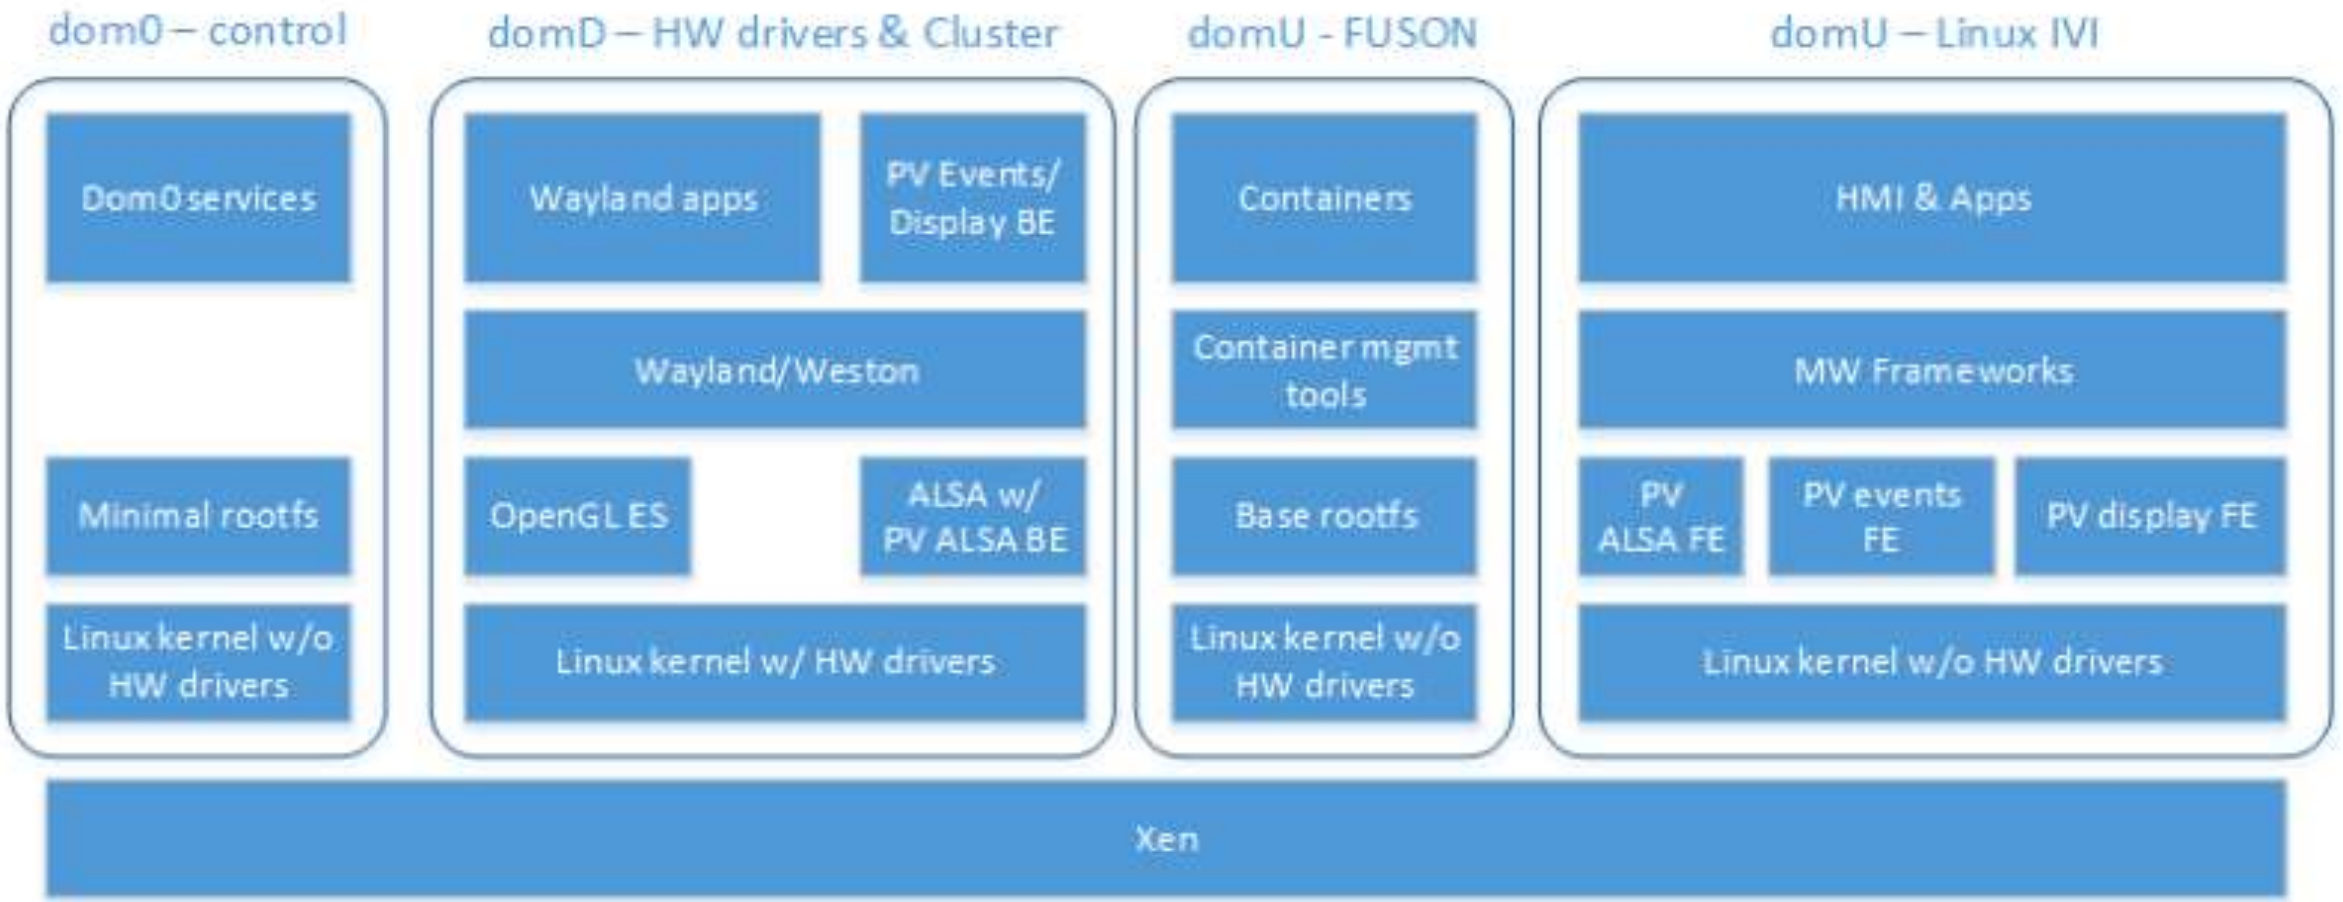

Xen Project 4.11

- **Highlights**
  - Regression testing and hardware validation completed successfully
  - Enormous work for the Meltdown and Spectre mitigations
  - Configurable SErrors handling
  - Many reliability fixes, especially in the interrupt handling path (GIC, vGIC)
  - SMCCC 1.1

- **Highlights (cont.)**
  - RTDS scheduler improvements
  - "null" scheduler improvements: tracing, soft affinity
  - VPL011
  - Mem_Access improvements
  - new PV Drivers: PV Display, PV Audio, PVCalls, PV 9pfs

- **Features and Status**
  - Xen Project 4.11 Feature List

Mem_Access

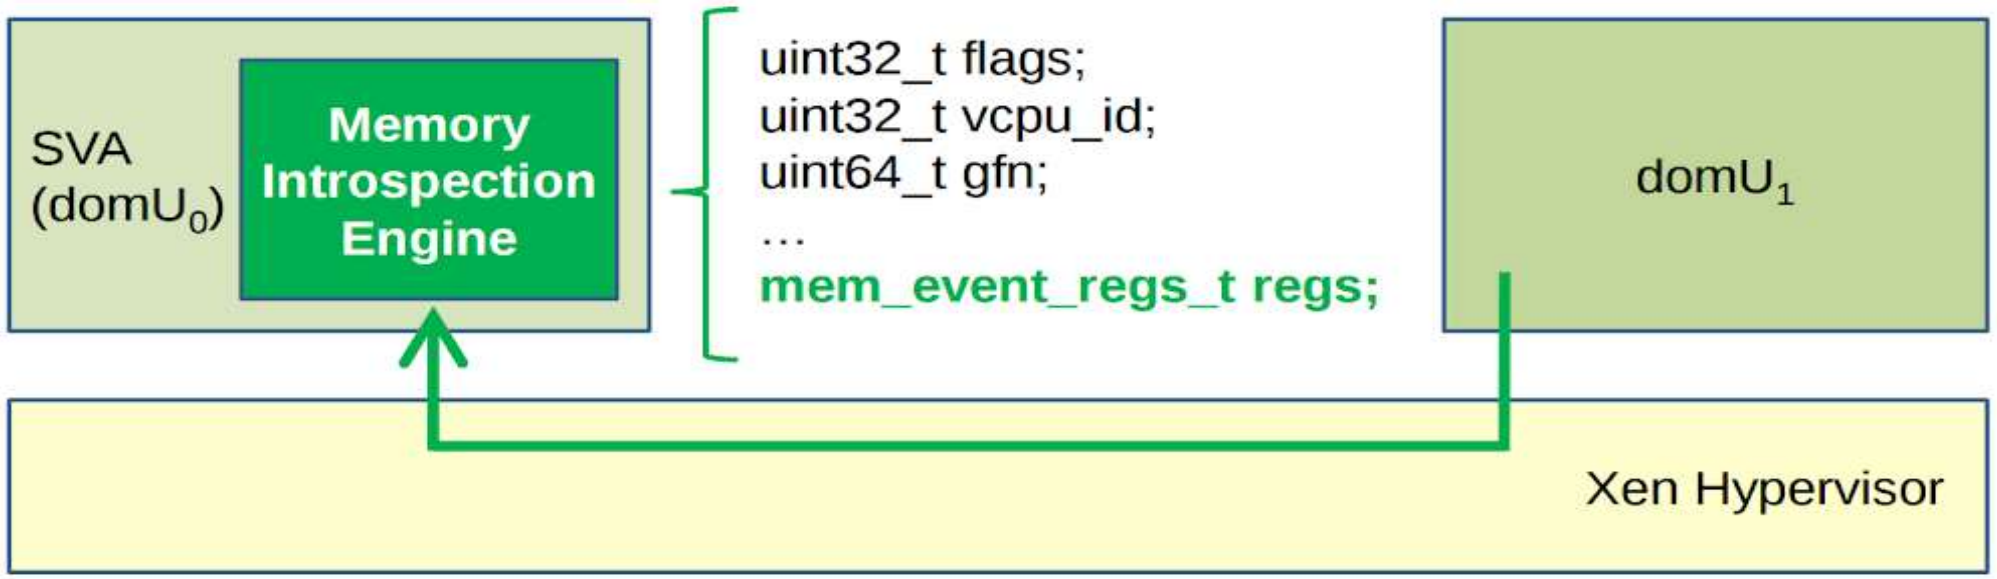

PV Drivers

- Existing: net, block, console, keyboard, mouse, framebuffer
- New: 9pfs, PVCalls, Multi Touch, Sound, Display
- Prerequisites: xenstore, grant table and event channels support (BSD code available)

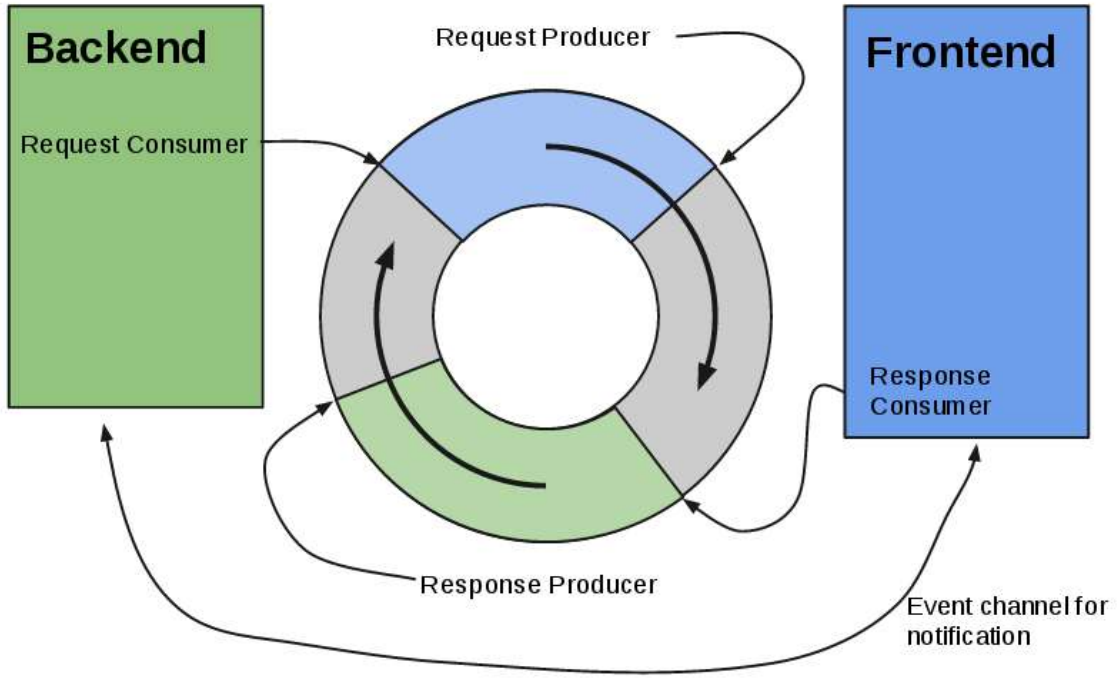

Static Partitioning Use-Case

sched=null vwfi=native

Static Partitioning Use-Case

sched=null vwfi=native

2.5 us

Static Partitioning Latency

Xilinx Zynq Ultrascale+ MPSoC
Physical Timer

Xen with phys_timer patch
vwfi=native

dom0_mem=1G
max_dom0_vcpus=2
1 vcpu TBM ctest

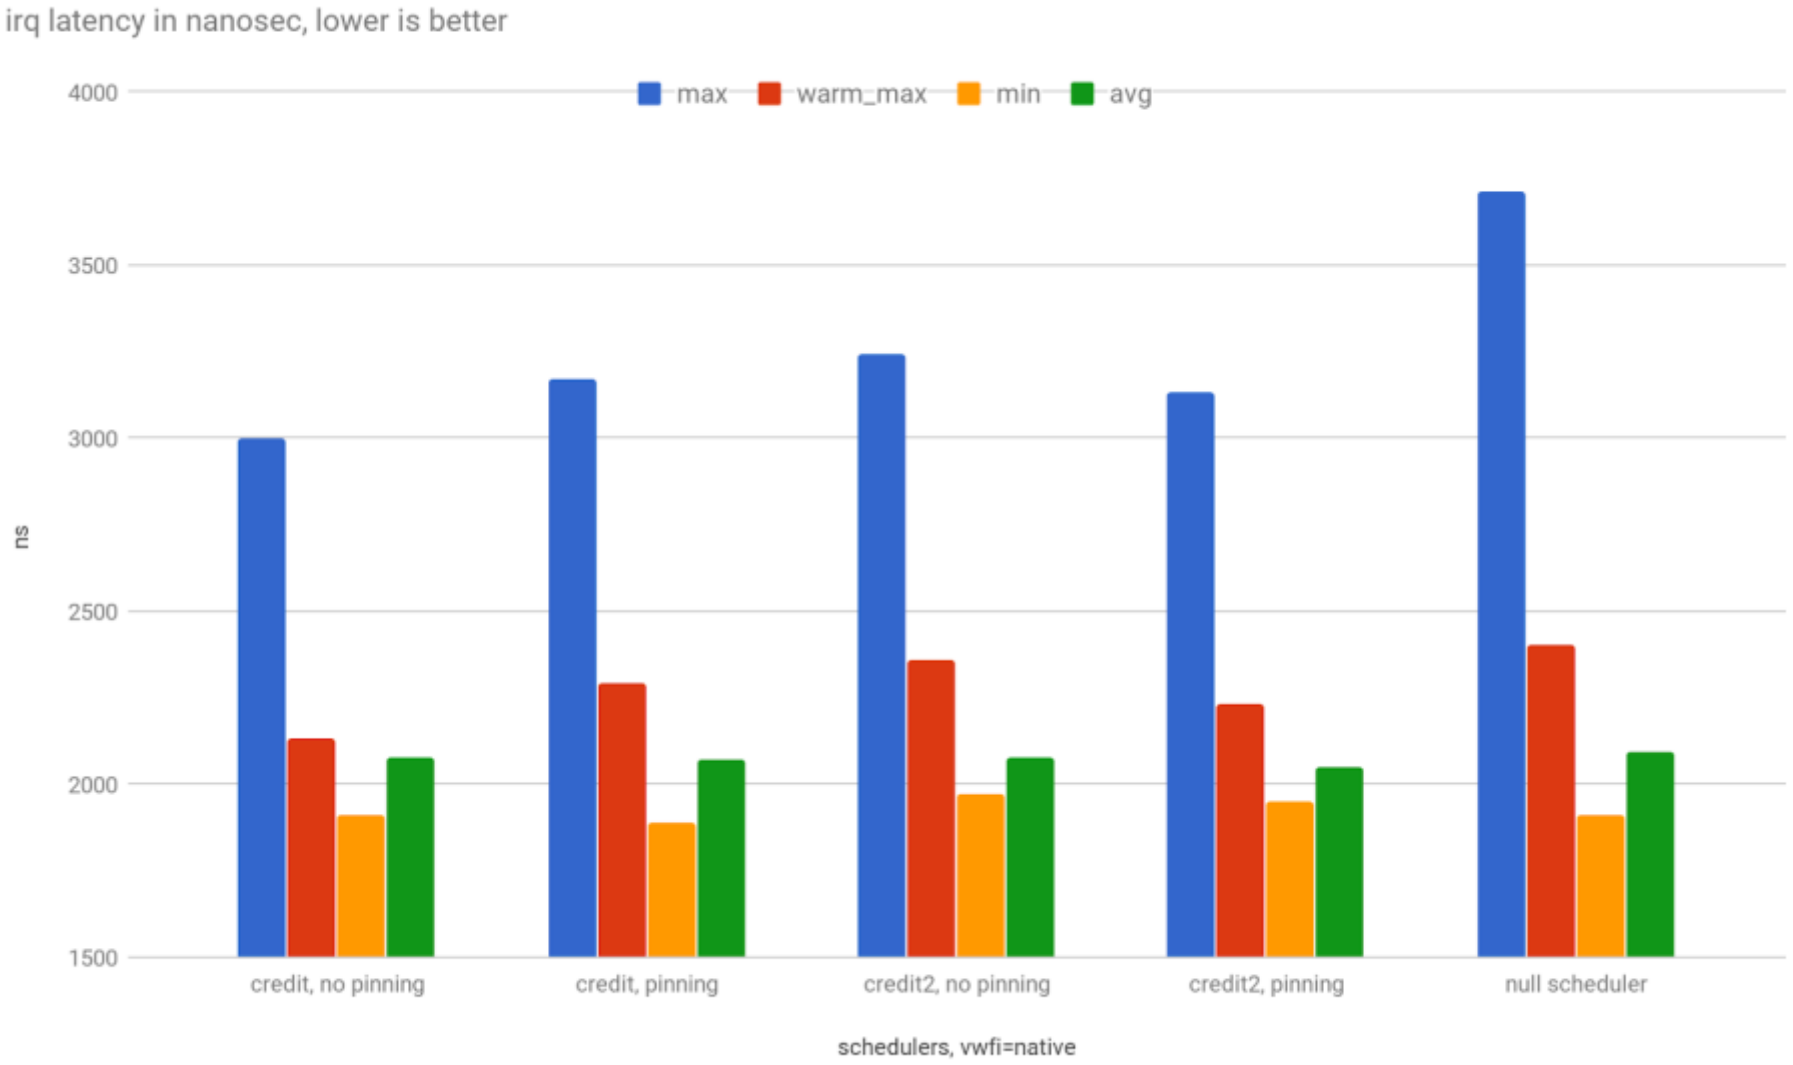

Xen Schedulers

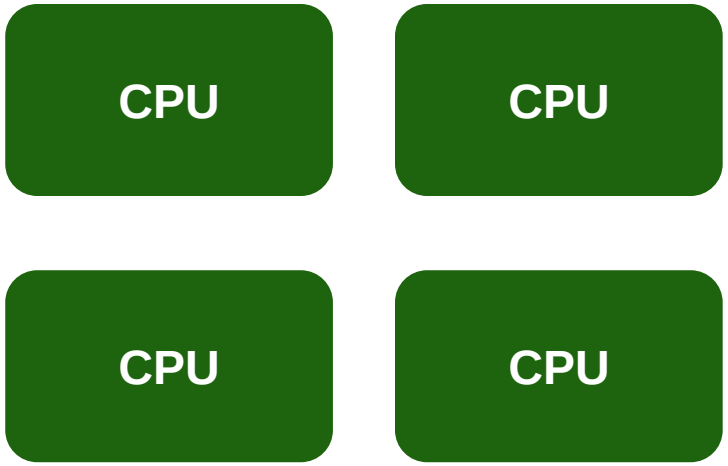

Xen Schedulers

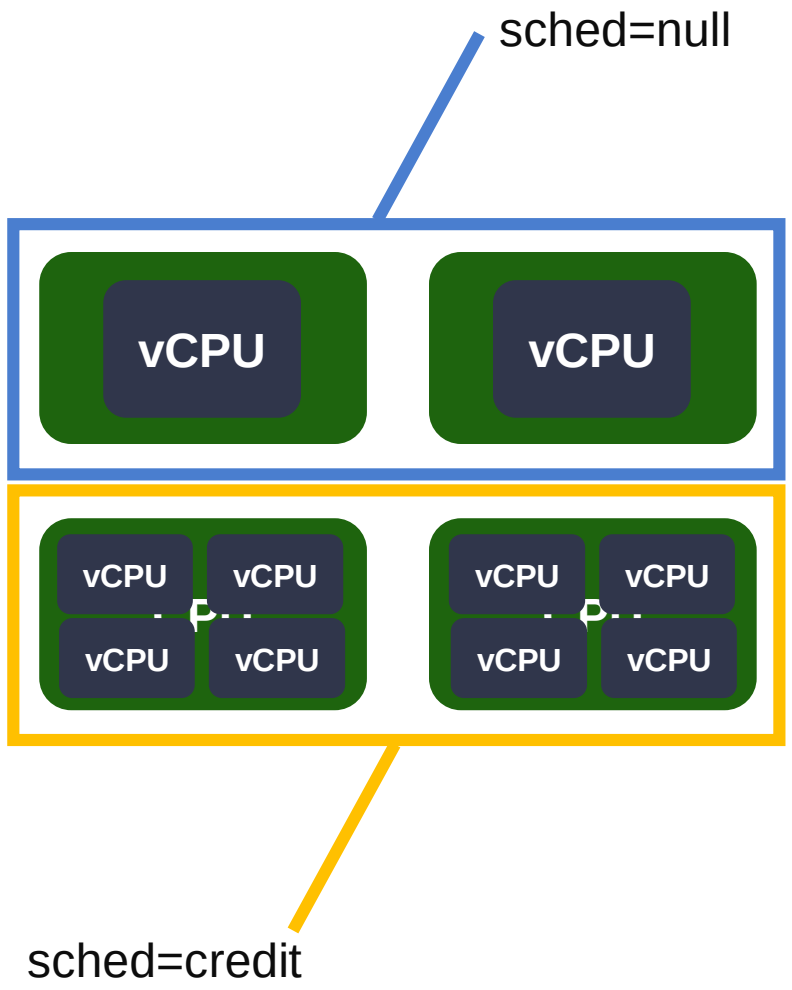

Xen VM-to-VM communication mechanisms

- **Libvchan**
  - Linux library
  - Direct VM to VM communication channel based on a ring on shared memory
  - libxenvchan_send and libxenvchan_recv
- **PVCalls**
  - Socket API virtualization
  - VM to VM communication mediated by the backend domain (typically dom0)
  - "lo" becomes a inter-VMs communication namespace
- **V4V**
  - Linux library and hypercall, kernel space and user space
  - VM to VM communication mediated by Xen
  - Trivial to implemented on your own kernel
  - Not fully upstream

Brand New Features

Introduction Slide

Shared Memory

- **Completely Configurable**
  - Support any memory attributes, including cacheable memory (default)
- **No need for Xen support to use it**
- **Can export the memory to Linux userspace and use OpenAMP**

```
static_shm = ["id=ID1, begin=0x40000000, size=0x1000, role=master"]
```

```
static_shm = ["id=ID1, offset=0, begin=0x48000000, size=0x1000, role=slave"]
```

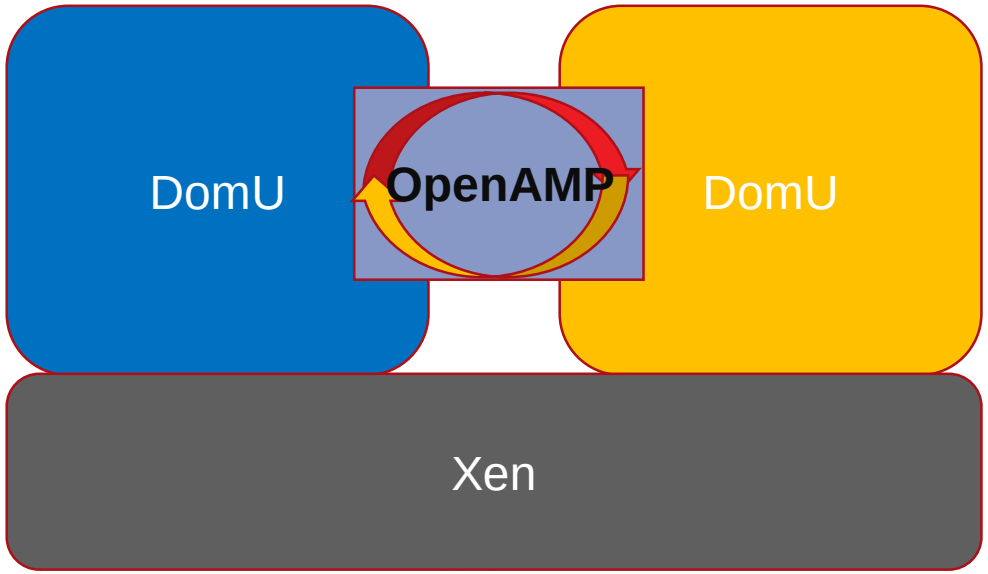

Reducing Code Size

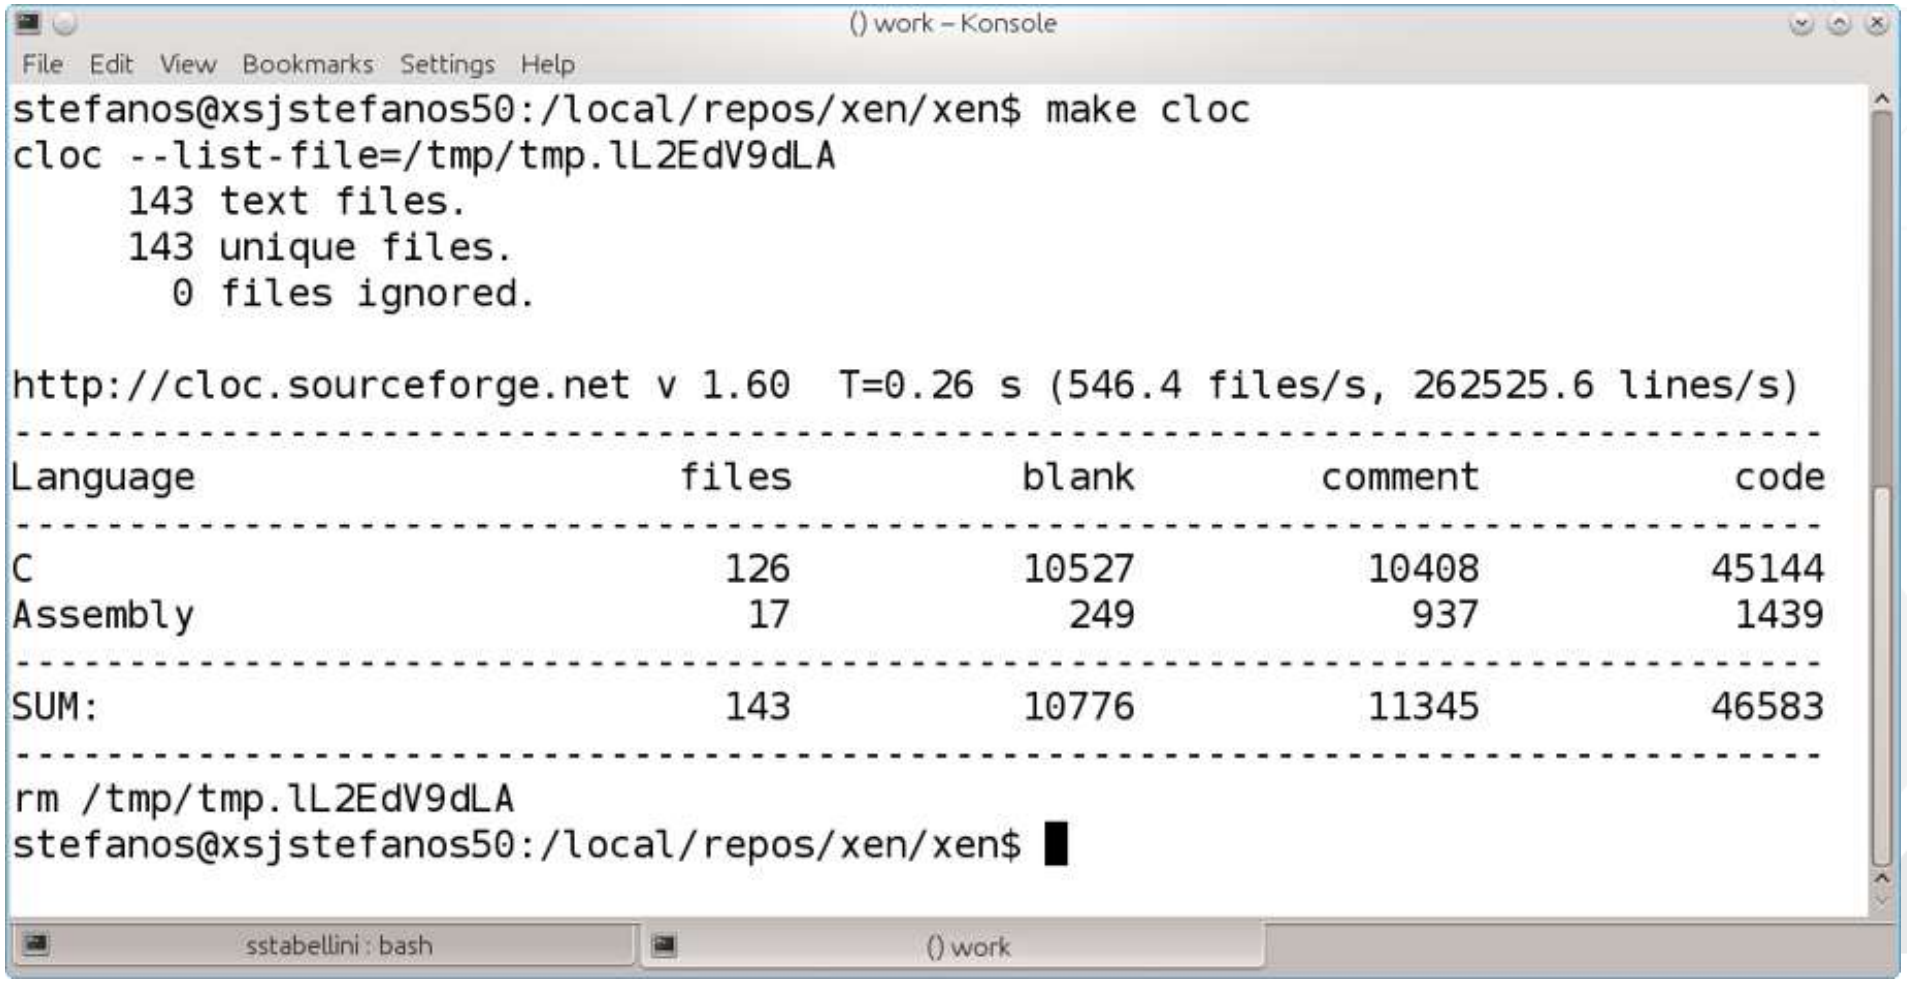

Certifications

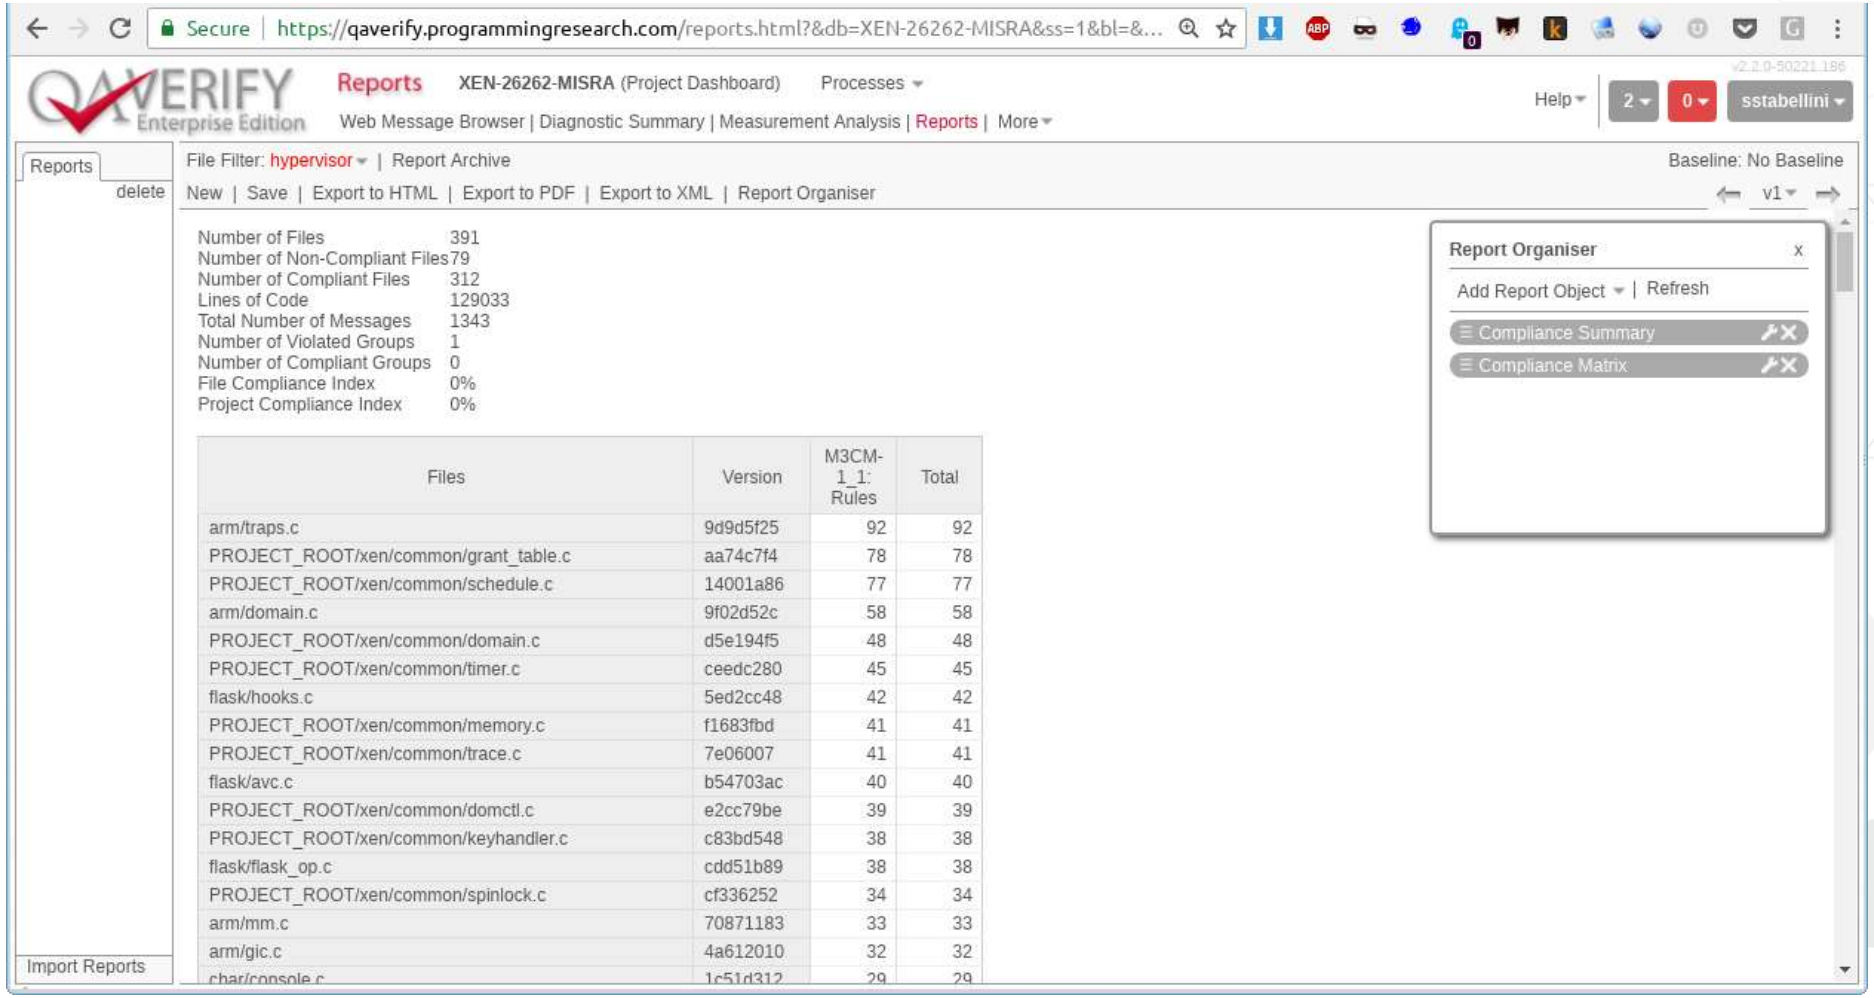

Files	Version	M3CM-1_1: Rules	Total
arm/traps.c	9d9d5f25	92	92
PROJECT_ROOT/xen/common/grant_table.c	aa74c7f4	78	78
PROJECT_ROOT/xen/common/schedule.c	14001a86	77	77
arm/domain.c	9f02d52c	58	58
PROJECT_ROOT/xen/common/domain.c	d5e194f5	48	48
PROJECT_ROOT/xen/common/timer.c	ceedc280	45	45
flask/hooks.c	5ed2cc48	42	42
PROJECT_ROOT/xen/common/memory.c	f1683fbd	41	41
PROJECT_ROOT/xen/common/trace.c	7e06007	41	41
flask/avc.c	b54703ac	40	40
PROJECT_ROOT/xen/common/domctl.c	e2cc79be	39	39
PROJECT_ROOT/xen/common/keyhandler.c	c83bd548	38	38
flask/flask_op.c	cdd51b89	38	38
PROJECT_ROOT/xen/common/spinlock.c	cf336252	34	34
arm/mm.c	70871183	33	33
arm/gic.c	4a612010	32	32
char/console.c	1c51d312	29	29

Dom0-less

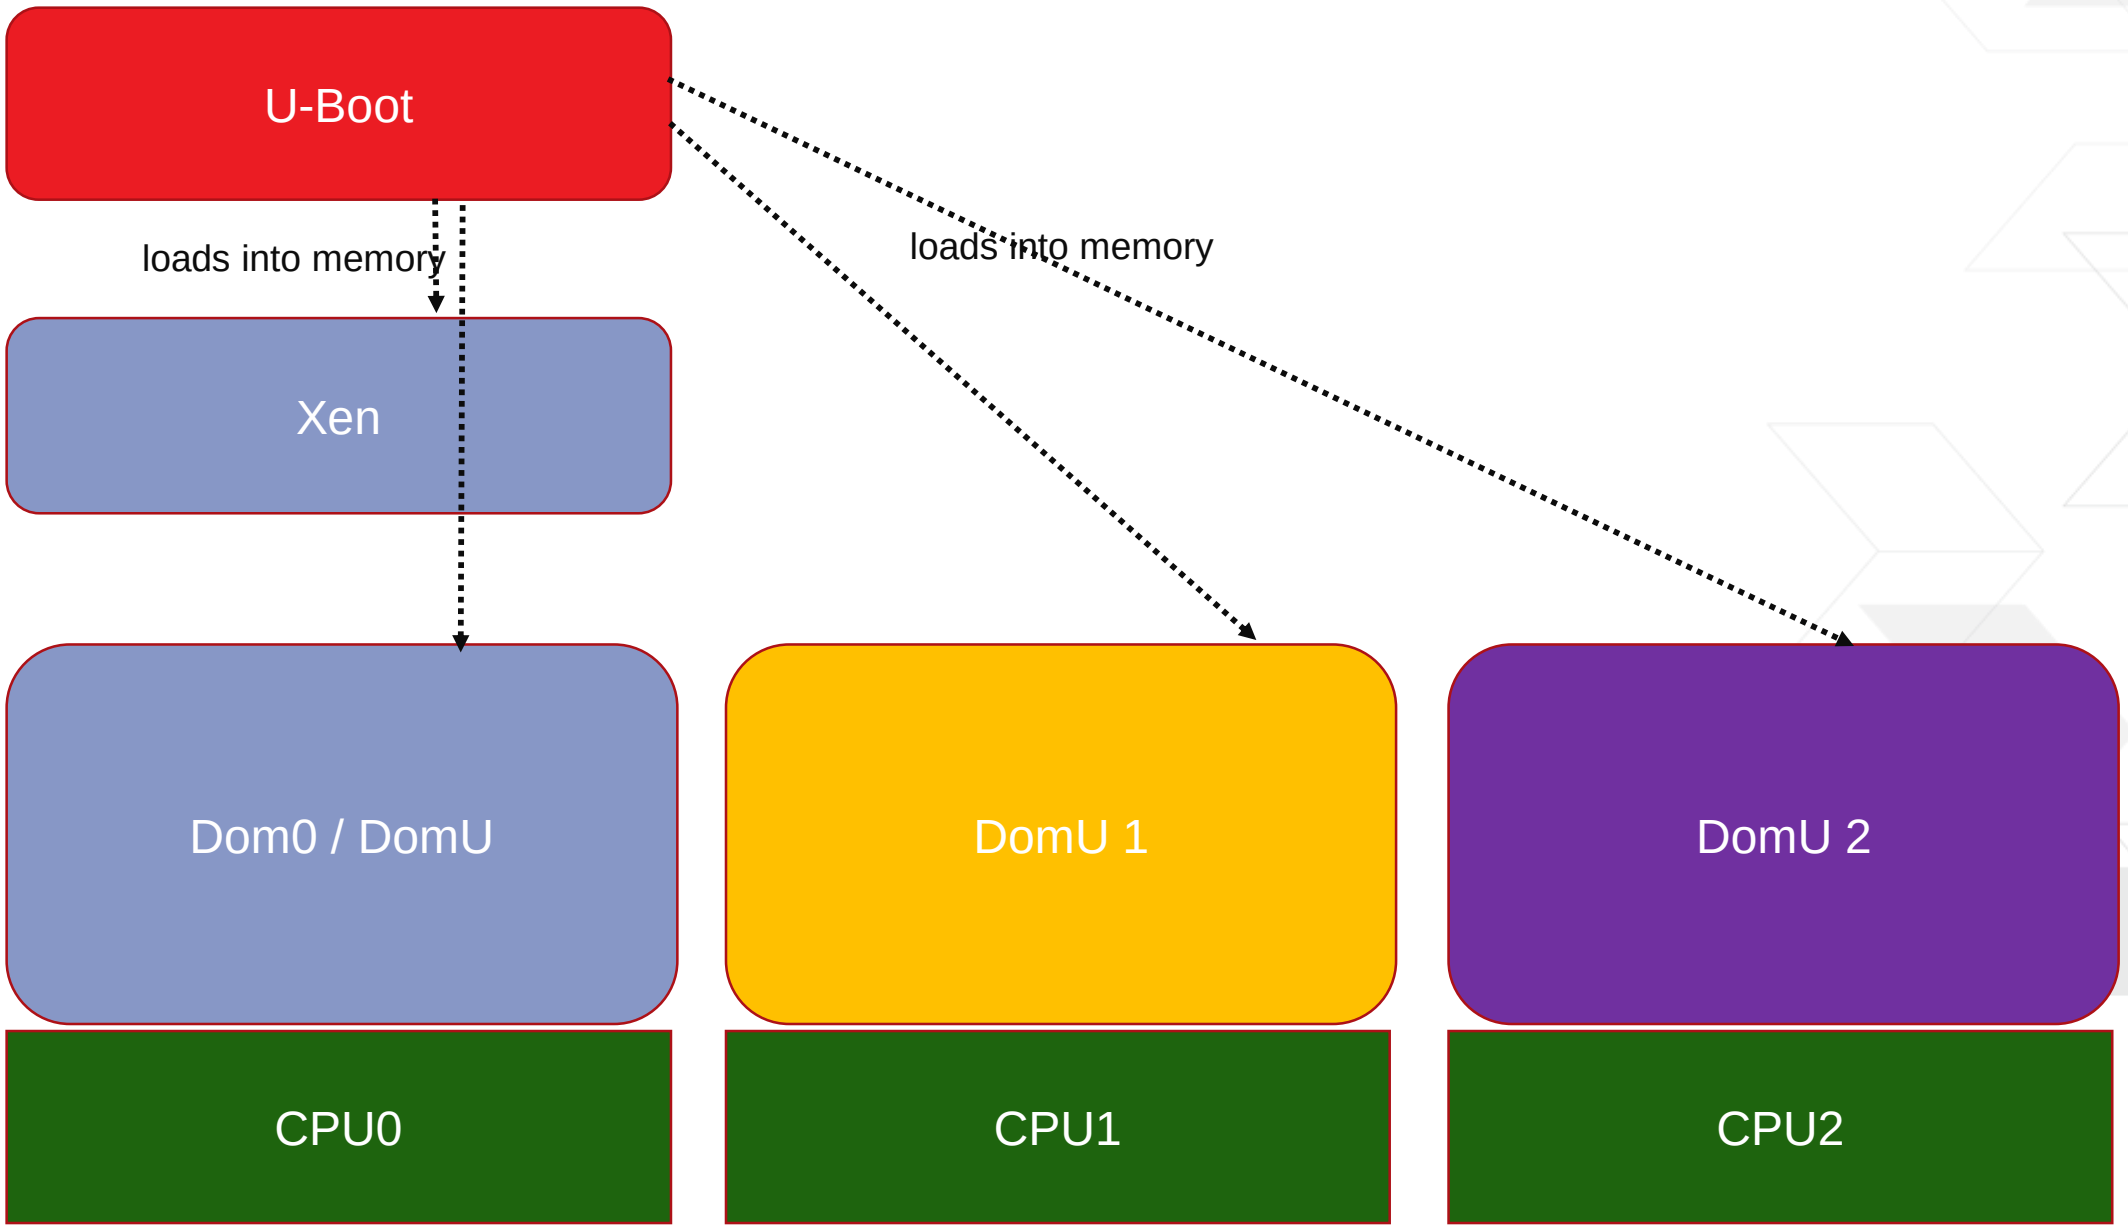

Dom0-less

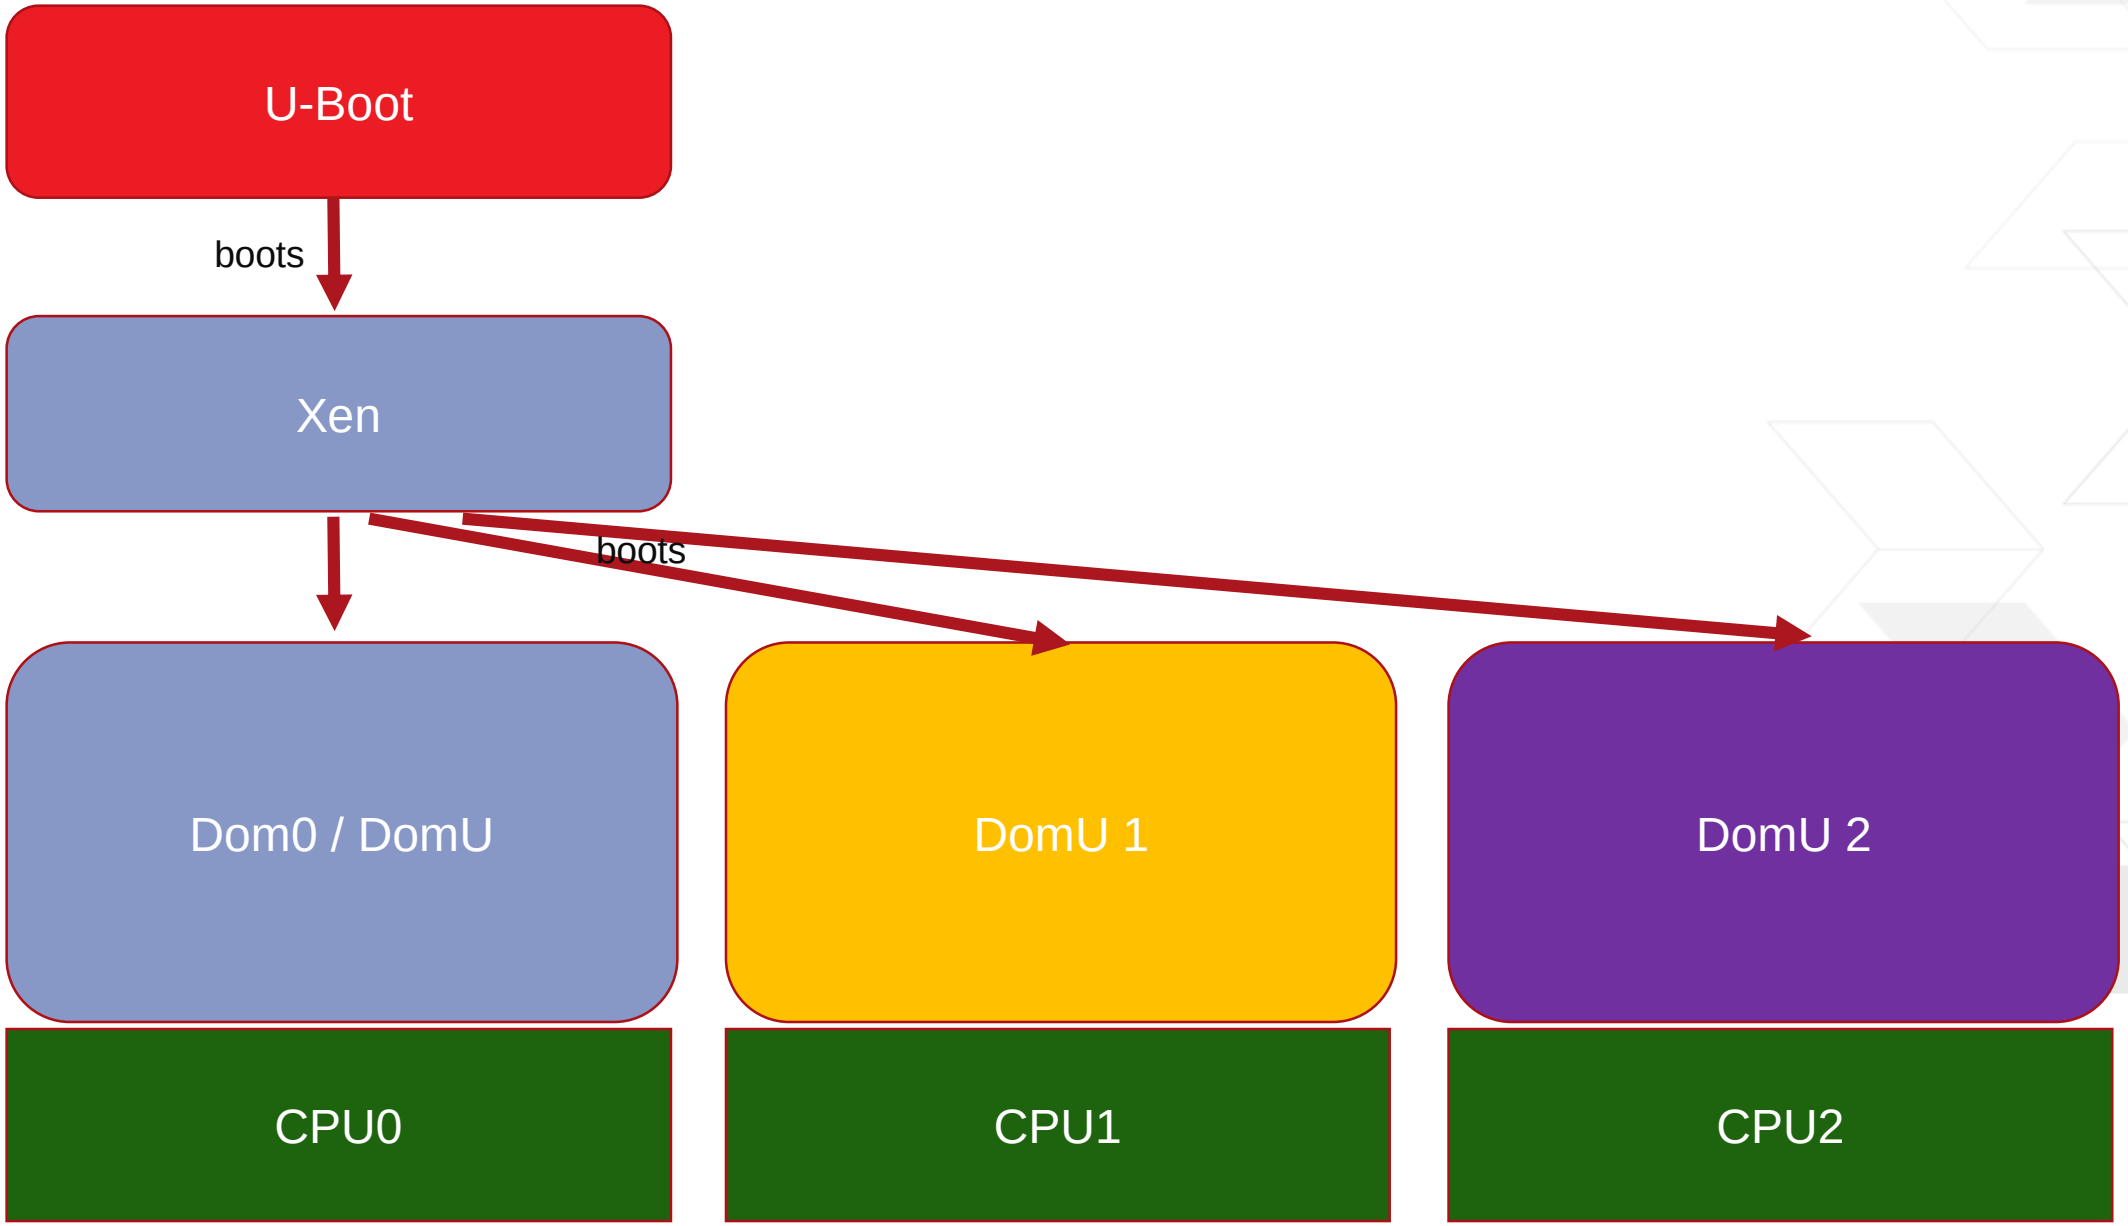

Xen Project "OSSTests"

- **OSSTests: Xen Project official CI-loop**
  - Run 24/7
  - Commits move to master only after passing the CI-loop tests
  - Based in Boston, MA
  - Only accept off-the-shelf hardware
- **Xilinx MPSoC ZCU102 coming to Xen Project!**
  - Will validate master on Xilinx hardware
  - Every Xen release will be checked against Xilinx hardware
  - Increase overall quality
  - Reduce risks of rebasing Xen in Petalinux

"The best security process in the industry"

- A very transparent process
- Responsible disclosure
- Only few security issues for Xen on ARM
- Xen stable trees maintained for security for 3 years

Commercial Xen Support

- **DornerWorks**
  - Xilinx Premier Design Services Partner
  - Hardware, software and systems expertise
  - Xilinx partner for Xen support and design customization services
- **Community Support**
  - Free Community Support is available to the entire Zynq UltraScale+ MPSoC community.
  - This support includes all software for Virtuosity™, plus all supported configurations or workflows that are documented by the distribution.
- **DornerWorks Xen commercial support**
  - Custom hardware porting
  - New guest OS support
  - Custom device drivers
  - Programmable Logic integration
  - System architecture design
  - Scheduling and partitioning for ARINC 653 and FACE
- **http://dornerworks.com/xen**

Other Hypervisors

Jailhouse

- **Open source hypervisor**
  - https://github.com/siemens/jailhouse
- **Lightweight implementation**
  - Focus on resource partitioning and not on virtualization
    - No schedulers, no PV devices, no Driver Domains, etc.
- **Features**
  - Optimized for simplicity rather than feature richness
  - Relatively new ARM64 support
- **Linux used for bootstrap and control of partitions**
- **Commercially supported on Zynq UltraScale+ MPSoC by Enea**

Commercial Hypervisors

- DornerWorks (Xen, seL4)
- General Dynamics Mission Systems (OKL4 Microvisor®)
- Green Hills Multivisor®
- Lynx LynxSecure®
- Mentor Embedded Hypervisor
- BlackBerry QNX® Hypervisor
- Sysgo PikeOS® Hypervisor
- Wind River Virtualization Profile