Intro to QEMU

Alex Bennée, Linaro

accellera SYSTEMS INITIATIVE™

SYSTEMC™ EVOLUTION DAY
OCT 17, 2024 | MUNICH | GERMANY

What is QEMU?

- “A generic and open source machine emulator and virtualizer”
- 21 years old
- 19 guest architectures including
  - ARM
  - Hexagon
  - Loongarch
  - Risc V
  - PowerPC
  - s390x

How is QEMU?

- Collaboratively developed online
  - qemu-devel@nongnu.org
  - https://gitlab.com/qemu-project/qemu
  - #qemu on OTFC IRC!
- In the last year
  - 270 kloc of code
  - 8157 commits
  - 471 developers
  - 254 employers

Size of QEMU (loc)

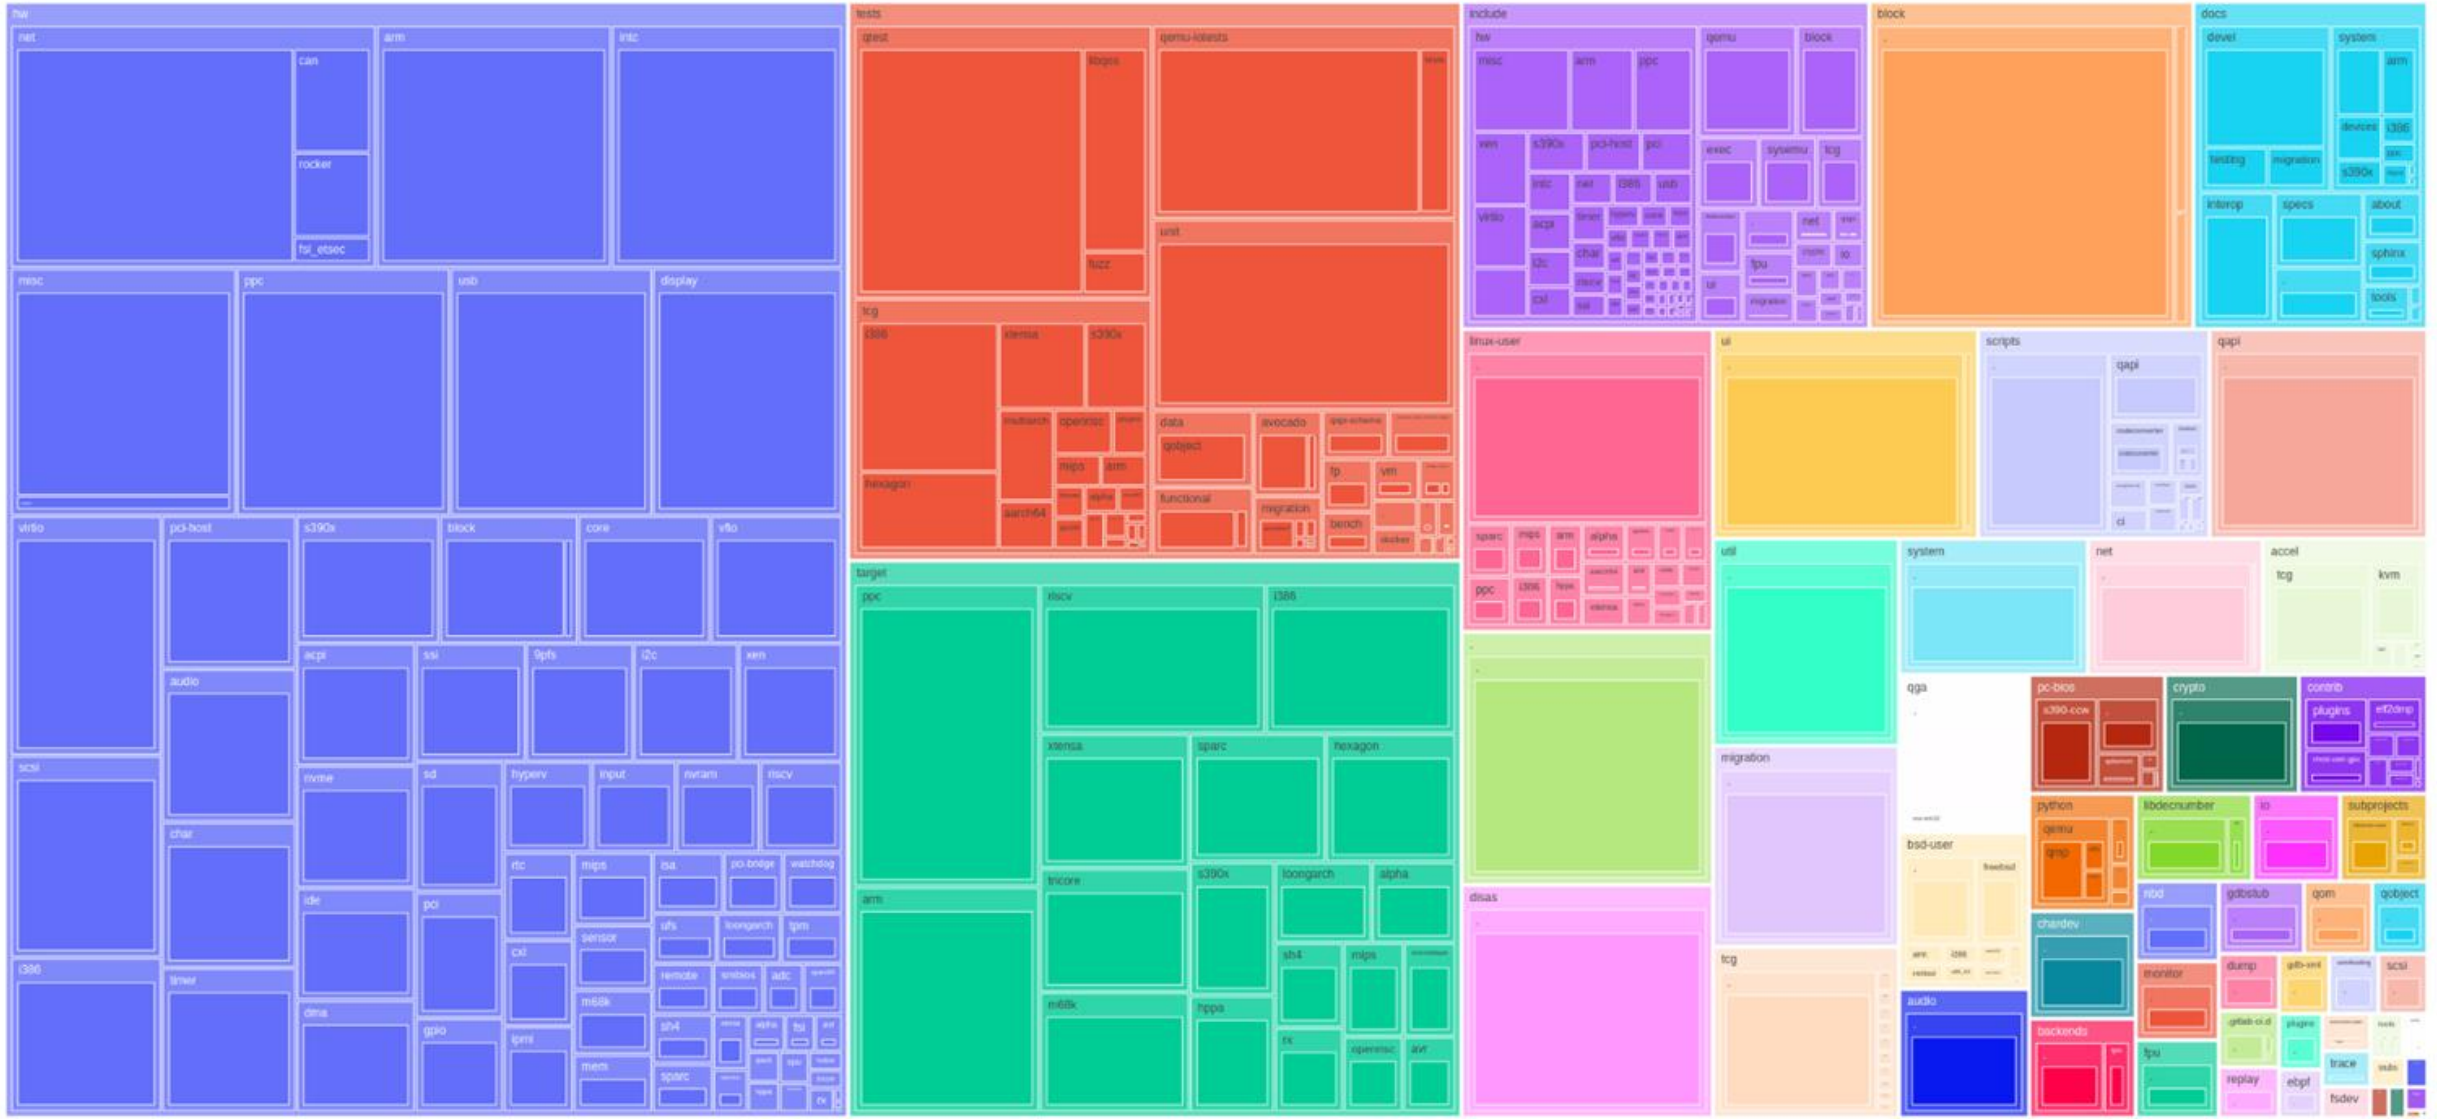

QEMU System Emulation Features

- Wide variety of board and device models
- Pause/Resume/Save/Restore/Migrate
- Flexible block, network and character backends
- QEMU API (QAPI)
- gdbstub
- Record/replay
- Instruction instrumentation framework (TCG plugins)

Threading Model

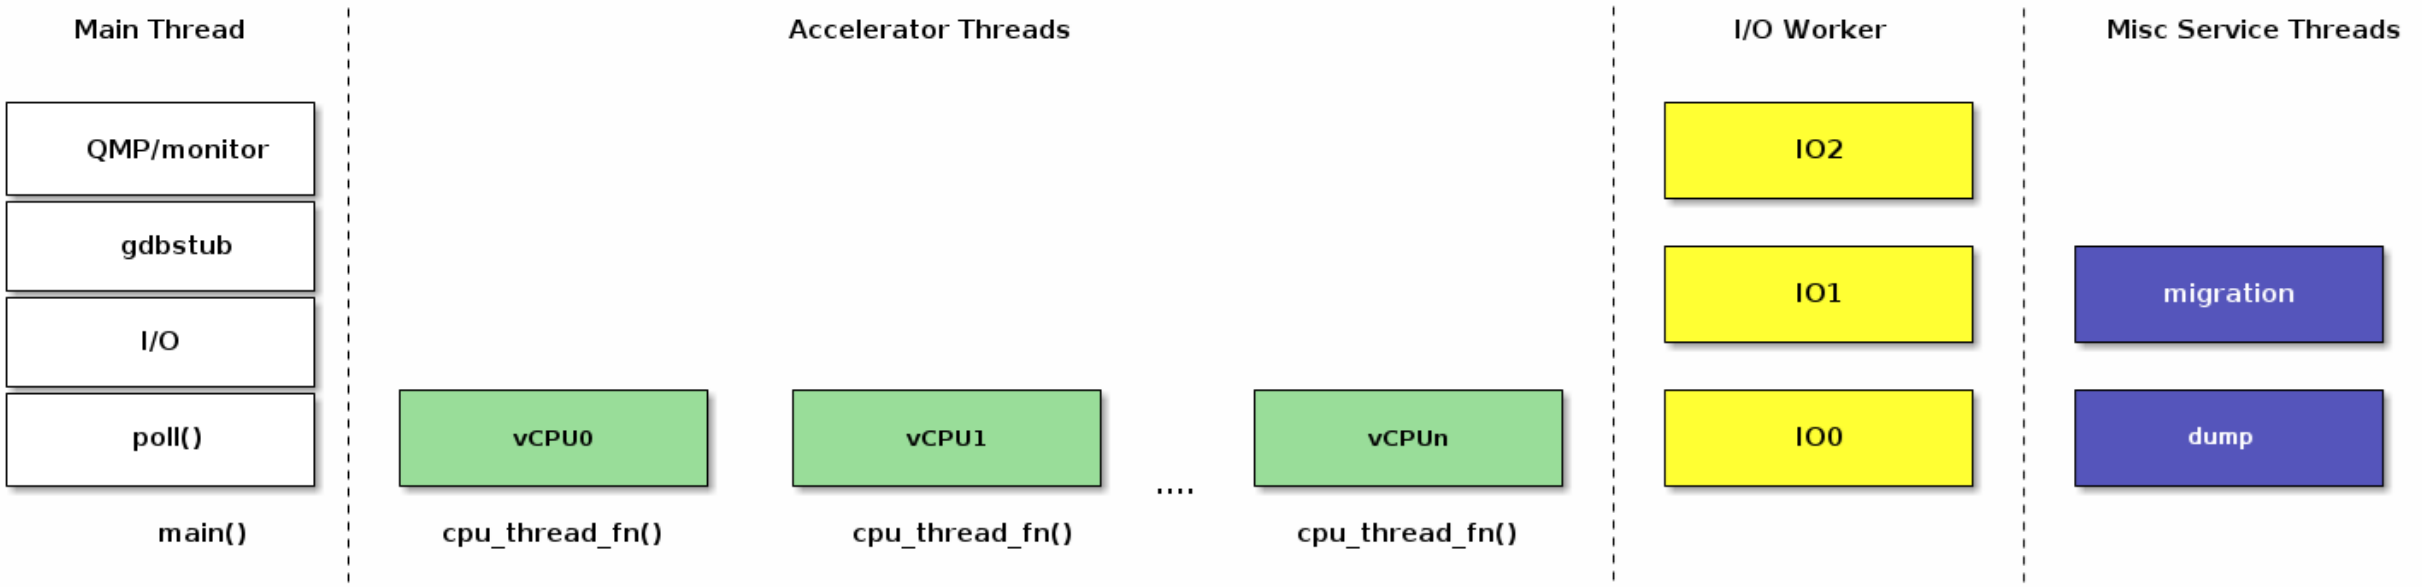

QEMU user-mode

- Faster than full system emulation
- Normally for userspace binaries with libc
- Can run nostdlib binaries with semihosting
  - useful for testing

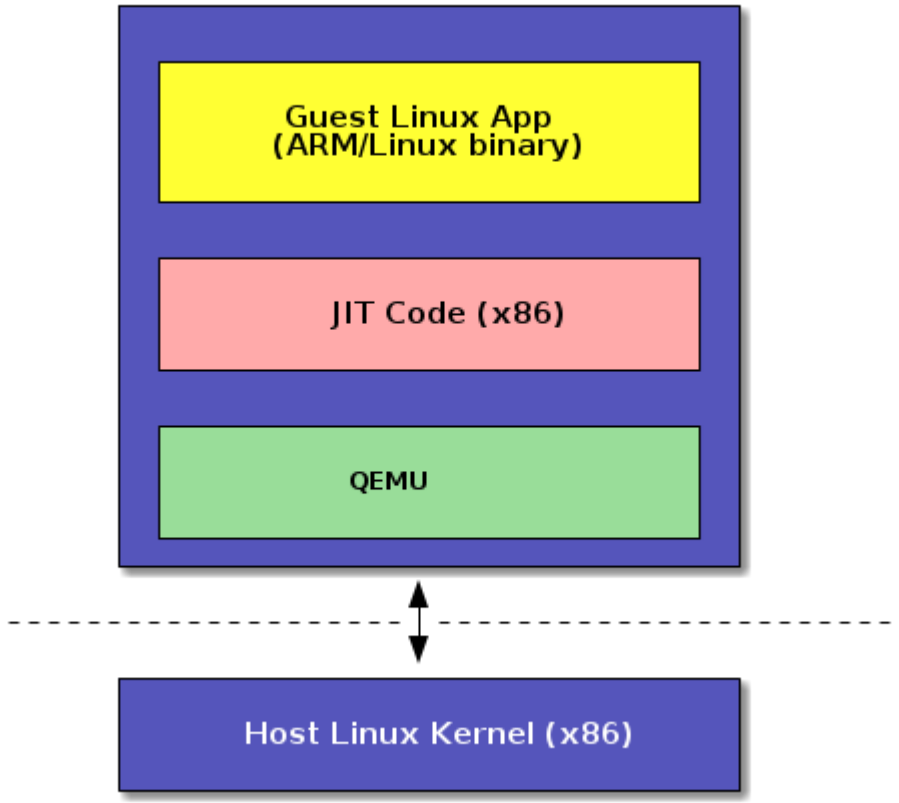

QEMU does not…

- Have a carefully thought out internal API
- Care about micro-architecture details
- Think too hard about time
- Currently work with SystemC

QEMU Internal API History

- Organic development
- Incomplete conversions
- Copy and tweak

Micro-architecture details

- Costly
- Off-load to plugins
- Cache modelling
- Driving simulation
  - TCG Plugin in Practice: A Case of Microarchitecture Research (KVM Forum 2024)

Time

- CLOCK_[REALTIME|VIRTUAL|HOST|VIRTUAL_RT]
- Wall-clock time
  - effect of emulation speed visible
- icount
  - Time == instructions executed
- TCG plugins
  - defer time to plugin

Example out-of-process Devices

- VirtIO (vhost-user)
- vfio-user (out-of-process PCI emulation)
- TPM
- ivshmem

Heterogeneous VirtIO Example

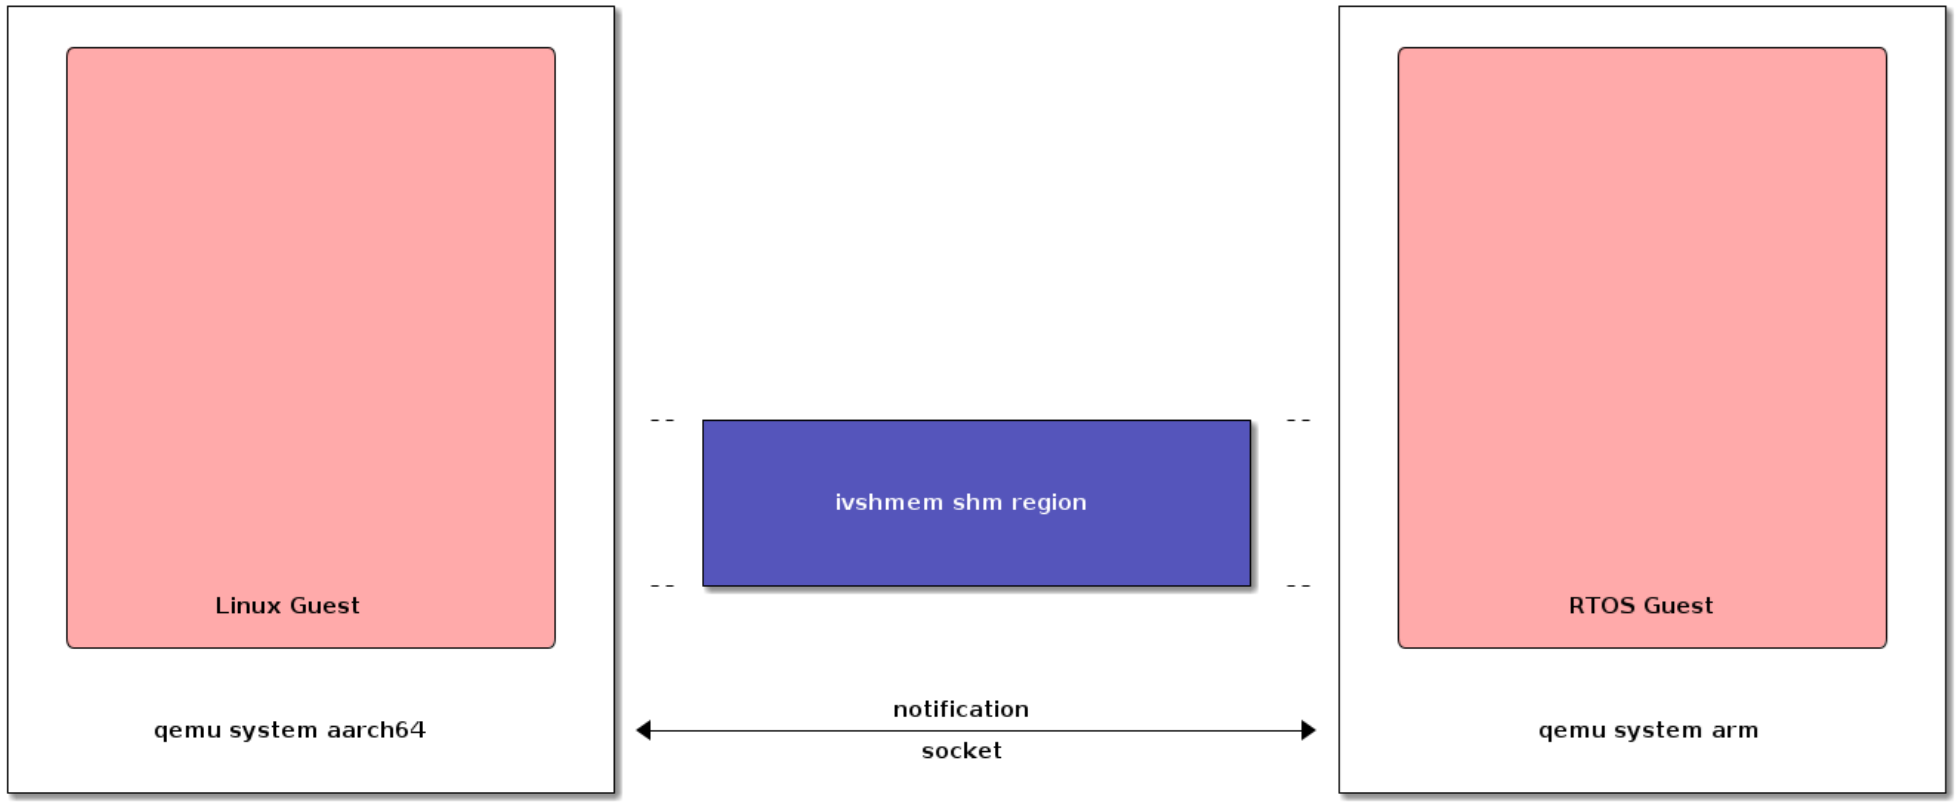

QEMU Downstreams

- Lots
  - not supported by upstream
  - face a continuous re-base battle
  - however show potential approaches
- Xilinx QEMU
  - heterogeneous modelling
  - utilises shared memory (SHM)
  - hwdtb machines
- Qualcomm QQVP
  - SoC modelling
  - glued with SystemC

Current upstream QEMU work

- Continued evolution of CPU models
  - confidential computing
  - ISA hardening extensions
- Heterogeneous Modelling
  - two different ISA vCPUs in the same process
- Dynamic machine types
  - beyond what “Virt” handles for hotplug
  - data driven board models
- Rust
  - currently experimental, community is keen

Any Questions?