QEMU - The Building Block of Open Source Virtualization

Glauber Costa
glommer@redhat.com

October 29, 2009

What's QEMU?

- A Code translator
- *NOT* cycle accurate.
- A System Emulator
- A niche-specific software, rapidly gaining attention under the spot

linux-user

- i386-user, x86_64-user, arm-user, whatever
- syscall mapping
- code translation (tcg)
- not that interesting for virtualization users

linux-user schematic

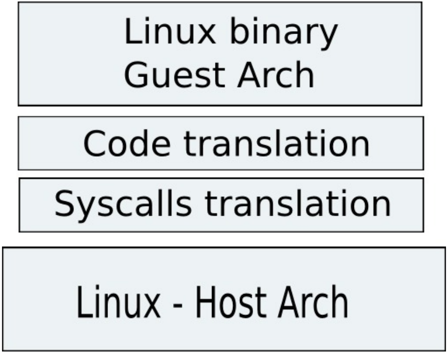

Linux binary
Guest Arch

Code translation

Syscalls translation

Linux - Host Arch

System Emulator

- Goal is to emulate a full machine
- PCI Bus, PCI devices, disk controllers, etc
- CPU.

qemu-system schematic

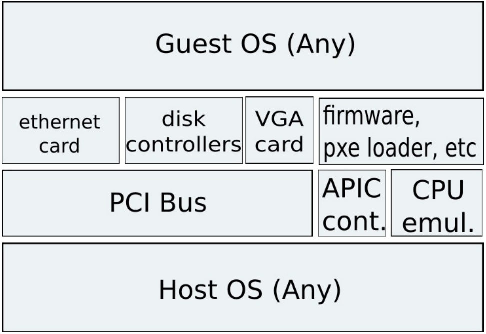

Virtualization

- First approach, emulate everything but the CPU
- Use of special devices, like virtio

Alternative CPU models

- Kqemu (thankfully dead)
- Xen, both PV and HVM
- KVM
- VirtualBox

Comparisons

- KVM: each cpu is a linux thread, linux schedules it: a lot of state in qemu's cpu
- Xen: have its own schedulers: just a few state in qemu's cpu

Qemu problems

- Qemu suffered from the commit access disease
- git was the cure
- Absurd lack of structure and patch review
- Version 0.9.1 lasted for very, very long: no useful bug reports from users

Led to... forking

- Patches were largely ignored, but life had to move on
- kvm, xen, maemo, had different forks
- Some forked last release, some forked svn
- kvm + xenner and linux user forks on its way to inclusion

new qemu people

- Many current qemu developers came from a linux kernel background
- Brings the kernel culture.
- kernel and qemu has 50 % of overlap in terms of developers (meaning more than half of qemu developers wrote something for the kernel)

What is missing from KVM front

- kernel based irqchip devices (i8259, APIC and IOAPIC controllers, etc)
- smp support

Questions?

Feel free.

Questions?

Feel free.