Konrad Rzeszutek Wilk Software Developer Manager Oracle

4.5 Roadmap



Copyright © 2014, Oracle and/or its affiliates. All rights reserved.

•Time vs Features

•Hypervisor changes for ARM, x86, AMD and Intel

Toolstack

•Linux, FreeBSD and MiniOS

•GRUB2



Release Manager/Coordinator History:

- 4.3: George Dunlap
 - Tracking blocking features and bugs
 - Guiding community discussions
- 4.4 : George Dunlap
 - All above and:
 - Updates on xen-devel + Wiki
 - Roadmap schedule
 - Bugs after feature freeze, clamping down features
 - Process documented

Release Manager/Coordinator History (cont.):

- 4.5 : Konrad Rzeszutek Wilk
 - All above
 - Performance regression testing
 - Fill out the process documentation as needed
 - Non Citrix employee



Xen Project 4.5: When: Time vs Feature

- 9 months release (4.4 was aiming to be 8 months, came out to be 9):
 - September 10 feature freeze
 - October 10th RC1
 - December 10th release
- Other codebases related to release:
 - Hypervisor: ARM and x86
 - Toolstack (libxl, libxc), libvirt, QEMU
 - -Linux, FreeBSD, QNX, MiniOS, rump
 - -GRUB

Xen Project: hypervisor: ARM Support for ARM IP and other standard interfaces:

• GICv3

- MSI, Power Management support (PSCI)
- Support for more than 8 CPUs
- Interrupt Translation (ITS)
- Emulation of GICv2 in guests
- Power State Coordination Interface (PSCI)
 - CPU down/up/suspend, reset, affinity, migrate
- UEFI booting
- IOMMU support (SMMU)
- Super Pages (2MB)

ORACLE

Xen Project: hypervisor: ARM

• IOMMU translation for remote processors

- Graphic rendering (GPUs)
- High Quality Video Decoding (IPUs)
- Passthrough:
 - Device assignment (non-PCI)
 - MMIO addresses assignment
- Interrupts:
 - Interrupt latency reduction
 - No more maintenance interrupts
 - Interrupt migration
 - Physical follow virtual (or vice-versa)

ORACLE[®]

Xen Project: hypervisor: ARM

- ARM Intellectual Property (different vendors customize what they will have)
 - GICv2 (done)
 - -GICv2m (Xen 4.6)
 - GICv3 (in review)
 - -SMMU (in review)
 - SMMUv2 (Xen 4.6, compatible with v1)
 - PSCI (in review)
 - UEFI (in review)



Xen Project: hypervisor: ARM

That means following motherboards support Xen 4.5:

- Midway (Calxeda)
- Vexpress (ARM Ltd.)
- OMAP5, OMAP4, DRA7 (Texas Instrument)
- Exynos5 (Samsung chip on the Arndale and various smartphones and tablets)
- SunXI (Allwinner)
- Mustang (Applied Micro X-Gene, the armv8 SoC)



Xen Project: hypervisor: x86

- Multibootv2+EFI
- HPET interrupt fixes
- IOMMU ABI for guests to map their DMA regions
- VMWare backdoor calls
- VPMU 'perf' support in Xen
 - Requires Linux 3.17 or later.
- vNUMA

Xen Project: hypervisor: x86

- Introspection of guests:
 - -extending mem_access for PV guests
 - -Introspection of HVM guests and injecting instructions
- 1TB slow destruction of guests (aka PGC_need_scrub)
- NUMA memory scrubbing
- PVH dom0
 - Requires Linux 3.16 or higher
- Re-write of vHPET

Xen Project: hypervisor: x86

- alternative assembler
 - Patching of hypervisor based on CPU support
- Serial support:
 - Broadcom TruManage chip (Serial Over Lan on some AMD chipsets)
 - NetMoss chipsets aka Oxford chipset (PCIe serial cards)
- ioreq-server, aka secondary emulators
 - Multiple QEMU for a guest with each different role
- Real-time scheduler

Xen Project: hypervisor: Intel

- SandyBridge and later:
 - -vAPIC in PVHVM guests (less VMEXIT)
- Ivy Bridge new features:
 - Cache QoS Monitoring
- Broadwell
 - SMAP (kernel can't touch user-mode pages)

Xen Project: hypervisor: AMD

- Family 10h and later:
 - PVH AMD hardware support
 - Microcode fixes
- Kabini, Kaveri and further:
 - Data breakpoint Extensions
 - Masking MSR support

Xen Project: Toolstack

- VM Generation ID:
 - Windows 2012 Server and later domain controllers.
- Migration-v2:
 - Faster, easier to maintain, with a design document, fixes bugs, and supports older data stream (upgrades to new version when streaming)
- Remus:
 - libxl support
 - libxc migration-v

Xen Project: Toolstack (cont.)

• libxl:

- discard support (also in libvirt) to disable/enable support
- snapshotting disks
- JSON to keep track of guest configs
- Systemd support



Libvirt, libxen

- PCI/SR-IOV passthrough, including hot{un}plug
- Migration
- Improved concurrency through job support in the libxl driver no more locking entire driver when modifying a domain
- Improved domxml-{to,from}-native support, e.g. for converting between xl config and libvirt domXML and viseversa
- PV console support
- Improved qdisk support



Libvirt – (cont.)

- Support for:
 - -<interface type='network'> allows using libvirt-managed networks in the libxl driver.
 - PARAVIRT and ACPI shutdown flags.
 - PARAVIRT reboot flag.
 - Domain lifecycle event configuration, e.g. on_crash, on_reboot, etc.
 - Discard (enabled/disable)
- A few improvements for ARM
- lots of bug fixes



QEMU (upstream)

- Bigger PCI hole in QEMU
 - To have 3GB of MMIO space
- Intel IGD passthrough
 - Changes in Linux i915 and Windows driver needed
- Xen block PV driver in OVMF (UEFI in guest)
- QEMU for ARM
 - Enable PV backends such as qdisk, framebuffer, etc



Linux (3.15, 3.16, 3.17, 3.18)

- Xen-EFI hypercall support
 - Dom0 can boot under Xen-EFI
- Netback
 - Multiqueue
 - Grant copy the header instead of map and memcpy
- Block backend
 - multiqueue (NVMe)
- Remove _PAGE_IOMAP in Linux kernel
- PVH dom0 support



Linux (cont.)

- Bus/slot reset for PCI passthrough
- VPMU (perf)
 - To allow dom0/Xen/guest performance monitoring
- vAPIC
 - Interrupts for passthrough via vAPIC, not events



FreeBSD, MiniOS

• FreeBSD:

- PVH support already in 10.0 (January this year)
- dom0 for 11 (next year)
- multiboot support for bootloader

• MiniOS

 $-\operatorname{Make}$ it build and work under ARM



GRUB2 Multiboot v2 + EFI

- Xen.efi is an EFI application with EFI calling convention
 - GRUB2 can call it (chainloading), but no support for parameters all in a config file
 - Want GRUB2 menu with capability to edit the parameters (like legacy)
- Solution:
 - New multiboot2+EFI inside Xen.efi which looks like EFI and has multibootv2 structure
 - Multibootv2 would have new flags to tell GRUB2 to:
 - Load it in 64-bit mode.
 - Disable ExitBootServices() call.



Q&A

