

# What's Going on in Xen

Status update of the Xen Project

Roger Pau Monné

Barcelona – October 21, 2015




# Goals of this presentation




- ▶ Description of Xen.
- ▶ Understanding how the Xen community works.
- ▶ A peek into Xen's new features.

# Xen Architecture (type-1 hypervisor)



# Type-2 hypervisor architecture



# Paravirtualization



- ▶ Virtualization technique developed in the late 90s.
- ▶ Designed by:
  - ▶ XenoServer research project at Cambridge University.
  - ▶ Intel.
  - ▶ Microsoft labs.
- ▶ x86 instructions behave differently in kernel or user mode, options for virtualization were full software emulation or binary translation.
  - ▶ Design a new interface for virtualization.
  - ▶ Allow guests to collaborate in virtualization.
  - ▶ Provide new interfaces for virtualized guests that allow to reduce the overhead of virtualization.
- ▶ The result of this work is what we know today as paravirtualization.

# Paravirtualization



- ▶ All this changes lead to the following interfaces being paravirtualized:
  - ▶ Disk and network interfaces
  - ▶ Interrupts and timers
  - ▶ Boot directly in the mode the kernel wishes to run (32 or 64bits)
  - ▶ Page tables
  - ▶ Privileged instructions

# Full virtualization



- ▶ With the introduction of hardware virtualization extensions Xen is able to run unmodified guests
- ▶ This requires emulated devices, which are handled by Qemu
- ▶ Makes use of nested page tables when available.
- ▶ Allows to use PV interfaces if guest has support for them.

# New x86 virtualization mode: PVH



- ▶ PV in an HVM container.
- ▶ PVH should use the best aspects from both PV and HVM:
  - ▶ No need for any emulation.
  - ▶ Has a "native" MMU from guest point of view.
  - ▶ Has access to the same protection levels as bare metal.
- ▶ Written by Mukesh Rathor @ Oracle.
- ▶ Significant revisions by George Dunlap @ Citrix.
- ▶ Currently the internals of PVH in Xen are being rewritten.



# The virtualization spectrum



VS	Software virtualization
VH	Hardware virtualization
PV	Paravirtualized

	Poor performance
	Room for improvement
	Optimal performance

Disk and network  
 Interrupts and timers  
 Emulated motherboard  
 Privileged instructions  
 and page tables

HVM	VS	VS	VS	VH
HVM with PV drivers	PV	VS	VS	VH
PVHVM	PV	PV	VS	VH
PVH	PV	PV	PV	VH
PV	PV	PV	PV	PV

# Xen community overview



- ▶ The Xen Hypervisor was released under the GPL2 on 2003.
- ▶ The Xen Project became a Linux Foundation Collaborative Project in 2013.
- ▶ Xen governance similar to the Linux kernel.
- ▶ Xen Project teams:
  - ▶ Xen Hypervisor.
  - ▶ ARM Hypervisor.
  - ▶ XAPI.
  - ▶ Mirage OS.
  - ▶ Linux PVOPS.

# Xen governance



- ▶ Roles:
  - ▶ Maintainers: own one or more components in the Xen source tree.
  - ▶ Committers: maintainers that are allowed to commit changes into the source code repository.
  - ▶ Sub-projects and teams: run by individuals, projects are related or based on the Xen Project.
- ▶ See <http://www.xenproject.org/developers/governance.html> for more information.

# Xen Hypervisor



- ▶ Main project, contains the hypervisor and the toolstack.
- ▶ Led by 5 committers; 2 from Citrix, 1 from Suse, 1 from Oracle, 2 Independent.
- ▶ During the 4.6 release cycle the Xen Project had contributions from 96 individuals from 30 organizations.
- ▶ Organizations that contributed to the 4.6 release: Citrix, SUSE, Intel, Oracle, Linaro, Fujitsu...
- ▶ Full list can be found at [http://wiki.xen.org/wiki/Xen\\_Project\\_4.6\\_Acknowledgements](http://wiki.xen.org/wiki/Xen_Project_4.6_Acknowledgements).

# Xen's new features



- ▶ Xen 4.6 released one week ago:
  - ▶ Improved and extended the memory event subsystem.
  - ▶ Improved the default XSM policy and added support for testing it in OSStest.
  - ▶ vTPM 2.0 support.
  - ▶ Improved grant table scalability.
  - ▶ New locking mechanism to improve fairness (ticket locks).

# x86 specific improvements



- ▶ Support for several new features in upcoming Intel processors: alternate P2M, page modification logging, cache allocation and memory bandwidth monitoring.
- ▶ Support for profiling Xen using a performance monitoring unit (PMU).
- ▶ Virtual NUMA topology for Xen guests.

# ARM specific improvements



- ▶ Supported number of vCPUs increased from 8 to 128.
- ▶ Passthrough of non-PCI devices.
- ▶ GICv2 on GICv3 support.
- ▶ OVMF support (UEFI for guests).
- ▶ Support for guests running with 64KB pages.
- ▶ Support for new platforms: Renesas R-Car Gen2, Thunder X, Huawei hip04-d04 and Xilinx ZynqMP SoC.

# Toolstack updates



- ▶ New live migration implementation (Migration v2).
- ▶ Remus (HA on top of Xen) reworked on top of Migration v2.
- ▶ Improved Spice/QXL support.
- ▶ Xenalyze merged into Xen codebase.



# Test lab



- ▶ Moved the Xen test infrastructure to a new lab.
- ▶ Funded by the Xen Project advisory board, currently with 24 dedicated hosts.
- ▶ Expanded test coverage: XSM, libvirt...

# OS related updates



- ▶ Linux:
  - ▶ Added multipage support to blkfront.
  - ▶ Added xen-scsi support.
  - ▶ Performance improvements.
  - ▶ PV guests can support up to 512GB of RAM.
- ▶ FreeBSD:
  - ▶ PVH DomU/Dom0 support.
  - ▶ Removal of the i386 PV port.
  - ▶ Support for indirect descriptors in blkfront.

# The Xen Project status



- ▶ The 4.6 release has been the one with the best quality and quantity of commits so far.
- ▶ The Xen Project ecosystem is a very active community with contributions from a wide array of companies and individuals.
- ▶ First release to be completely on time.

## Q&amp;A



Thanks  
Questions?