

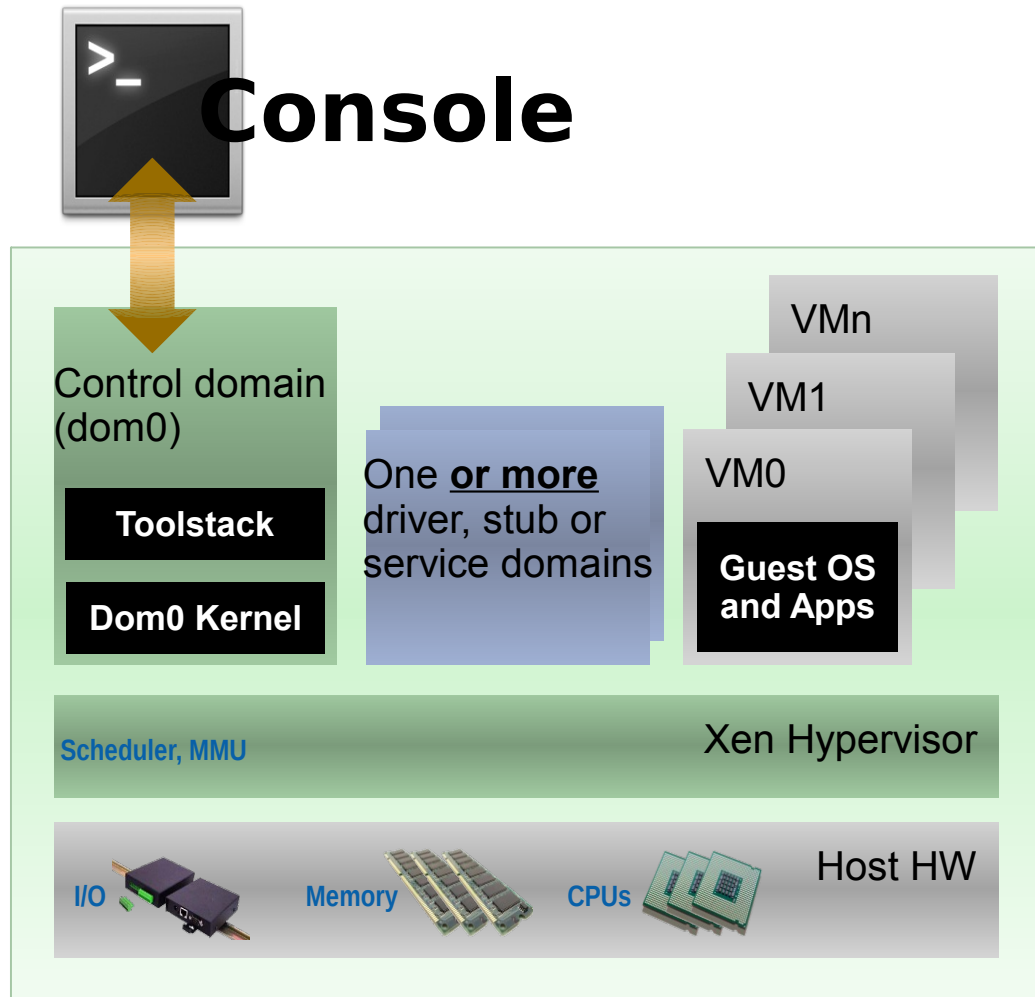
xen.org

Evolving Xen Paravirtualisation



Ian Campbell
ian.campbell@citrix.com

Basic Xen Concepts



Control Domain aka **Dom0**

- Dom0 kernel with drivers
- Xen Management Toolstack
- Trusted Computing Base

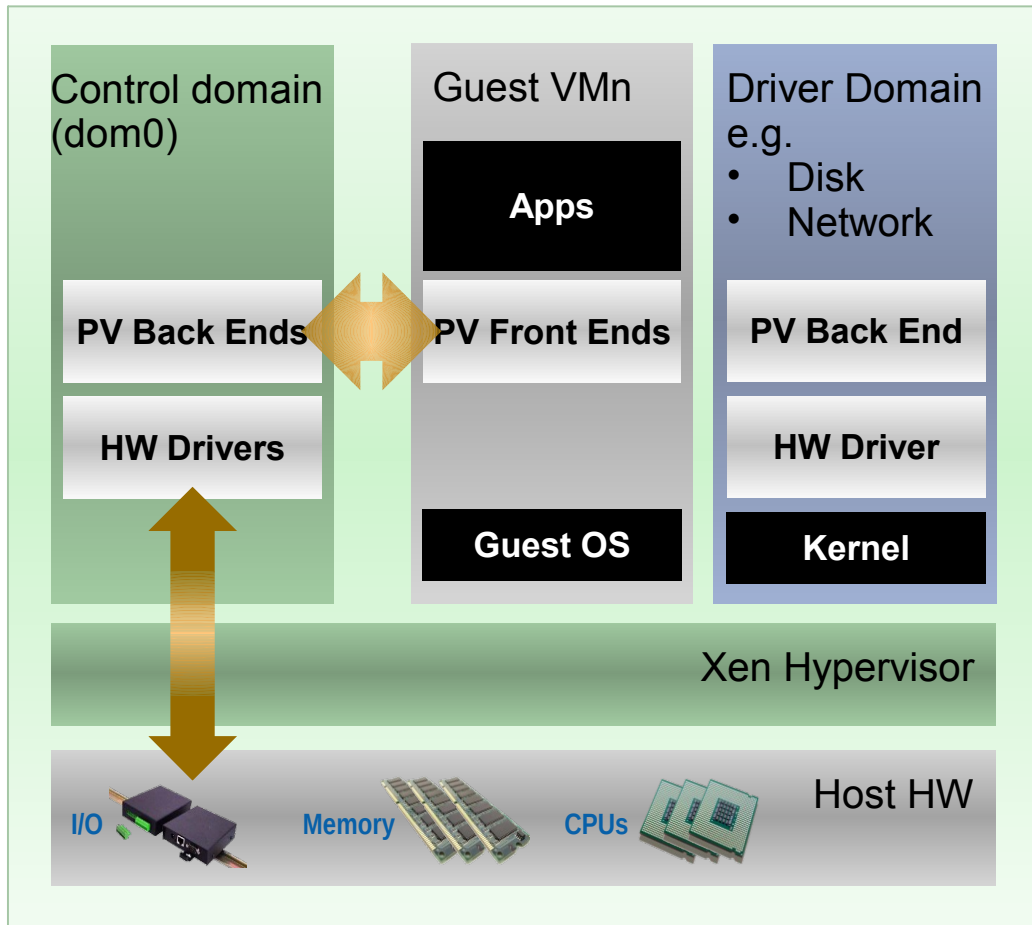
Guest Domains

- Your applications
- Your user's applications

Driver/Stub/Service Domain(s)

- A “driver, device model or control service in a box”
- De-privileged and isolated
- Lifetime: start, stop, kill

PV Domains



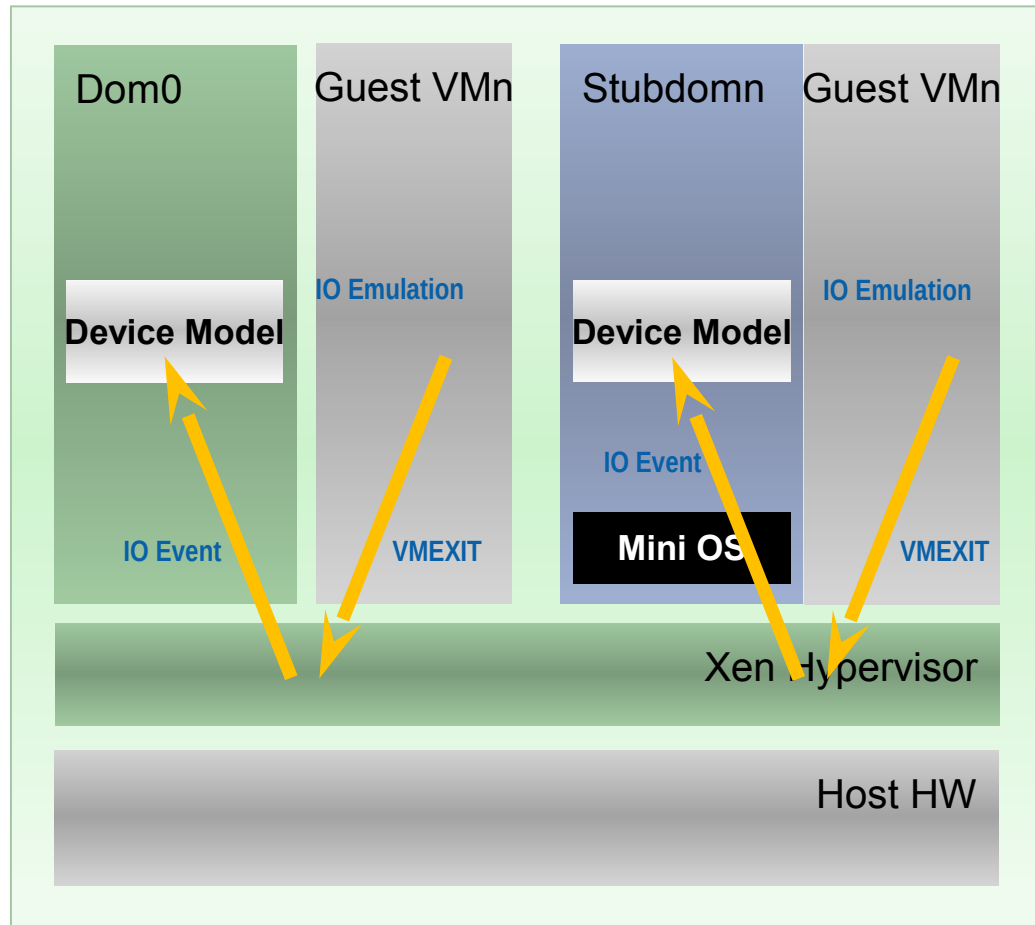
Limitations

- limited set of virtual hardware

Advantages

- Fast
- Works on any system
- (even without virt extensions)

HVM Domains



Disadvantages

- Slower than PV due to Emulation (mainly I/O devices)

Advantages

- Install the same way as native Operating System

Stub Domains


- Security
- Isolation
- Reliability and Robustness

PV Spectrum

Virtualisation Spectrum

	PV		PVHVM	PV I/O	HVM
Privileged Instructions	PV		HW	HW	HW
Platform. Motherboard "PC"	PV		EMU	EMU	EMU
Interrupts / Timers	PV		PV	EMU	EMU
I/O (Disk/Network)	PV		PV	PV	EMU

PV Evolution: PVH



- PV with Hardware
- Developed by Mukesh @ Oracle
- Targeting 4.3 Release


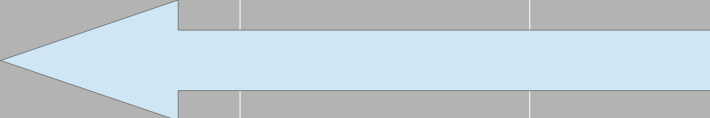
PVH: Why?

- System call performance on 64-bit x86
- PV MMU changes are large
 - Harder to port new guests OSes to Xen
 - Not always popular with maintainers
- Hardware has gotten better
 - New capabilities
 - Better capabilities

PVH: Spectrum



Virtualisation Spectrum

	PV	PVH	PVHVM	PV I/O	HVM
					
Privileged Instructions	PV	HW	HW	HW	HW
Platform. Motherboard "PC"	PV	PV	EMU	EMU	EMU
Interrupts / Timers	PV	PV	PV	EMU	EMU
I/O (Disk/Network)	PV	PV	PV	PV	EMU

Xen on ARM

- PVH like model from day 1
 - Uses new ARMv7 + v8 HYP mode
 - Hardware paging, interrupts
- Kernel side upstream from 3.7
 - ~1,000 LoC
 - took ~1 release cycle
- ARMv7 in 4.3, ARMv8 Preview



- **IRC:**
 - ##xen @ Freenode
- **Lists:**
 - xen-{users,devel,api}@lists.xen.org
- **Wiki:**
 - <http://wiki.xen.org/>

Questions ...



Slides available under CC-BY-SA 3.0



PVH: Still PV

- OS is still fundamentally aware
 - But needs to do less about it
- No emulation
 - ... or device models
- Mix & match hardware and software
 - Find the sweet spot



Conclusions

- PV Xen is continuing to evolve
- Take advantage of the best mix of software and hardware features
- Retain the advantages of a PV system