

Virtualization in the Cloud: Featuring Xen and XCP



FREENODE: lars_kurth

Lars Kurth
Xen Community Manager
lars.kurth@xen.org



@lars_kurth

A Brief History of Xen in the Cloud

Late 90s



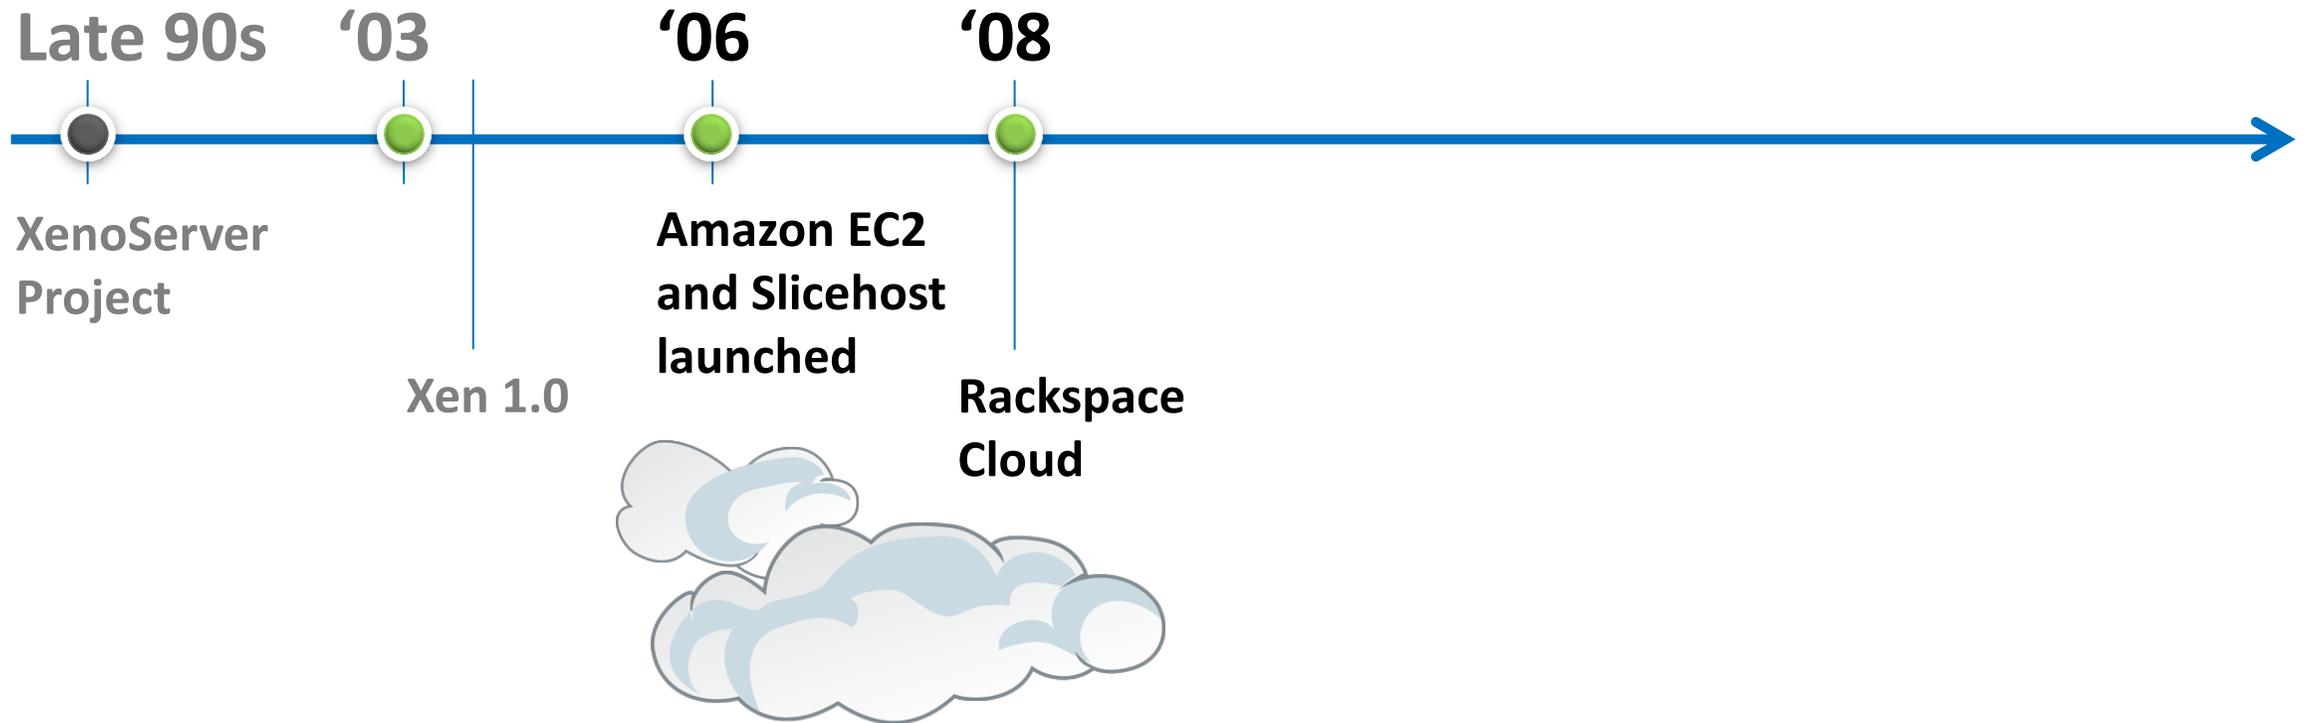
**XenoServer
Project**



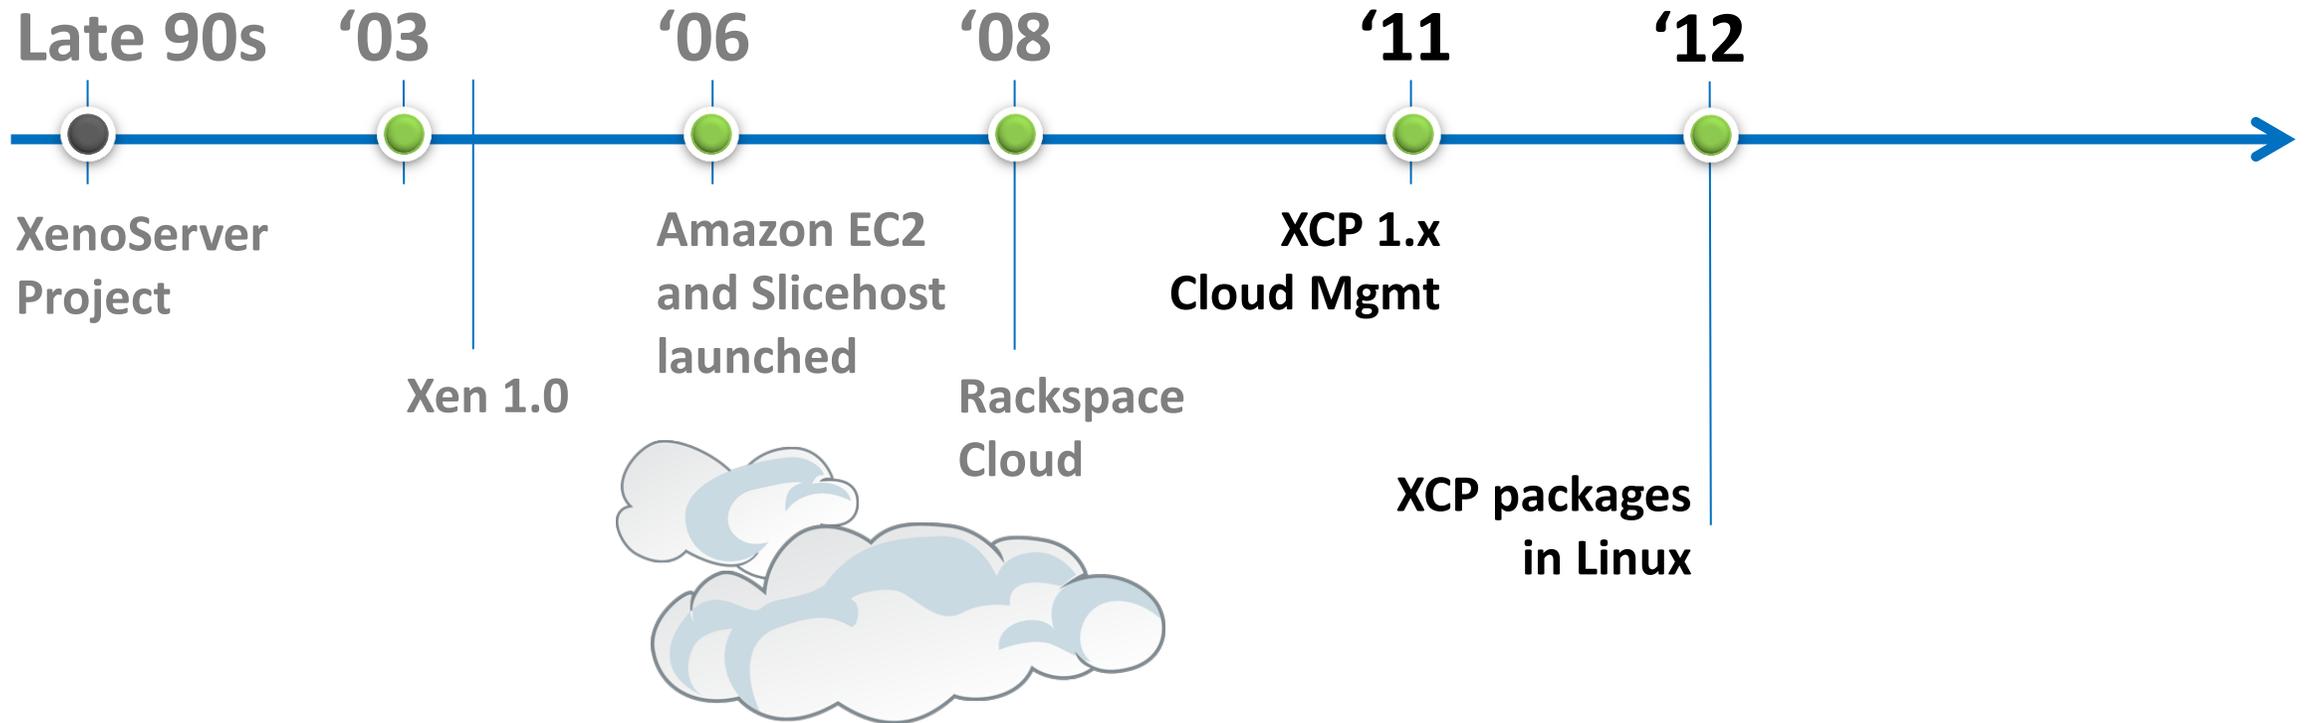
A Brief History of Xen in the Cloud



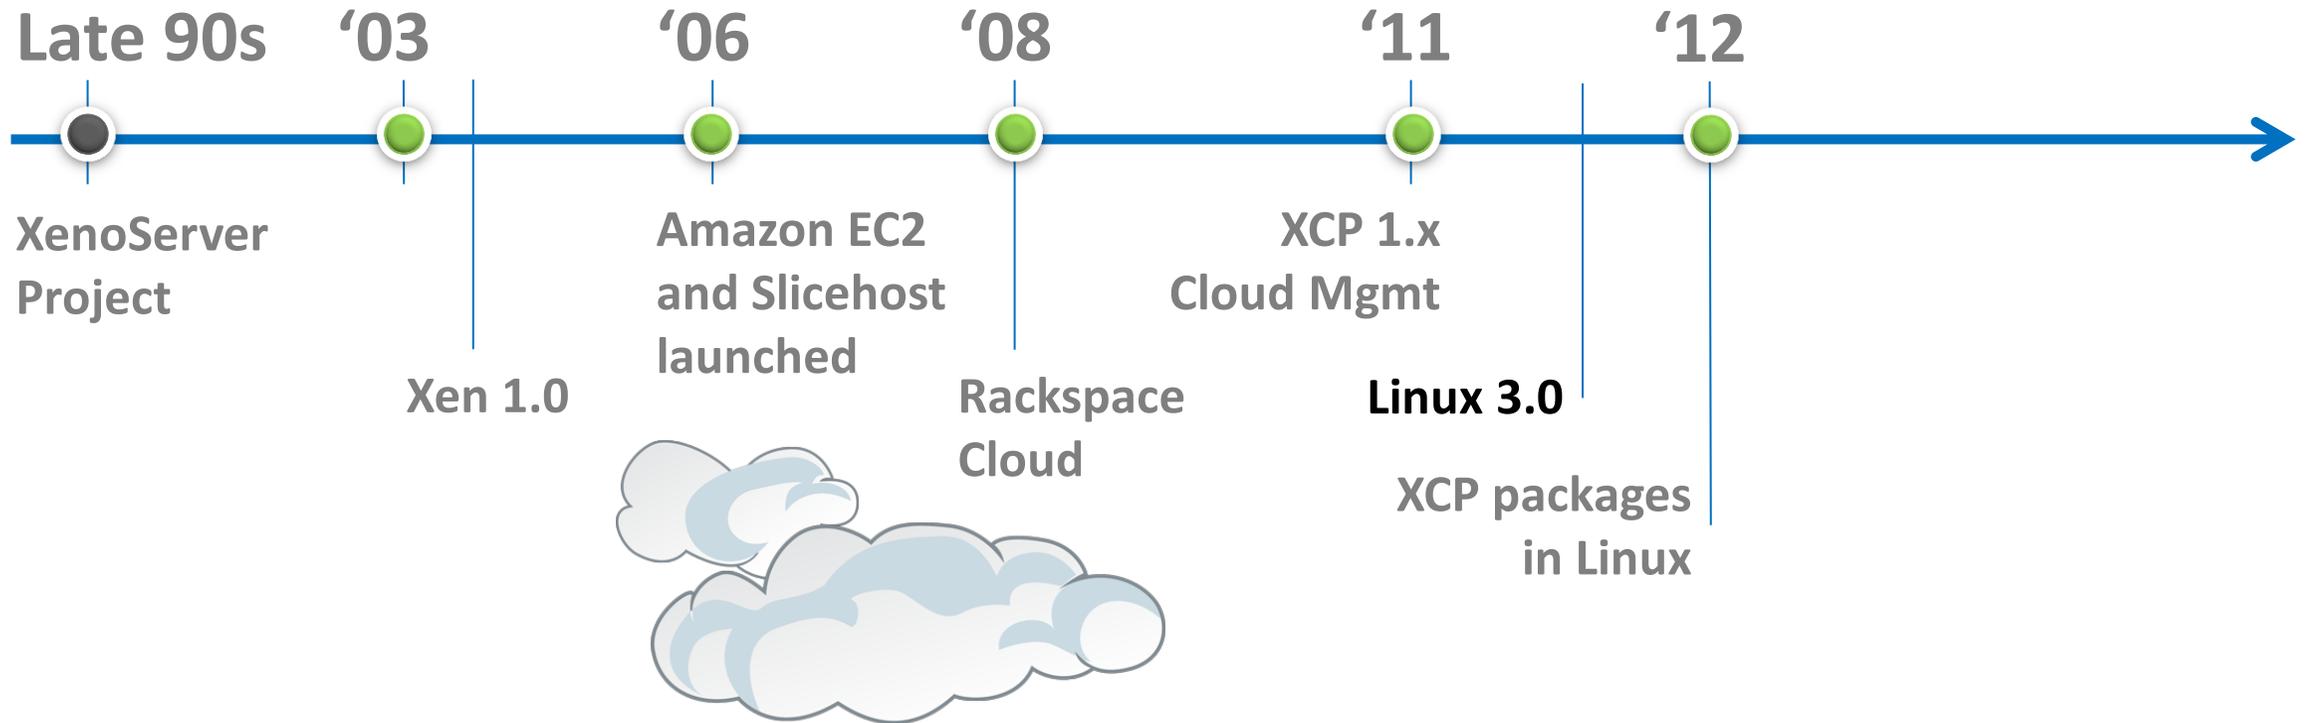
A Brief History of Xen in the Cloud



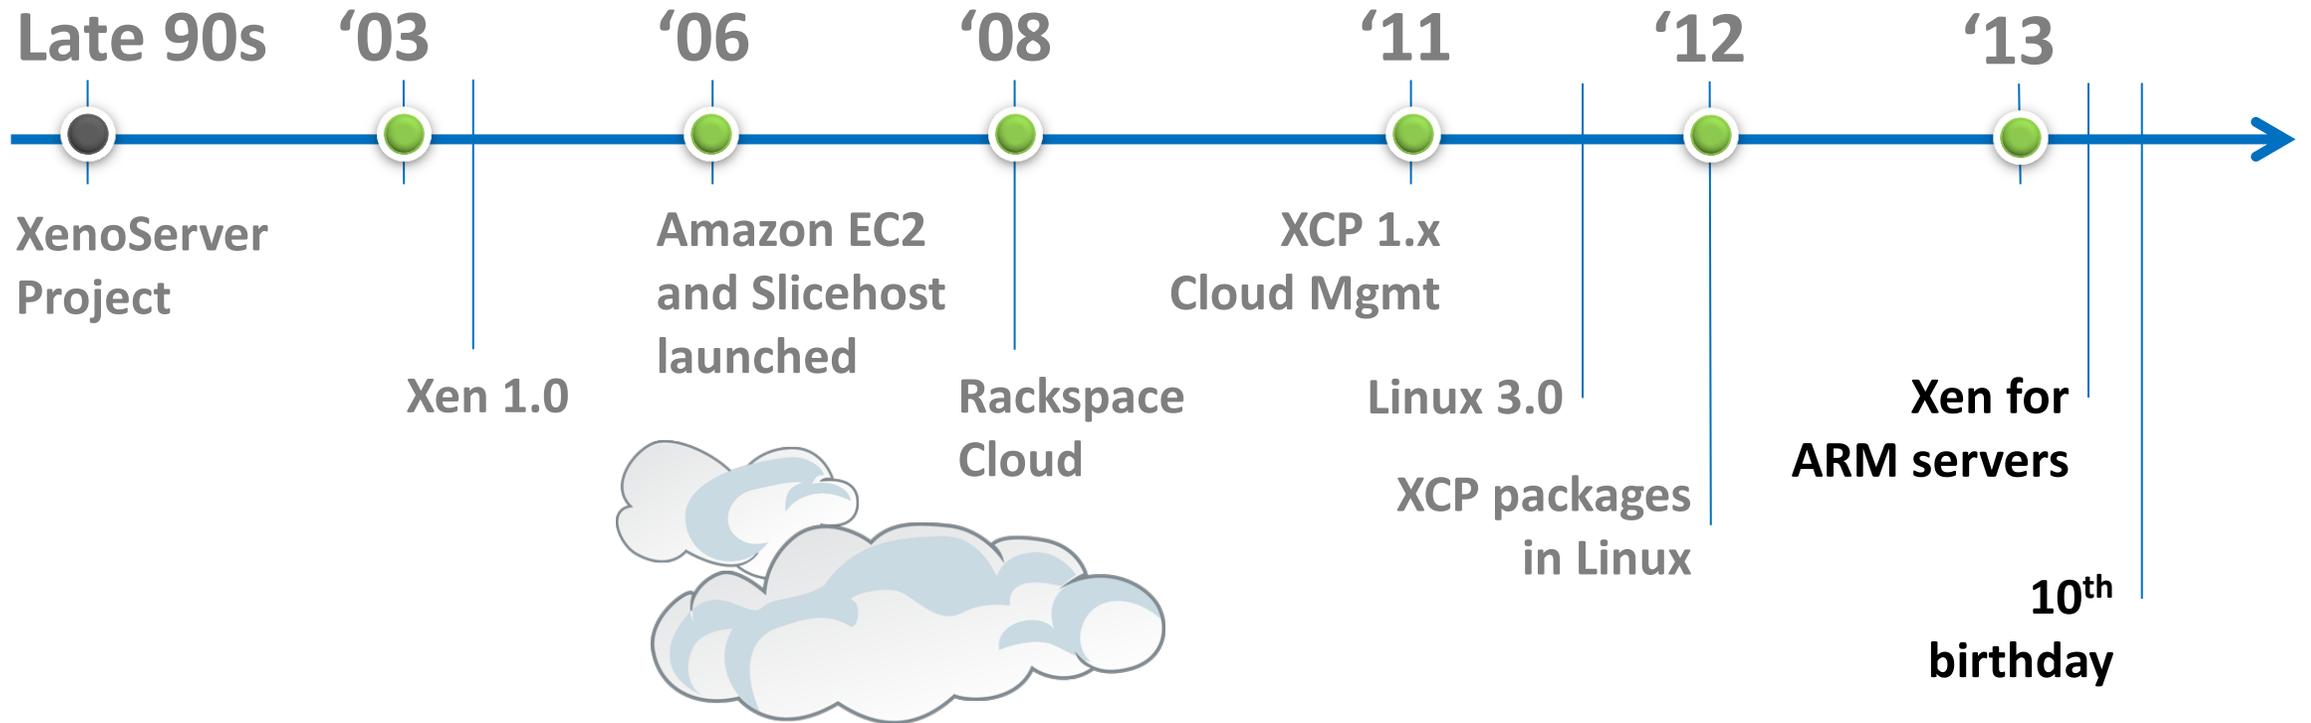
A Brief History of Xen in the Cloud



A Brief History of Xen in the Cloud



A Brief History of Xen in the Cloud



**The Xen Hypervisor was designed for
the Cloud straight from the outset!**

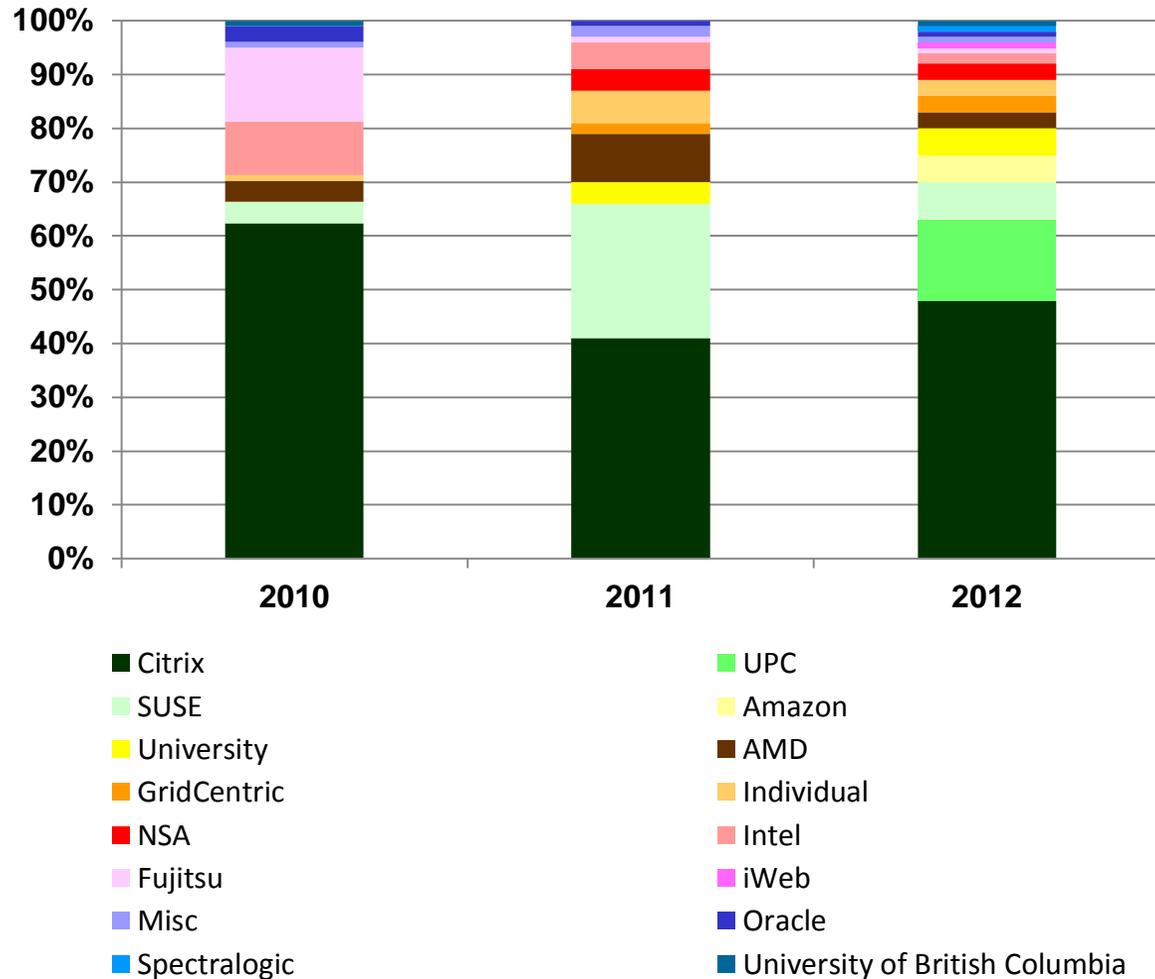


Xen.org

- Guardian of Xen Hypervisor and related OSS Projects
- Xen Governance similar to Linux Kernel
 - Plus project lifecycle and Project Management Committee (PMC)
- Projects
 - Xen Hypervisor
(led by 5 committers, 2 from Citrix, 1 from Suse, 2 Independent)
 - Xen Cloud Platform aka XCP (led by Citrix)
 - Xen ARM : Xen for mobile devices (led by Samsung)



Xen contributor community is diversifying



- The number of “significant” active vendors is increasing
- New feature development driving new participation



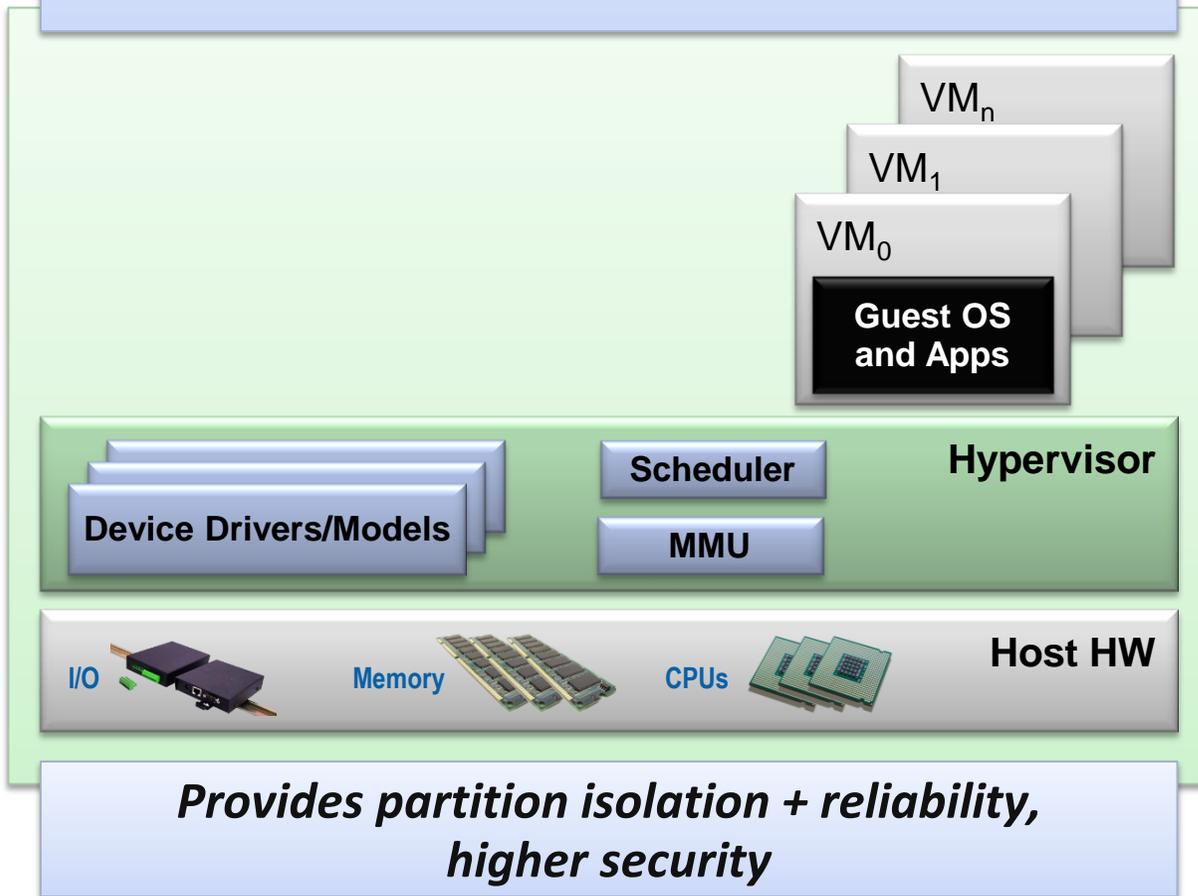
Xen Overview



Hypervisor Architectures

Type 1: Bare metal Hypervisor

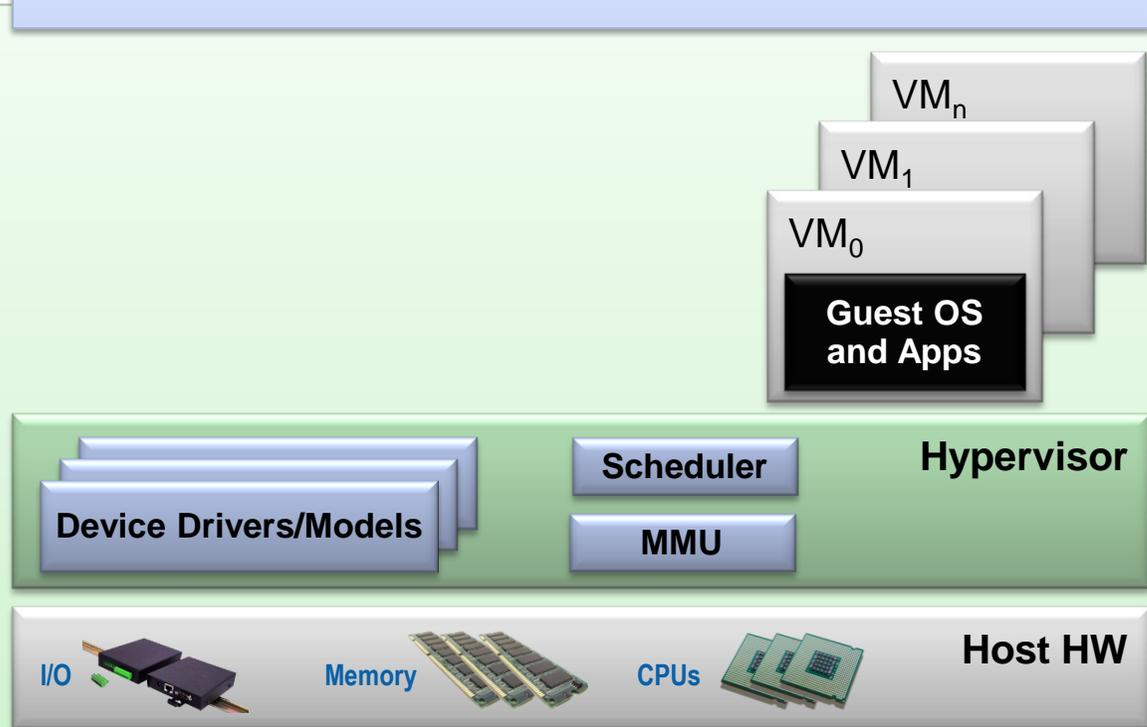
A pure Hypervisor that runs directly on the hardware and hosts Guest OS's.



Hypervisor Architectures

Type 1: Bare metal Hypervisor

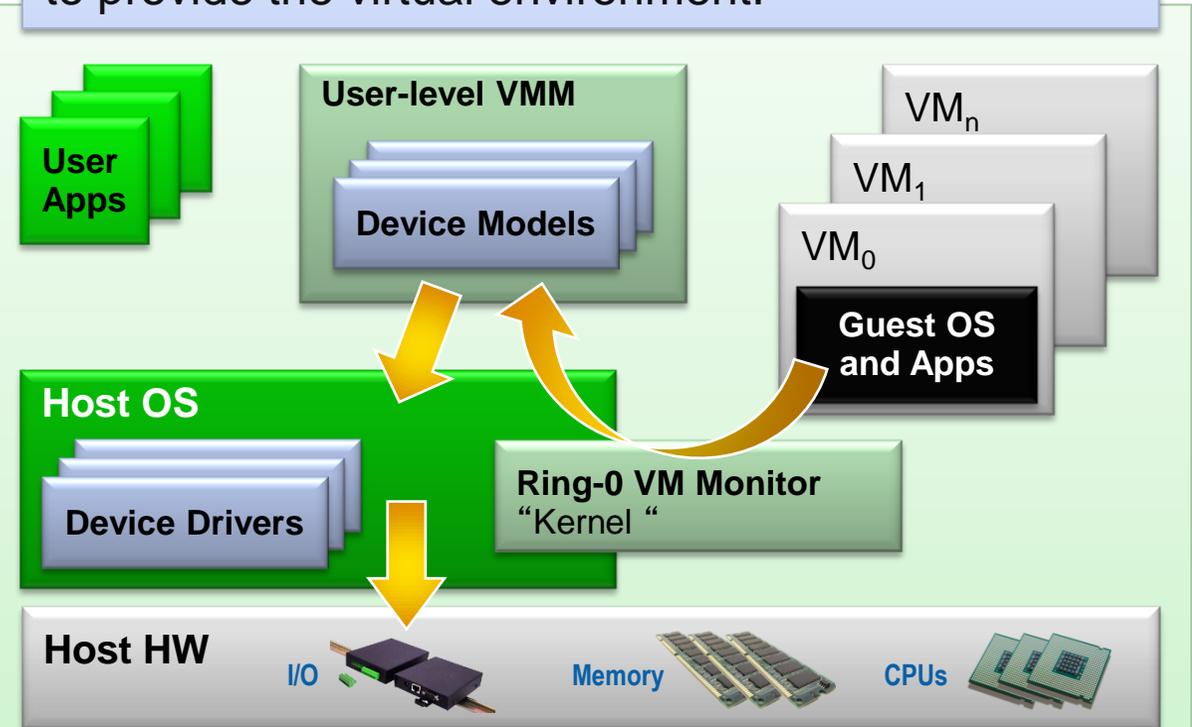
A pure Hypervisor that runs directly on the hardware and hosts Guest OS's.



*Provides partition isolation + reliability,
higher security*

Type 2: OS 'Hosted'

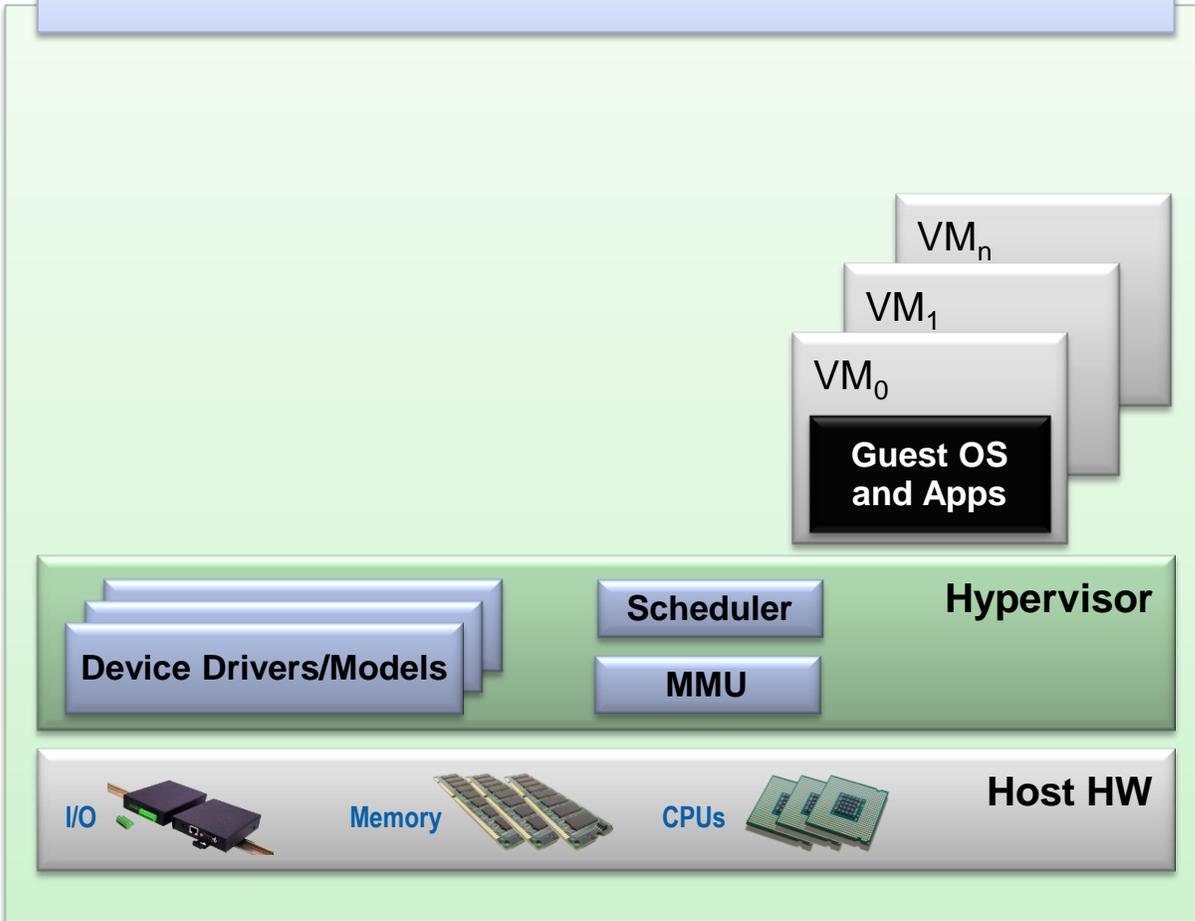
A Hypervisor that runs within a Host OS and hosts Guest OS's inside of it, using the host OS services to provide the virtual environment.



*Low cost, no additional drivers
Ease of use & installation*

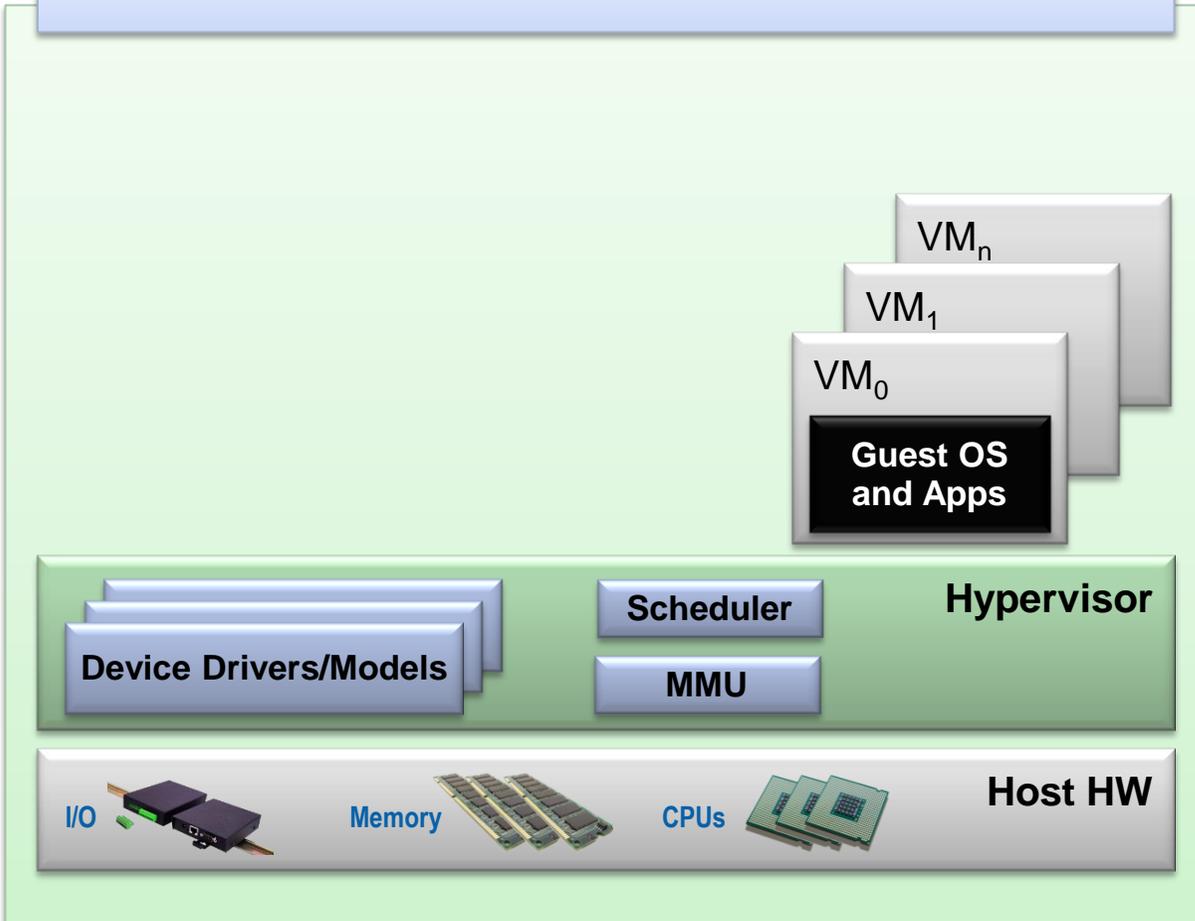
Xen: Type 1 with a Twist

Type 1: Bare metal Hypervisor

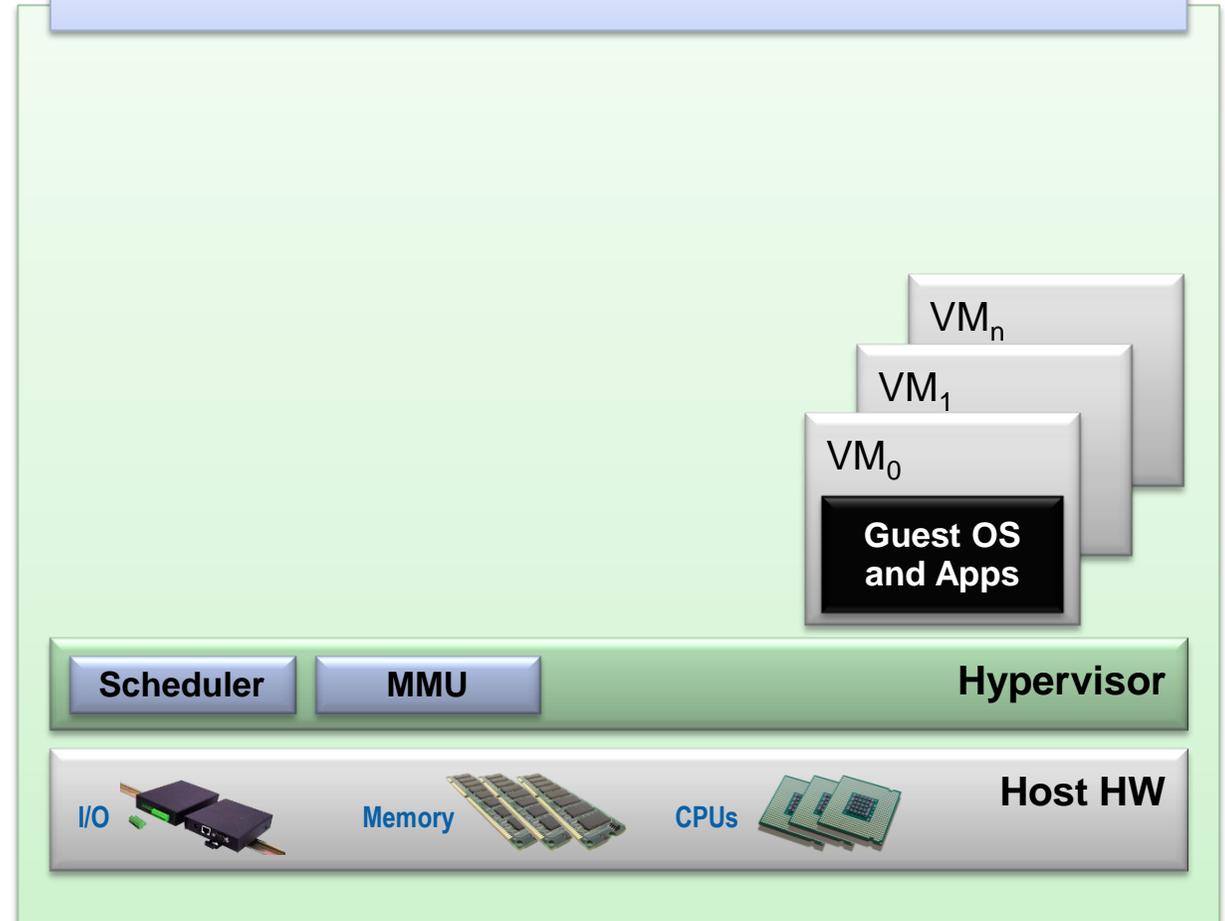


Xen: Type 1 with a Twist

Type 1: Bare metal Hypervisor

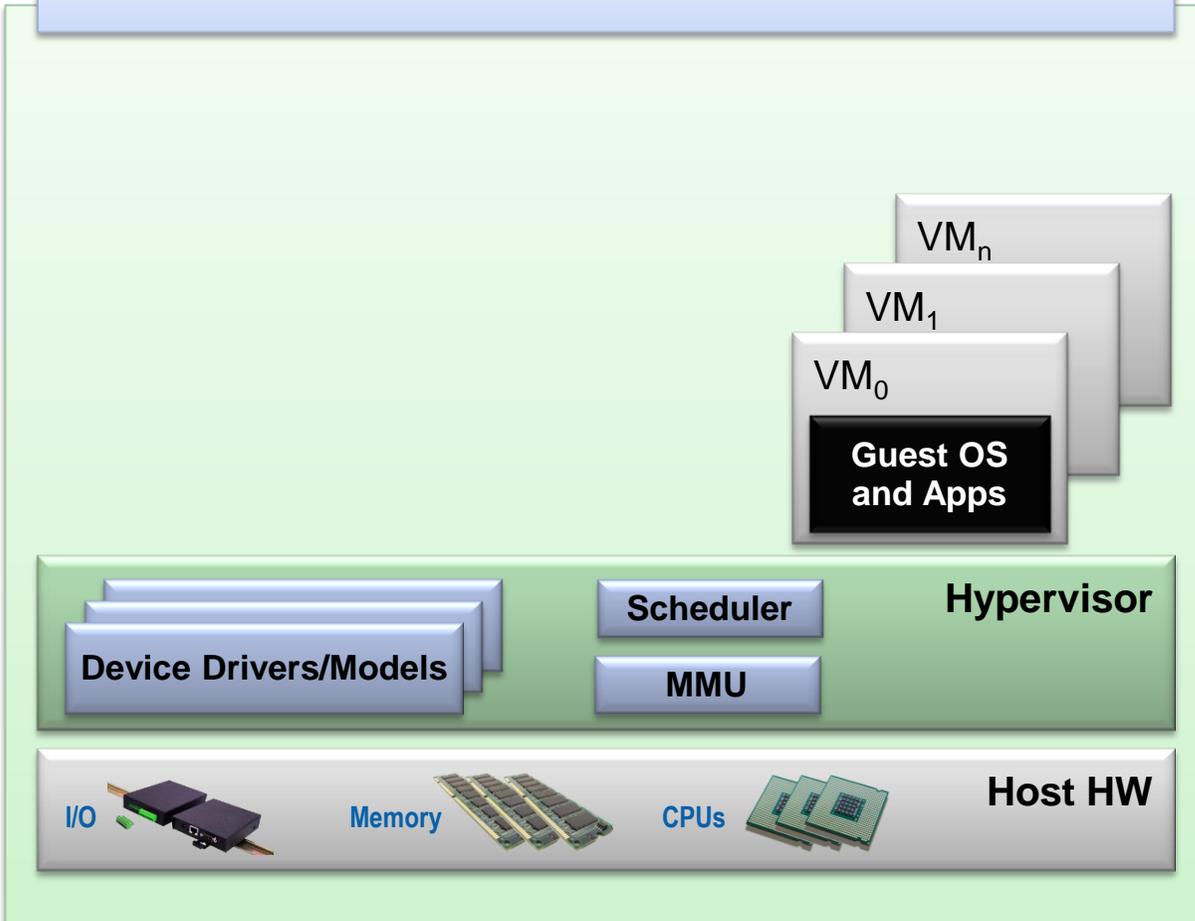


Xen Architecture

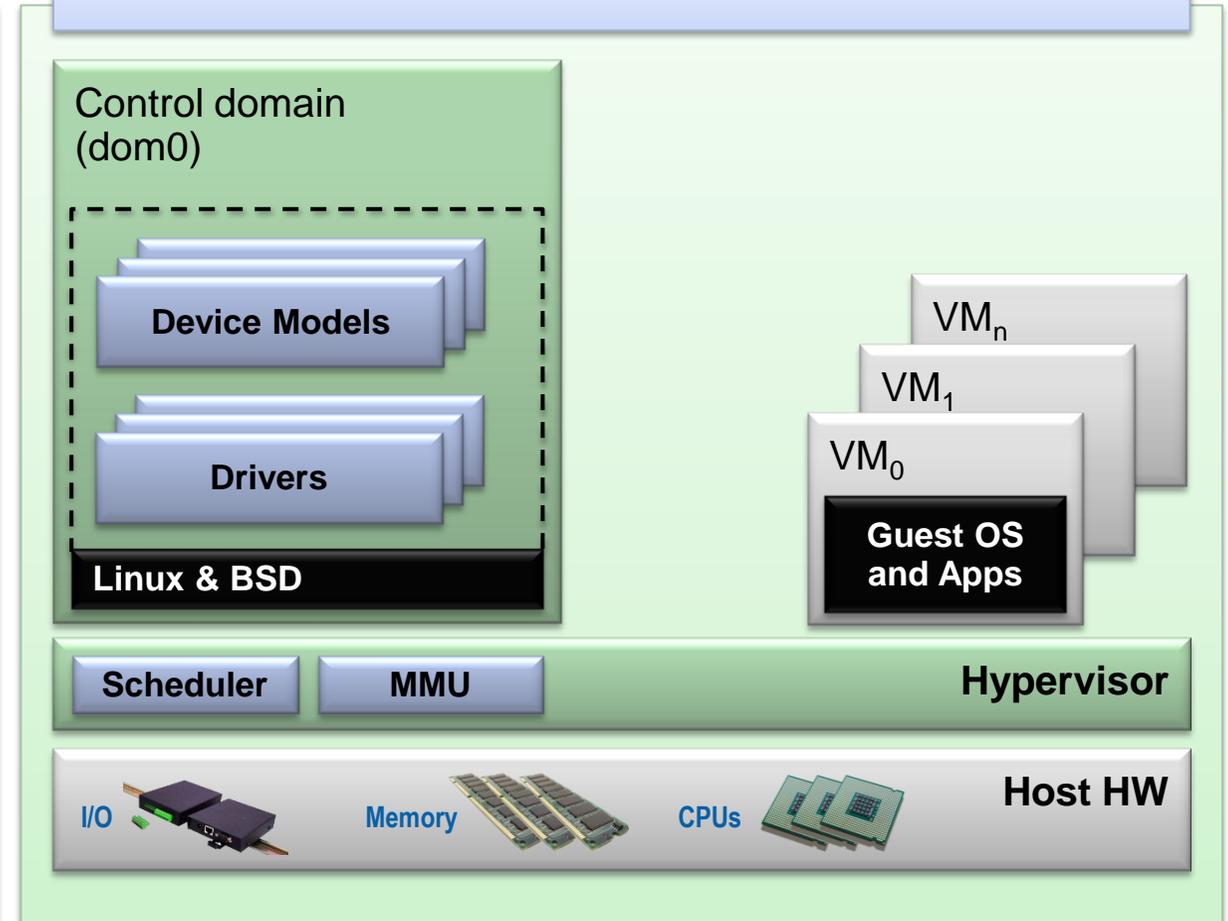


Xen: Type 1 with a Twist

Type 1: Bare metal Hypervisor



Xen Architecture



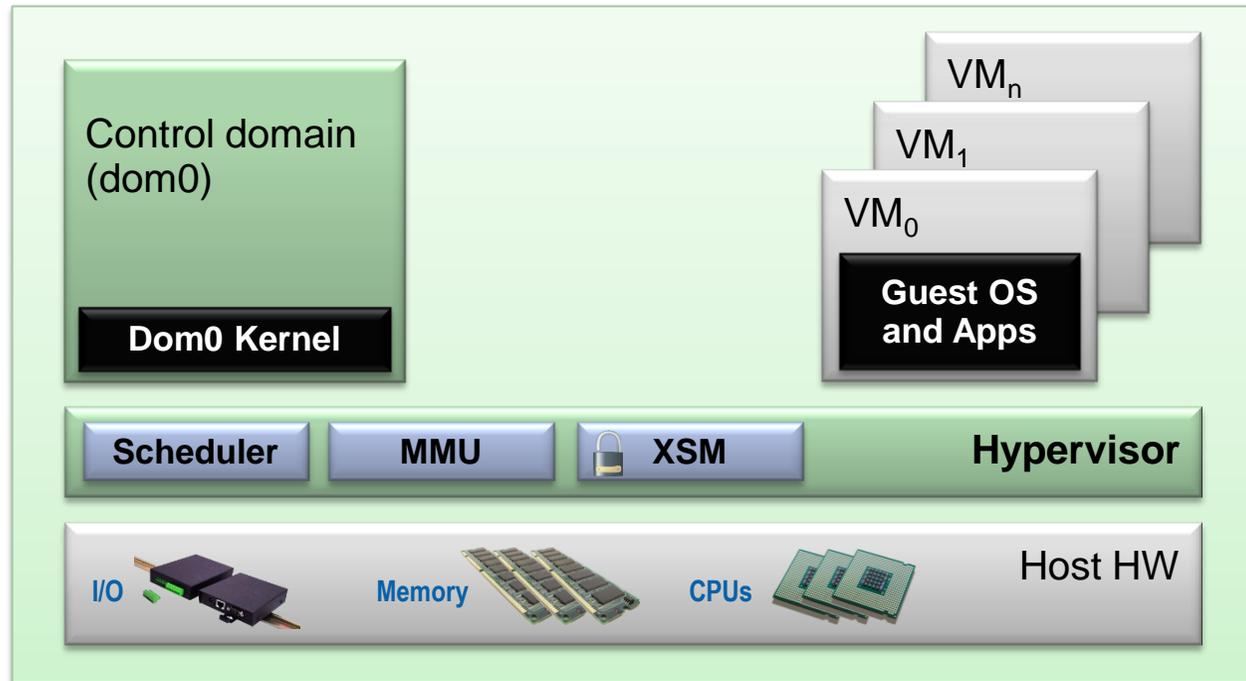
Xen and Linux

- Xen Hypervisor is **not** in the Linux kernel
- **BUT**: everything Xen and Xen Guests need to run is!
- Xen packages are in all Linux distros (except RHEL6)
 - Install Dom0 Linux distro
 - Install Xen package(s) or meta package
 - Reboot
 - Config stuff: set up disks, peripherals, etc.



[More info: wiki.xen.org/wiki/Category:Host_Install](http://wiki.xen.org/wiki/Category:Host_Install)

Basic Xen Concepts



■ Trusted Computing Base

Control Domain aka **Dom0**

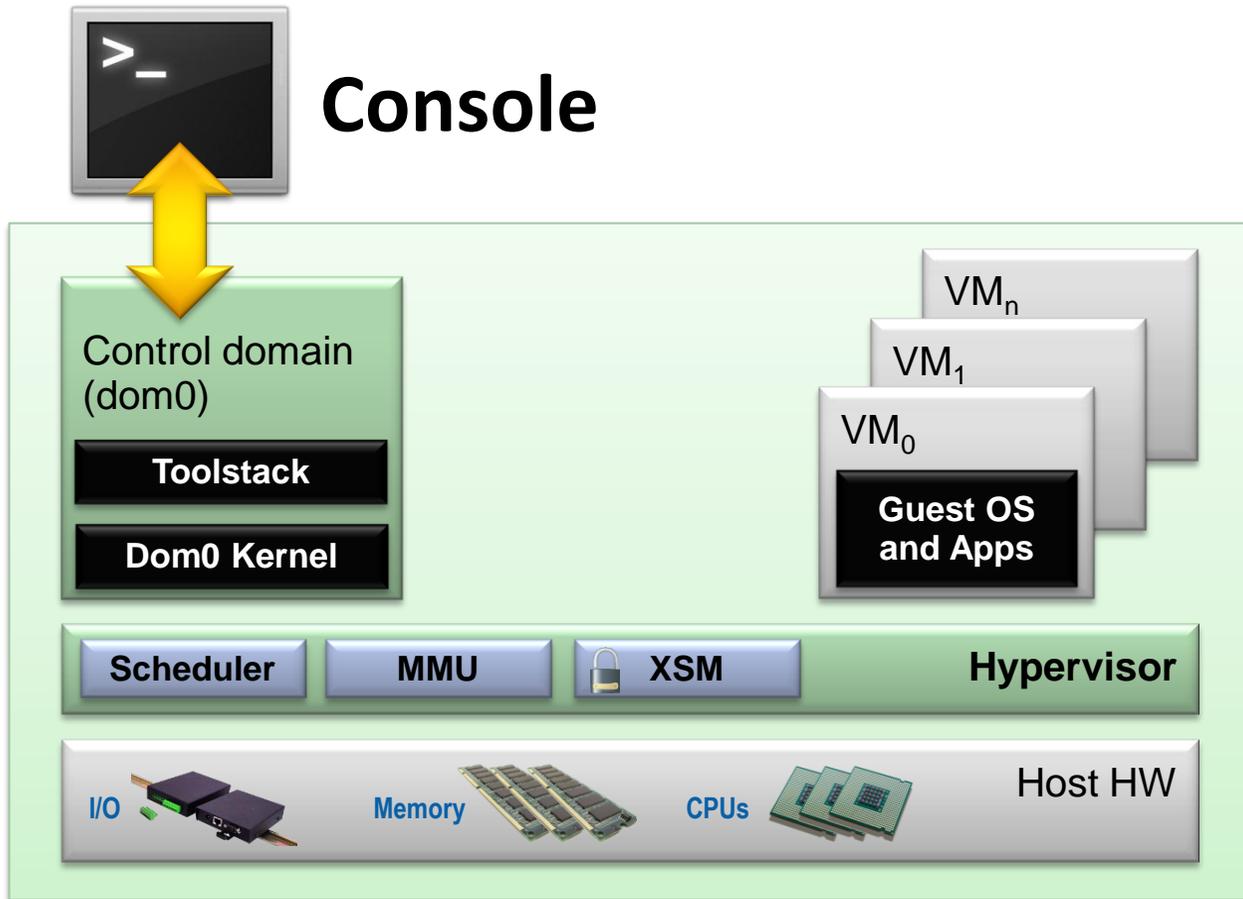
- Dom0 kernel with drivers

Guest Domains

- Your apps



Basic Xen Concepts



■ Trusted Computing Base

Console

- Interface to the outside world

Control Domain aka **Dom0**

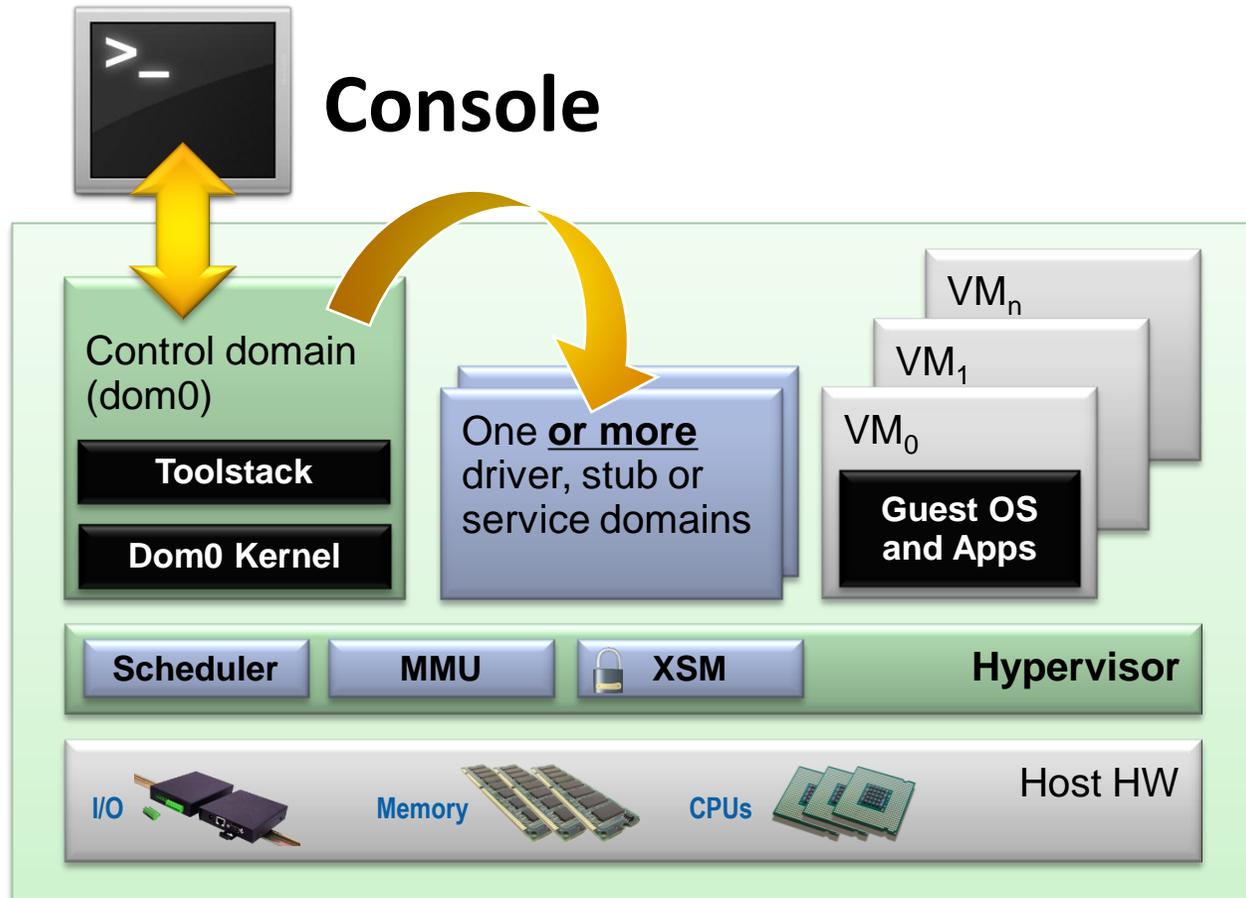
- Dom0 kernel with drivers
- Xen Management Toolstack

Guest Domains

- Your apps



Basic Xen Concepts



■ Trusted Computing Base

Console

- Interface to the outside world

Control Domain aka **Dom0**

- Dom0 kernel with drivers
- Xen Management Toolstack

Guest Domains

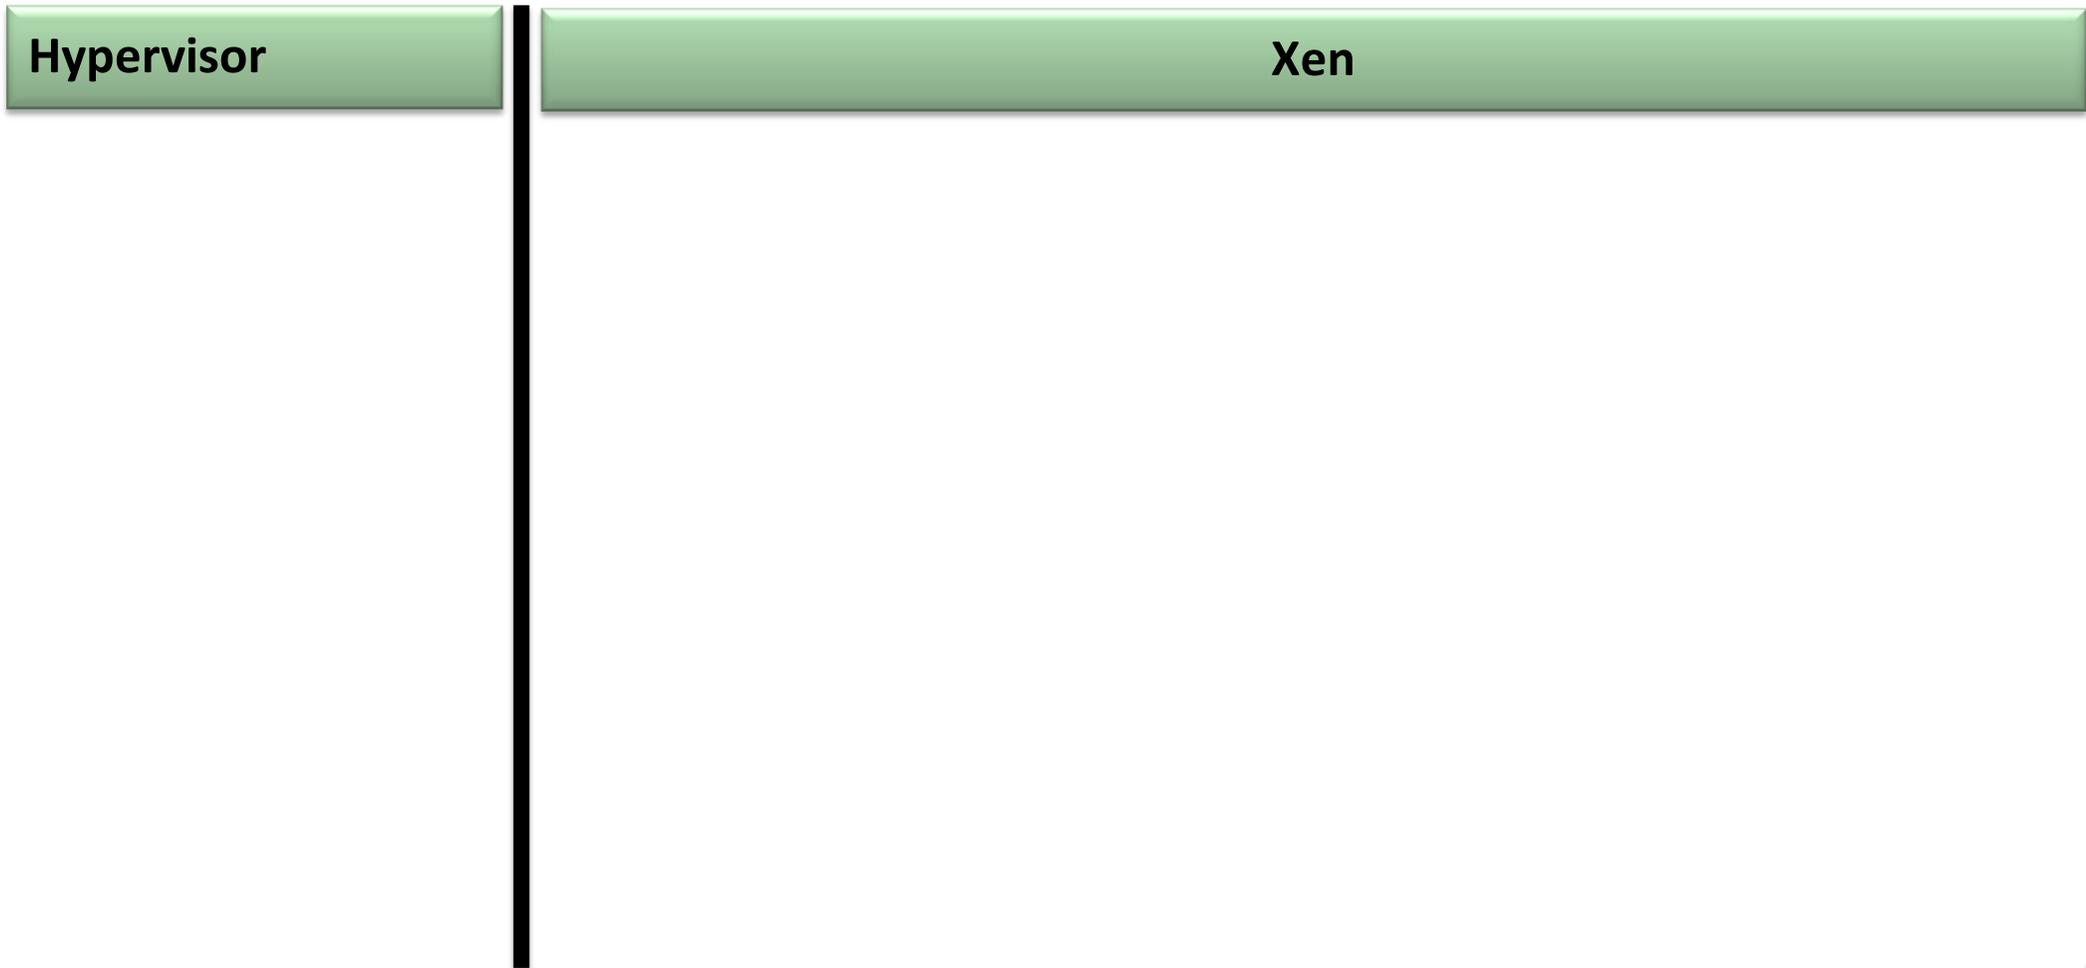
- Your apps

Driver/Stub/Service Domain(s)

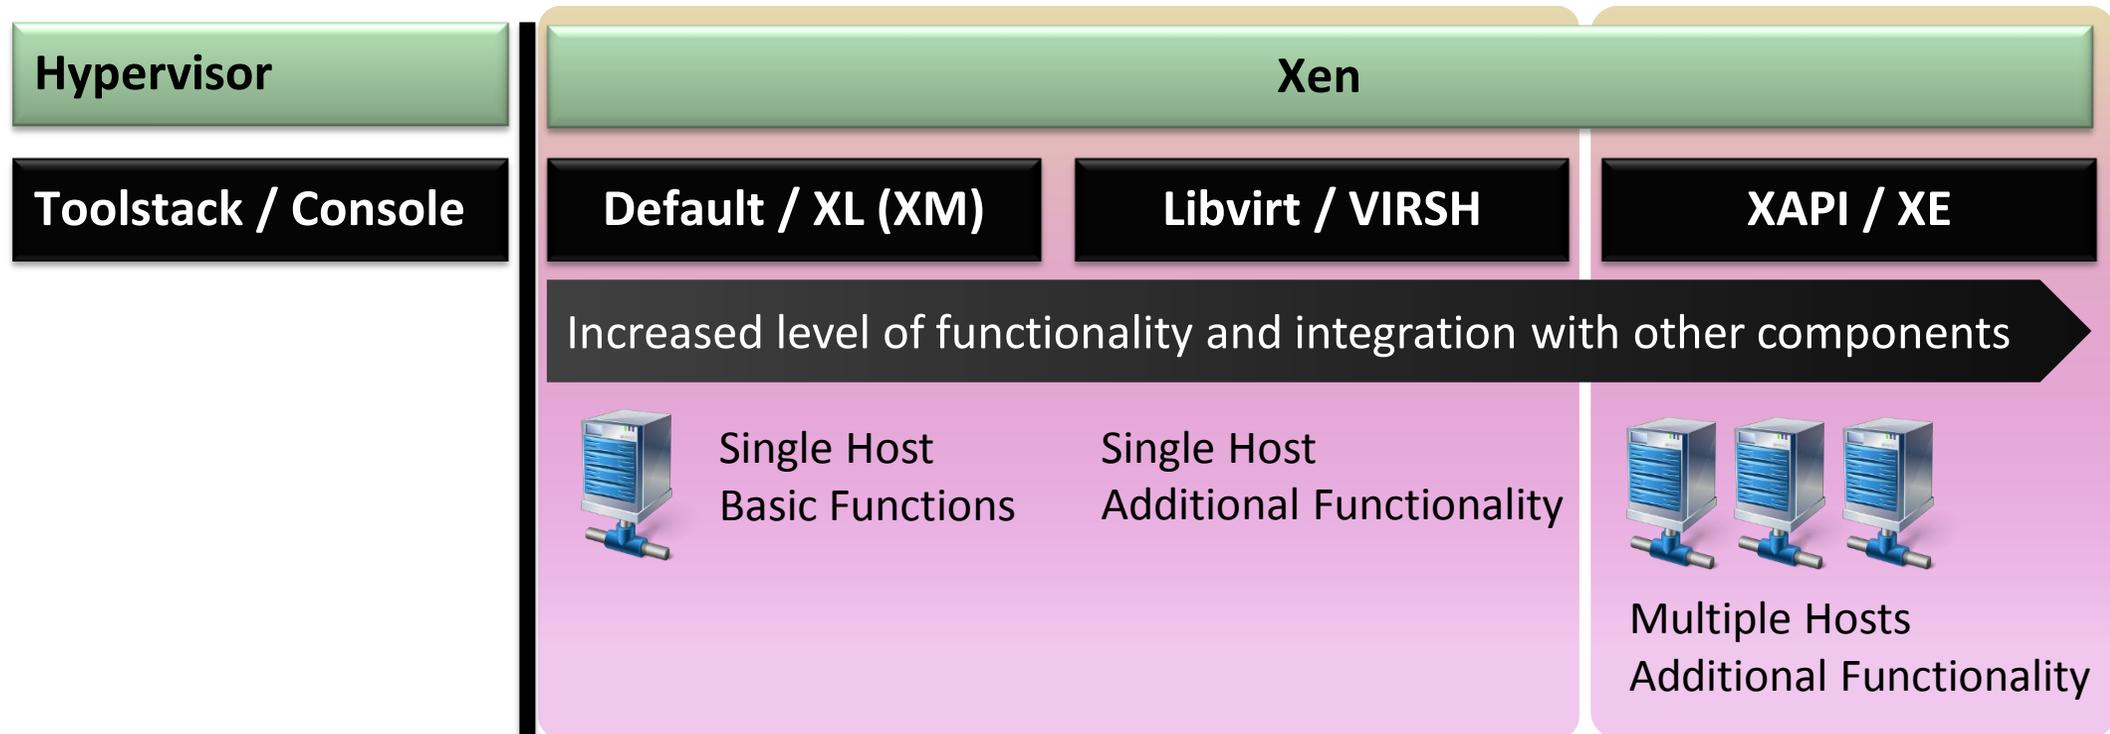
- A “driver, device model or control service in a box”
- De-privileged and isolated
- Lifetime: start, stop, kill



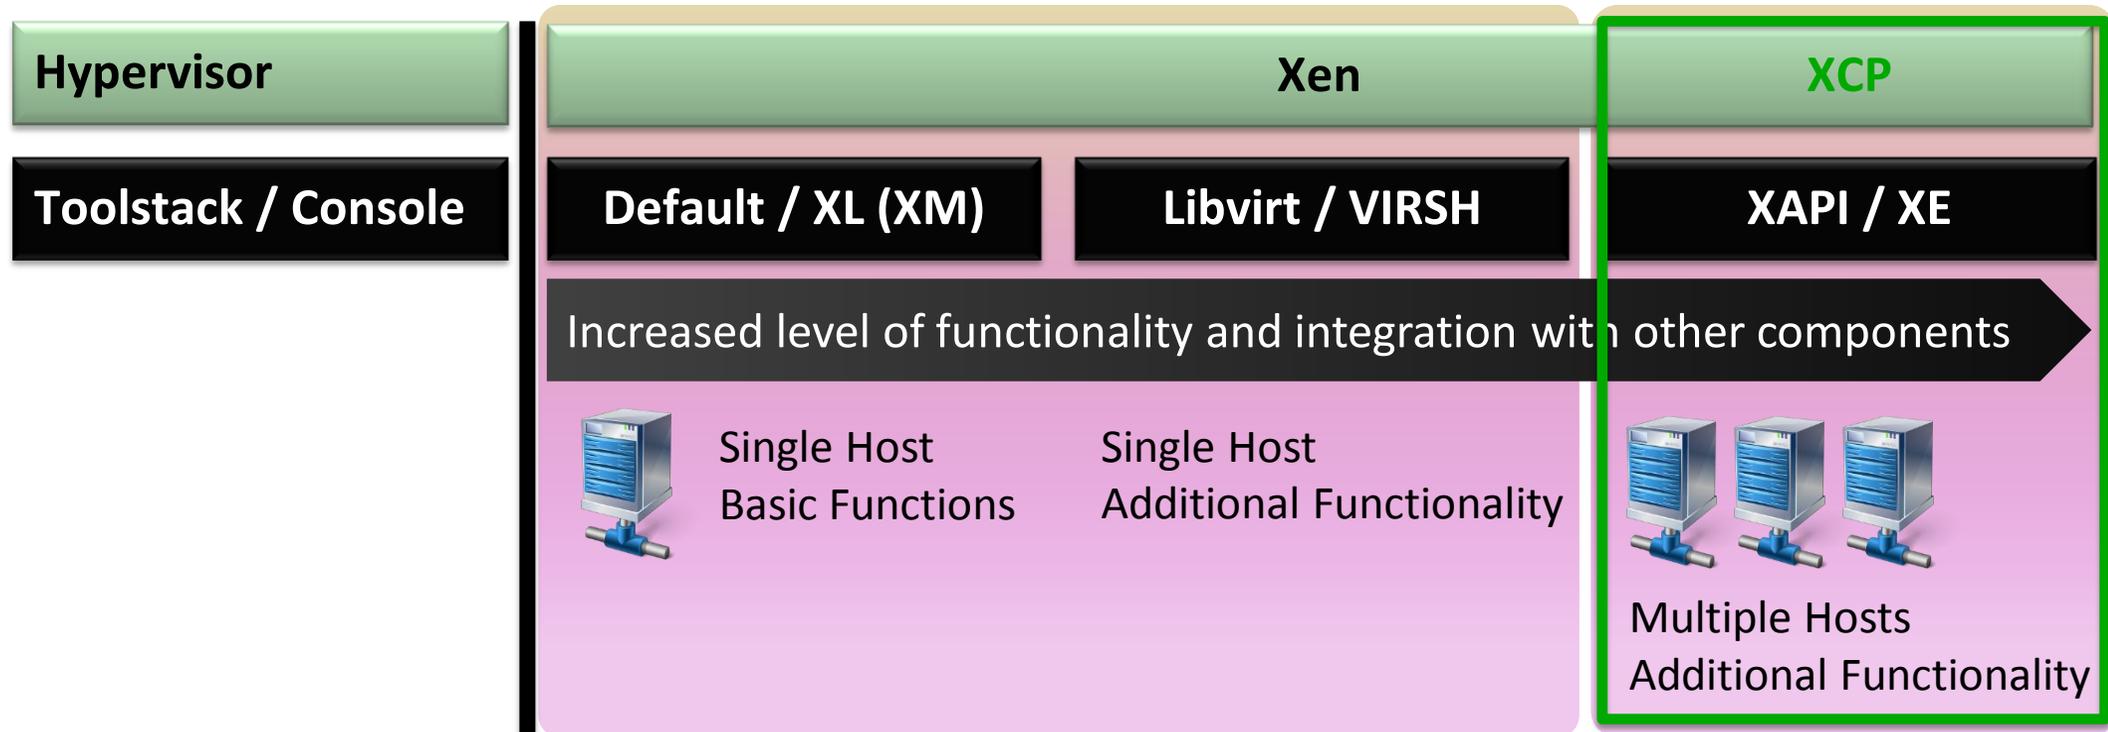
Xen Variants for Server & Cloud



Xen Variants for Server & Cloud



Xen Variants for Server & Cloud



Xen Variants for Server & Cloud

Project	Xen		XCP
Toolstack / Console	Default / XL (XM)	Libvirt / VIRSH	XAPI / XE
	Increased level of functionality and integration with other components		
Get Binaries from ...	Linux Distros	Linux Distros	Debian & Ubuntu ISO from Xen.org

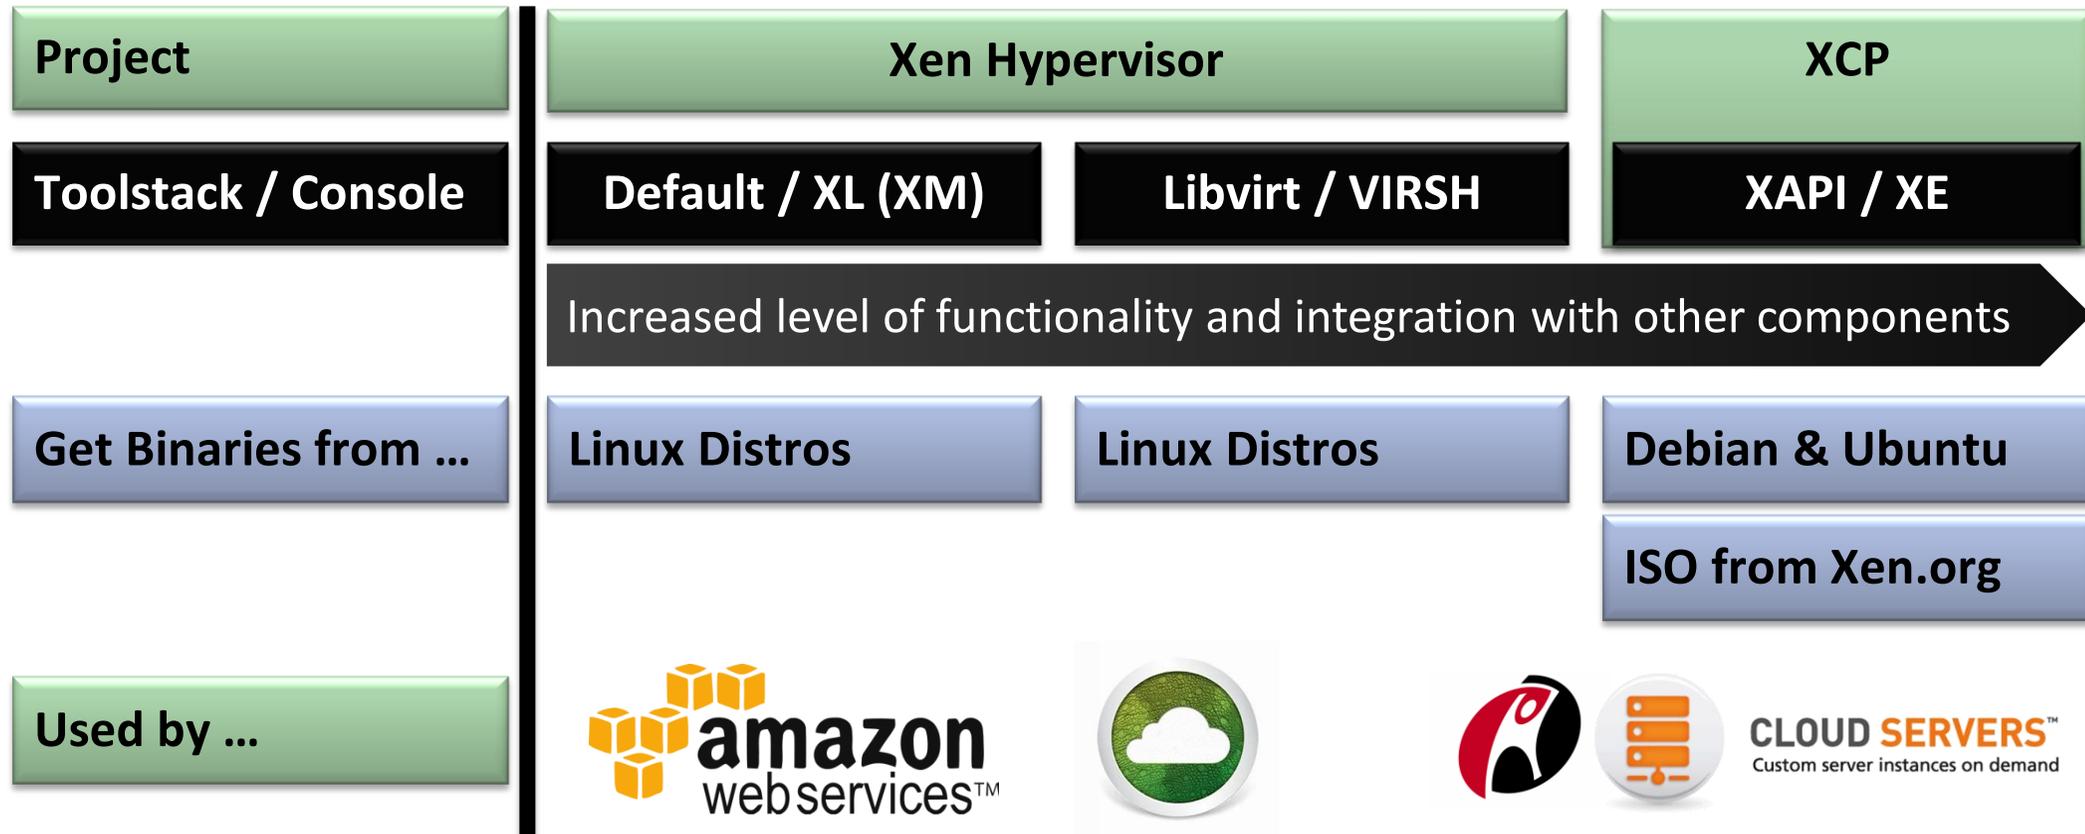


Xen Variants for Server & Cloud

Project	Xen Hypervisor		XCP
Toolstack / Console	Default / XL (XM)	Libvirt / VIRSH	XAPI / XE
	Increased level of functionality and integration with other components		
Get Binaries from ...	Linux Distros	Linux Distros	Debian & Ubuntu ISO from Xen.org
Products	Oracle VM	Huawei UVP	Citrix XenServer



Xen Variants for Server & Cloud



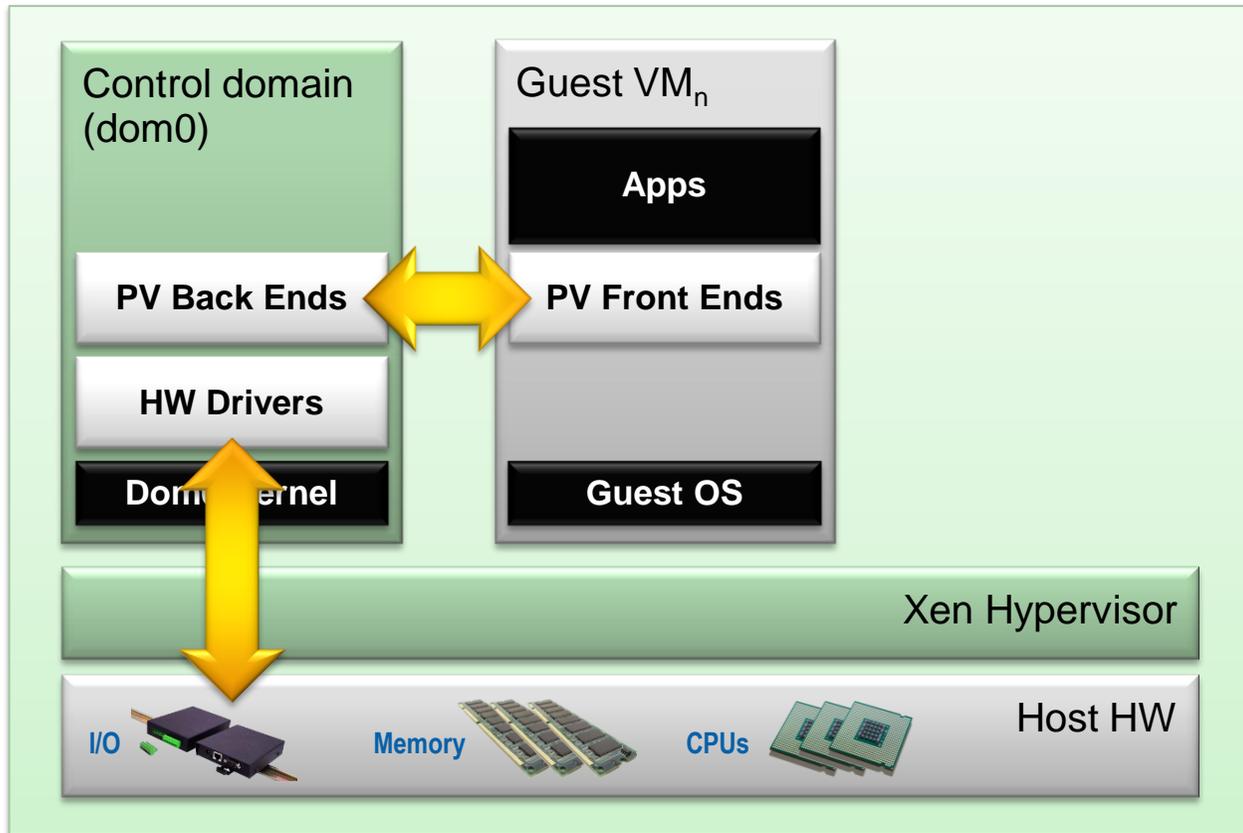
[More info:](http://xen.org/community/ecosystem.html) xen.org/community/ecosystem.html
xen.org/community/presentations.html
xen.org/products/case_studies.html



Xen : Types of Virtualization



PV Domains



Technology:

- Paravirtualization

Linux PV guests have limitations:

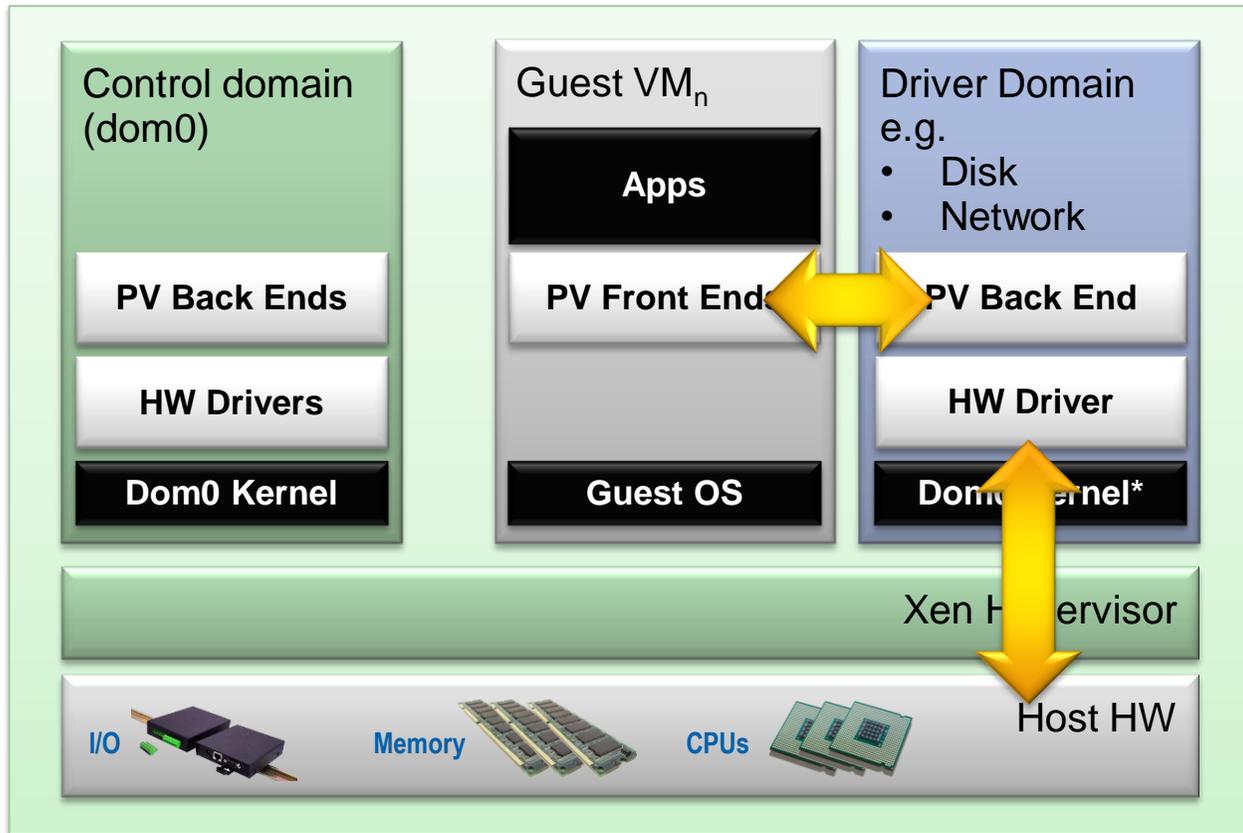
- limited set of virtual hardware

Advantages

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



PV Domains & Driver Domains



*) Can be MiniOS

Technology:

- Paravirtualization

Linux PV guests have limitations:

- limited set of virtual hardware

Advantages

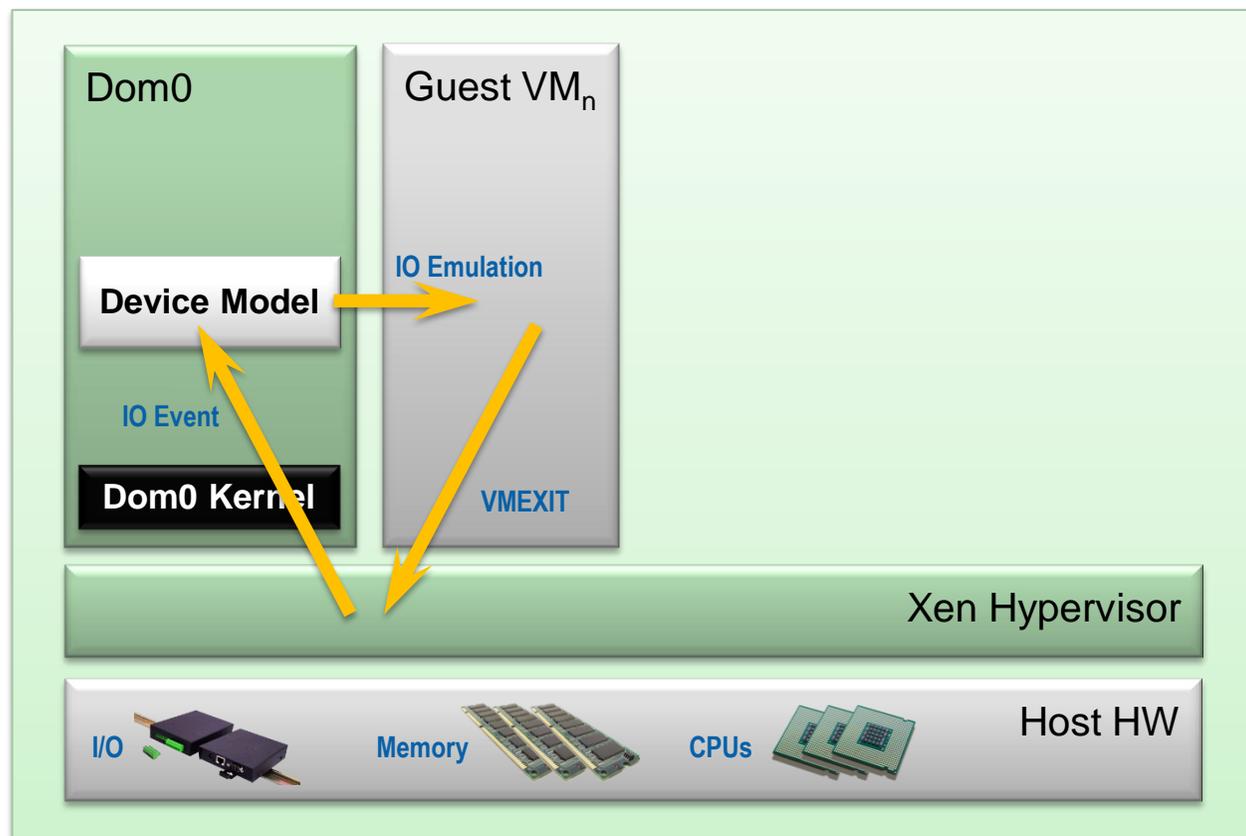
- Fast
- Works on any system
(even without virt extensions)

Driver Domains

- Security
- Isolation
- Reliability and Robustness



HVM & Stub Domains



Technology:

- Shows emulation using QEMU/Device Model (SW Virtualization)

- In other situation HW can be used

Disadvantages

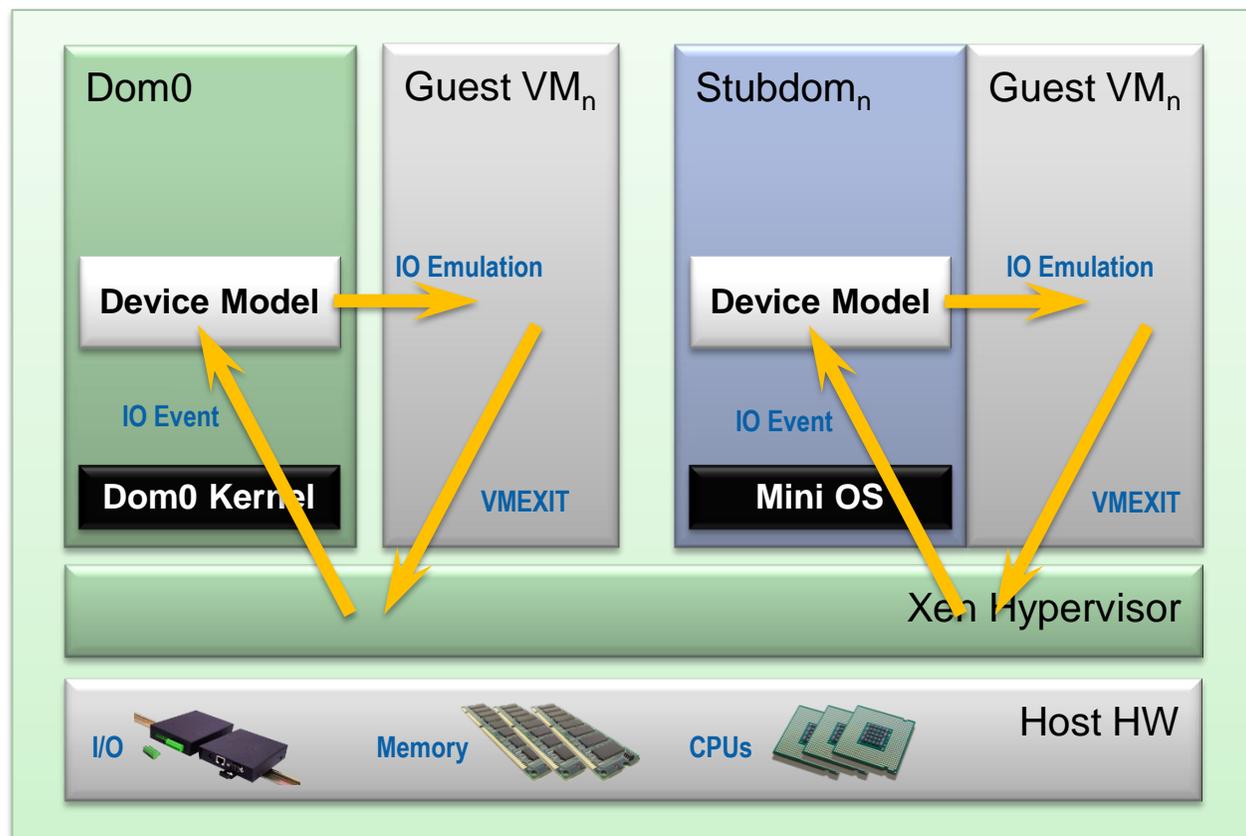
- Emulation slower than PV (mainly I/O devices)

Advantages

- No kernel support needed



HVM & Stub Domains



Technology:

- Shows emulation using QEMU/Device Model (SW Virtualization)

- In other situation HW can be used

Disadvantages

- Emulation slower than PV (mainly I/O devices)

Advantages

- No kernel support needed

Stub Domains

- Security
- Isolation
- Reliability and Robustness

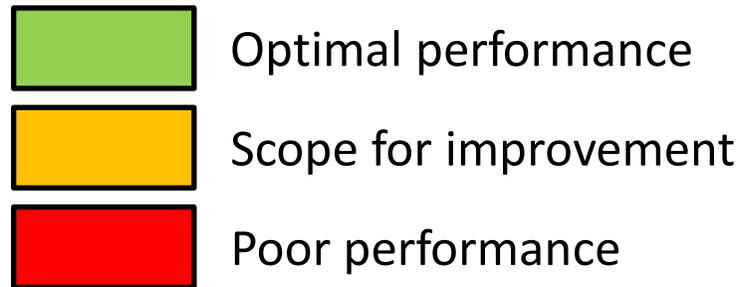


The Virtualization Spectrum

VS	Virtualized (SW)
VH	Virtualized (HW)
P	Paravirtualized

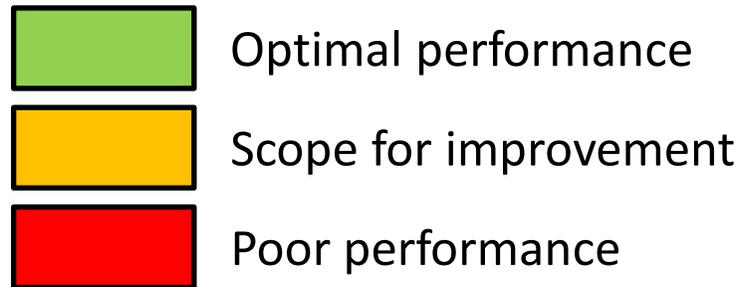
	Disk and Network	Interrupts, Timers	Emulated Motherboard, Legacy boot	Privileged Instructions and page tables	
Fully Virtualized (FV)	VS	VS	VS	VH	} HVM mode/domain
FV with PV for disk & network	P	VS	VS	VH	
PVHVM	P	P	VS	VH	
PVH NEW Xen 4.3	P	P	P	VH	} PV mode/domain
Fully Paravirtualized (PV)	P	P	P	P	

The Virtualization Spectrum



	Disk and Network	Interrupts, Timers	Emulated Motherboard, Legacy boot	Privileged Instructions and page tables	
Fully Virtualized (FV)	VS	VS	VS	VH	HVM mode/domain
FV with PV for disk & network	P	VS	VS	VH	
PVHVM	P	P	VS	VH	
PVH NEW Xen 4.3	P	P	P	VH	PV mode/domain
Fully Paravirtualized (PV)	P	P	P	P	

The Virtualization Spectrum



Important: Xen automatically picks the best option based on HW & OS capabilities and available drivers.
 As a Xen user I chose a HVM or PV domain.

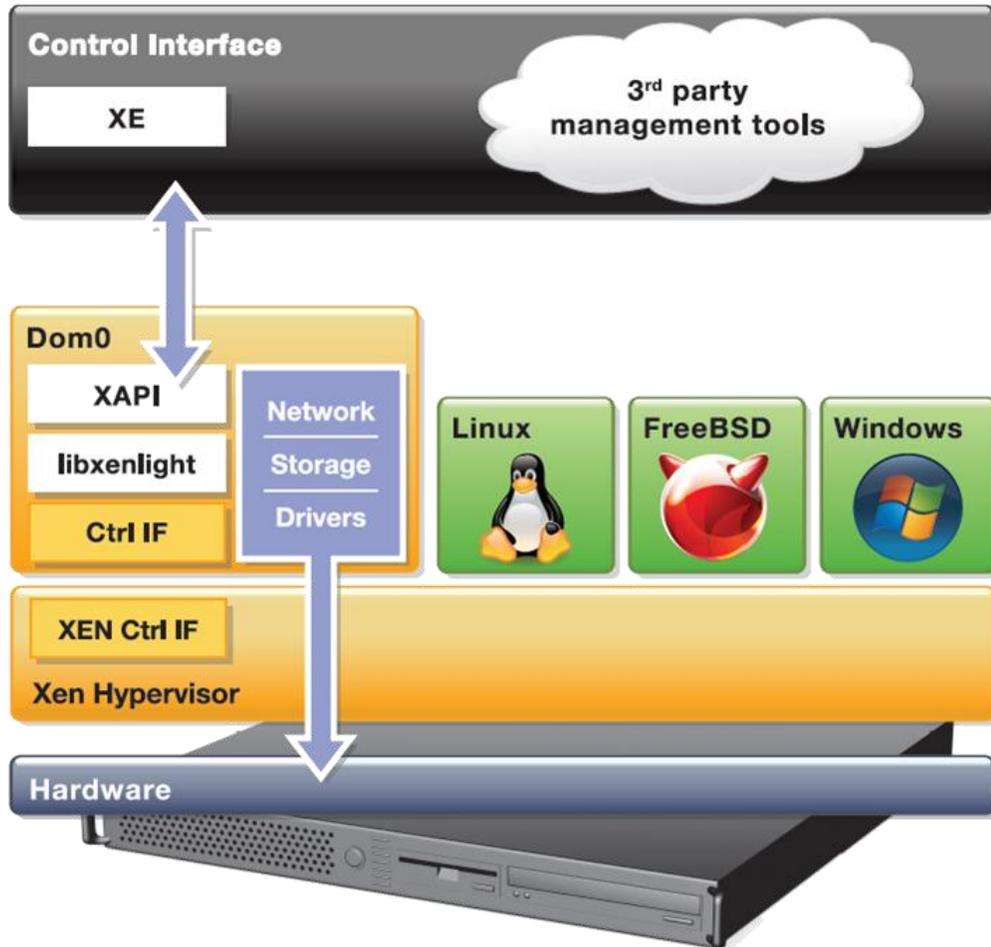
	Disk and Network	Interrupts, Tim	Emulated Mo	Legacy boot	Privileged Ins and page tar	
Fully Virtualized (FV)	VS	VS	VS	VH	} HVM mode/domain	
FV with PV for disk & network	P	VS	VS	VH		
PVHVM	P	P	VS	VH		
PVH NEW Xen 4.3	P	P	P	VH	} PV mode/domain	
Fully Paravirtualized (PV)	P	P	P	P		



XCP Project



XCP – Xen Cloud Platform



Complete stack for server virtualization

- Extends Xen to cover multiple hosts
- Adds further functionality and integrations for cloud, storage and networking to Xen HV
- GPLv2
- XenServer is a commercial XCP distro

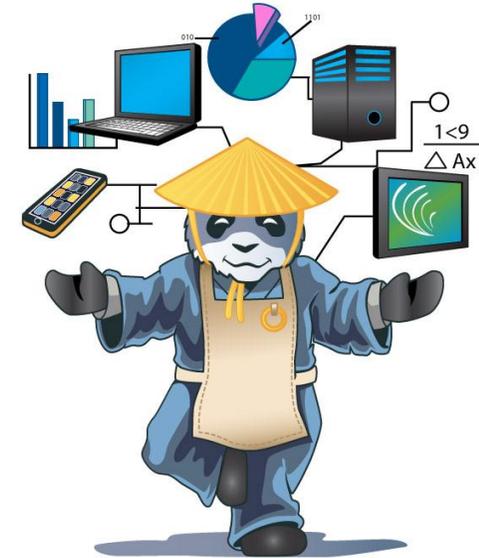
Two Flavours

- Appliance (ISO using CentOS Dom0)
- Packages in Debian & Ubuntu (more distros to come)



Major XCP Features

- VM lifecycle: live snapshots, checkpoint, migration
- Resource pools: flexible storage and networking
- Event tracking: progress, notification
- Upgrade and patching capabilities
- Real-time performance monitoring and alerting
- Built-in support and templates for Windows and Linux guests
- Open vSwitch support built-in (default)



[More info:](http://wiki.xen.org/wiki/XCP_Release_Features) wiki.xen.org/wiki/XCP_Release_Features



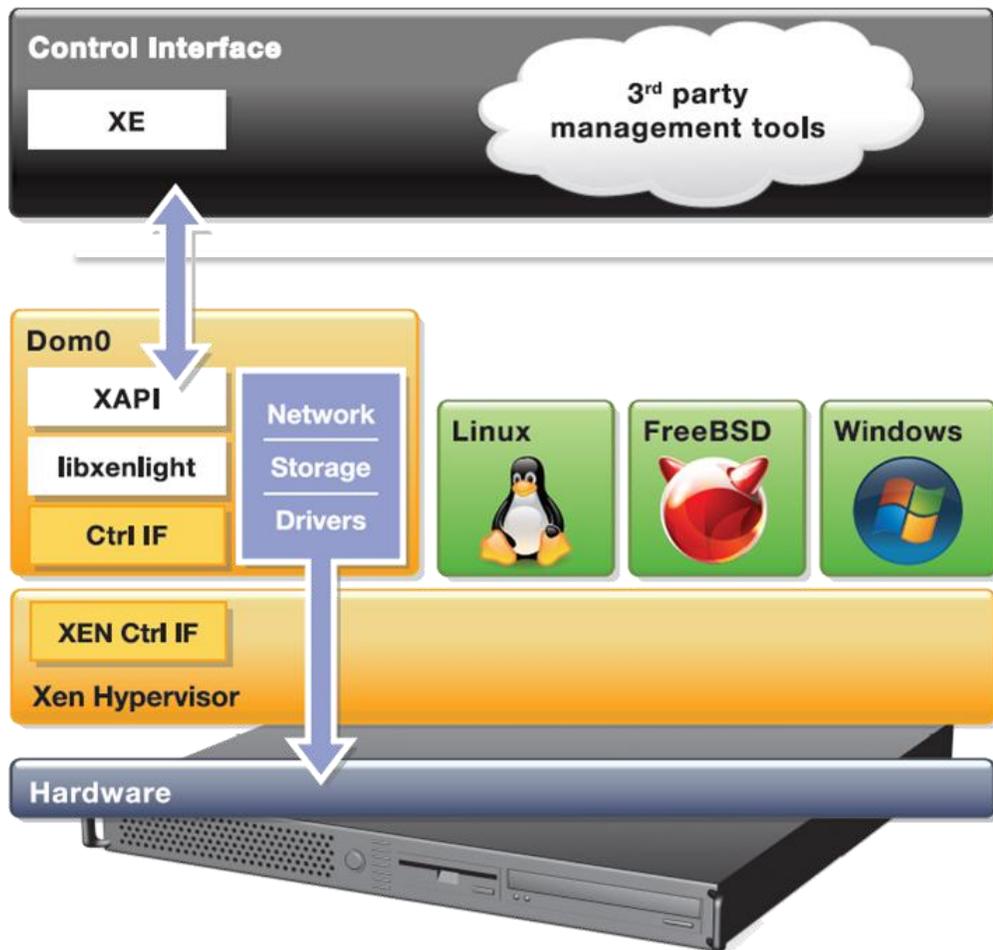
XCP 1.6

- **New format Windows drivers:**
installable by Windows Update Service
- **Networking:** Better VLAN scalability, LACP bonding, IPv6
- **Storage XenMotion:**
 - Migrate VMs between hosts or pools without shared storage
 - Move a VM's disks between storage repositories while the VM is running
- **Other: more templates, latest Xen, OVS, etc.**

[More info:](http://xen.org/download/xcp/releasenotes_1.6.0.html) xen.org/download/xcp/releasenotes_1.6.0.html &
xen.org/download/xcp/index_1.6.0.html



XCP and Cloud Orchestration Stacks



apache **cloudstack**
open source cloud computing

OpenNebula.org



Challenges for FOSS hypervisors



**“Security and QoS/Reliability are amongst
the top 3 blockers for cloud adoption”**

www.colt.net/cio-research



System characteristics cloud users care about: “Robustness, Performance, Scalability & Security”

Results XCP User Survey 2012 – 90% of users quoted these as most important attributes



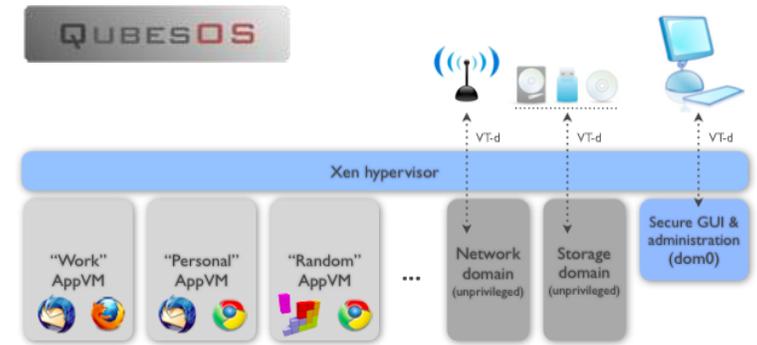
Disaggregation

Split Control Domain into Driver, Stub and Service Domains

- See: "[Breaking up is hard to do](#)" @ [Xen Papers](#)
- See: "[Domain 0 Disaggregation for XCP and XenServer](#)"

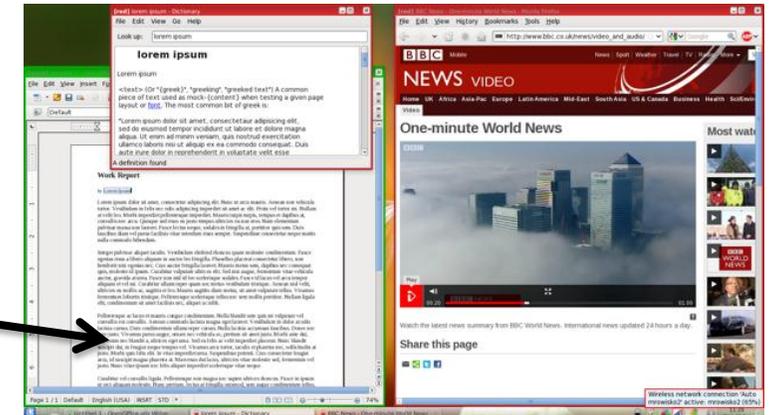
Used today by [Qubes OS](#) and Citrix XenClient XT

Prototypes for XCP



See qubes-os.org

Different windows run in different VMs



Benefits of Disaggregation

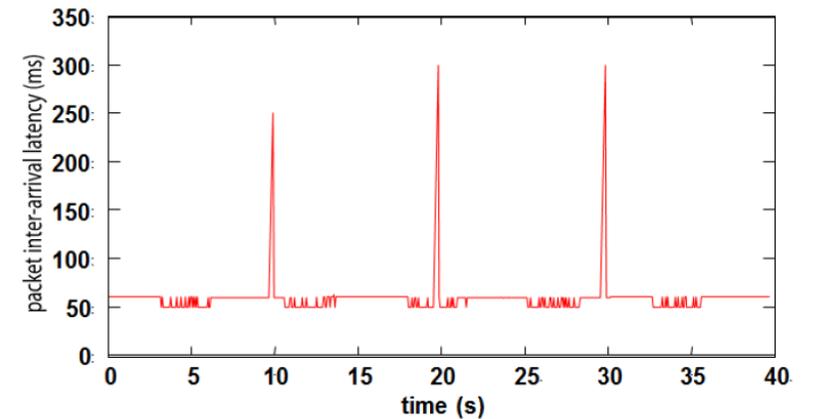
More Security

Increased serviceability and flexibility

Better Robustness

Better Performance

Better Scalability

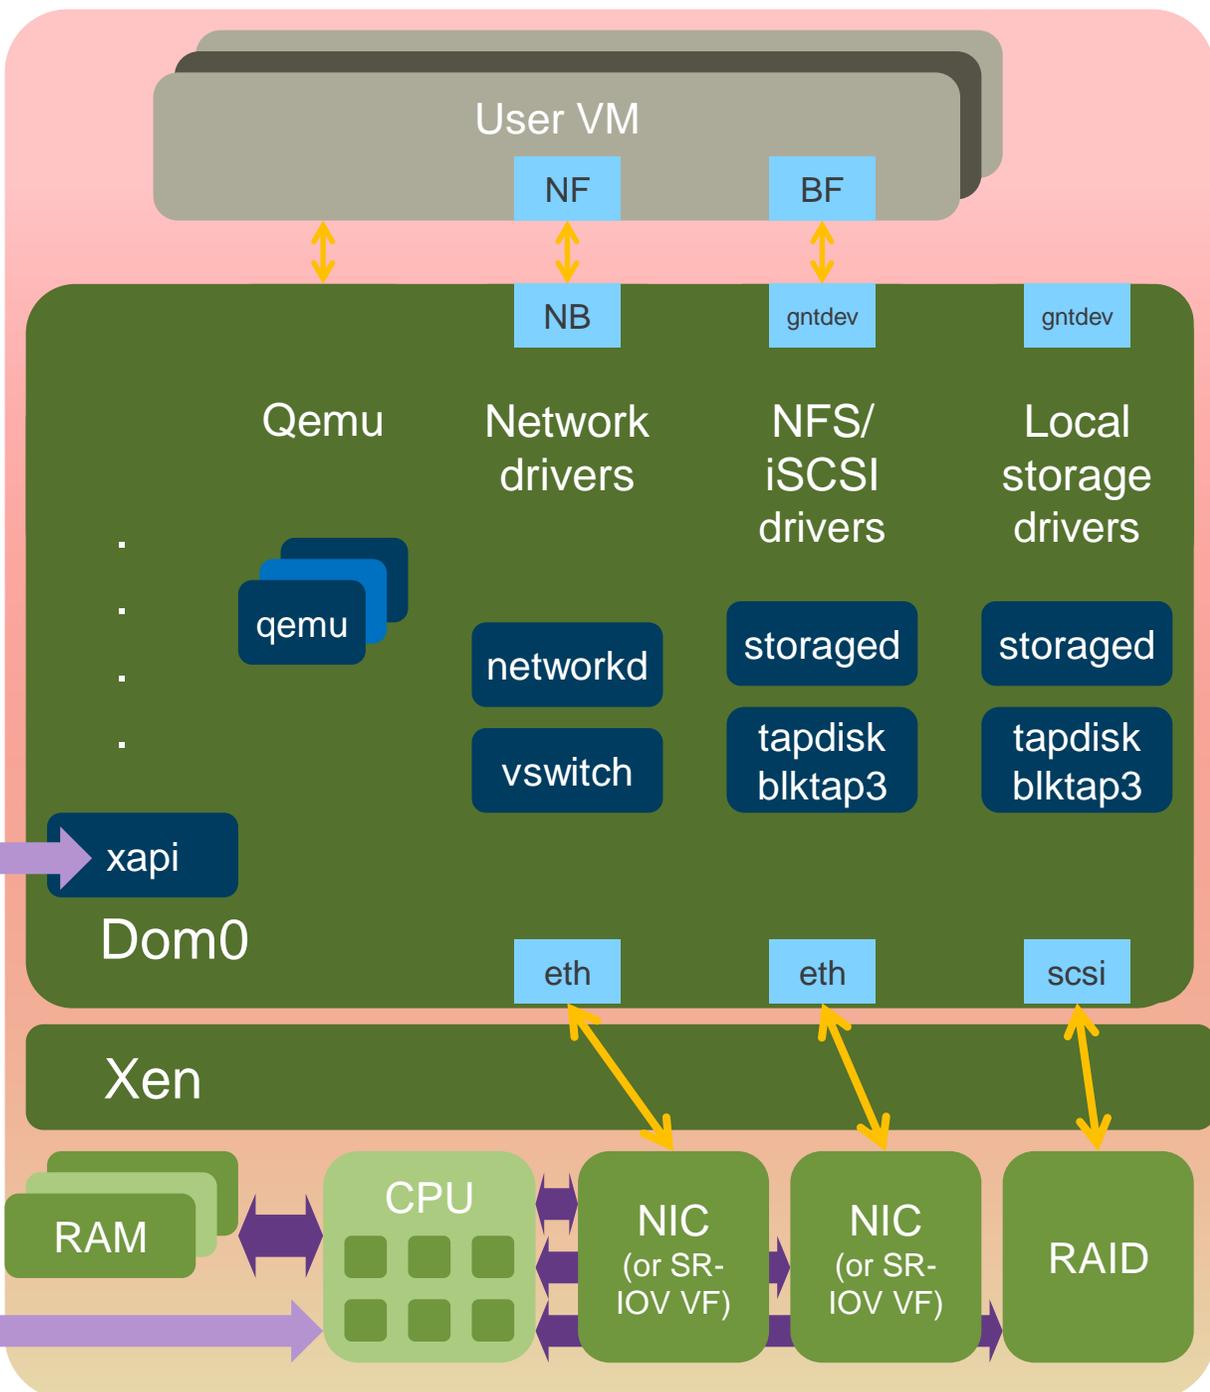
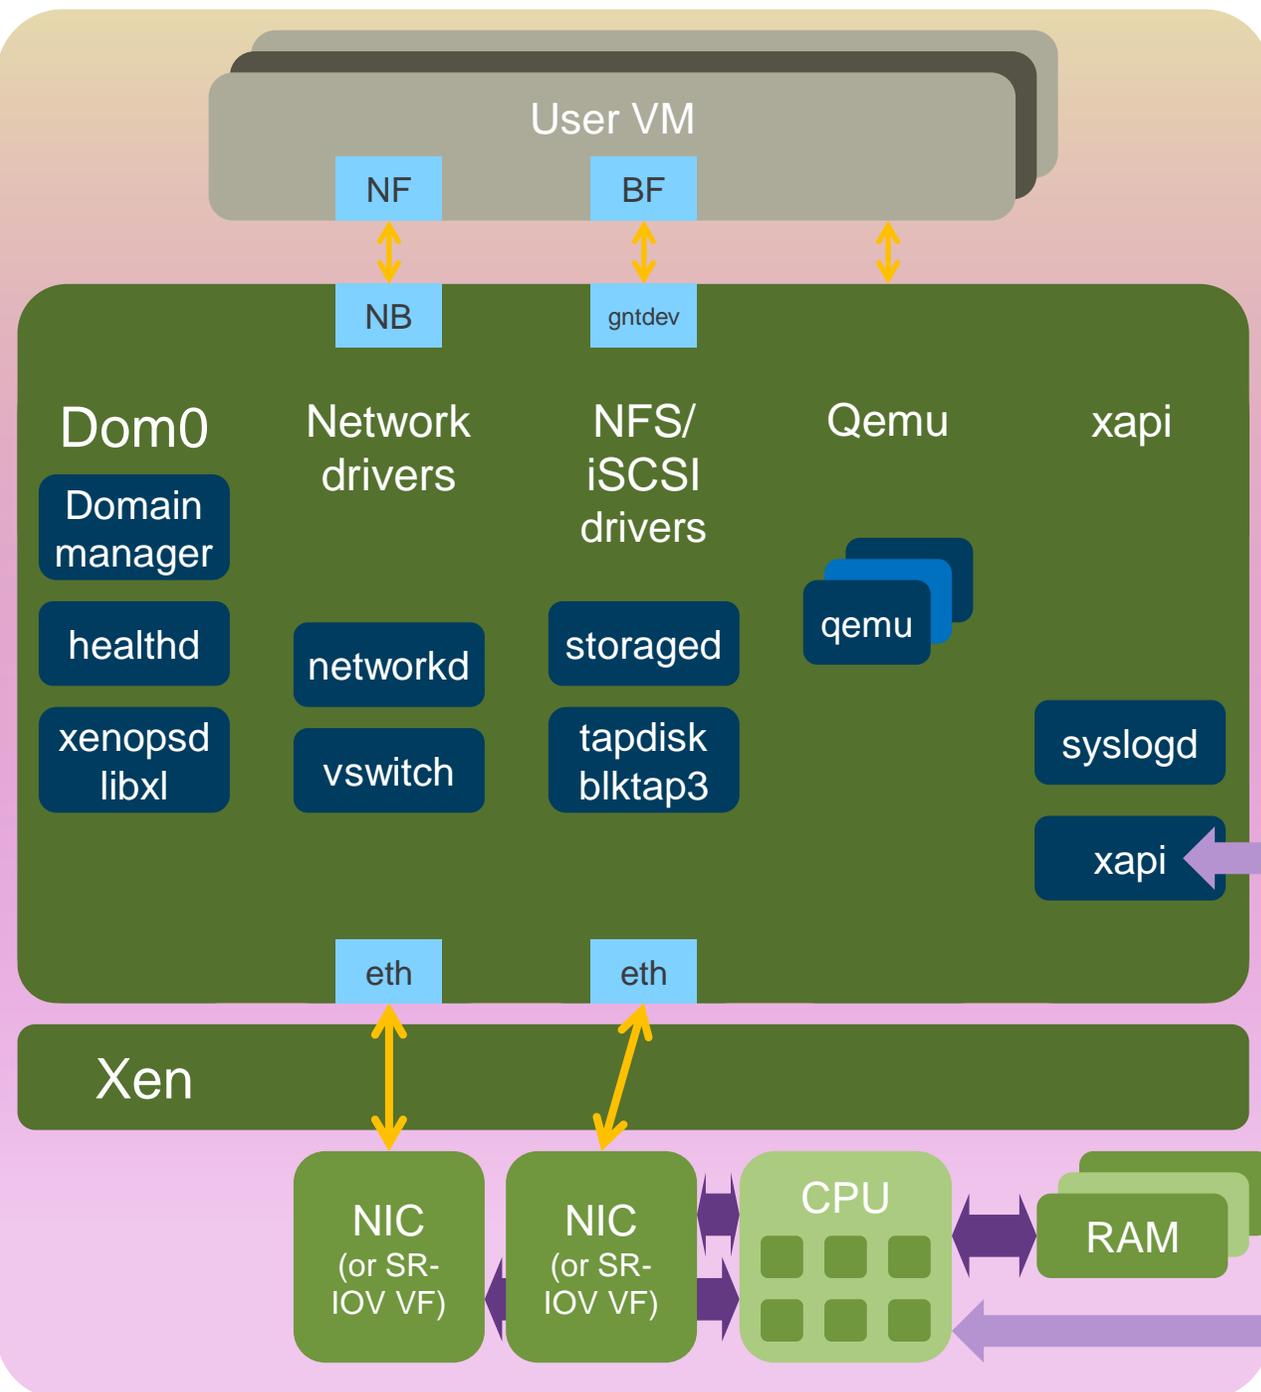


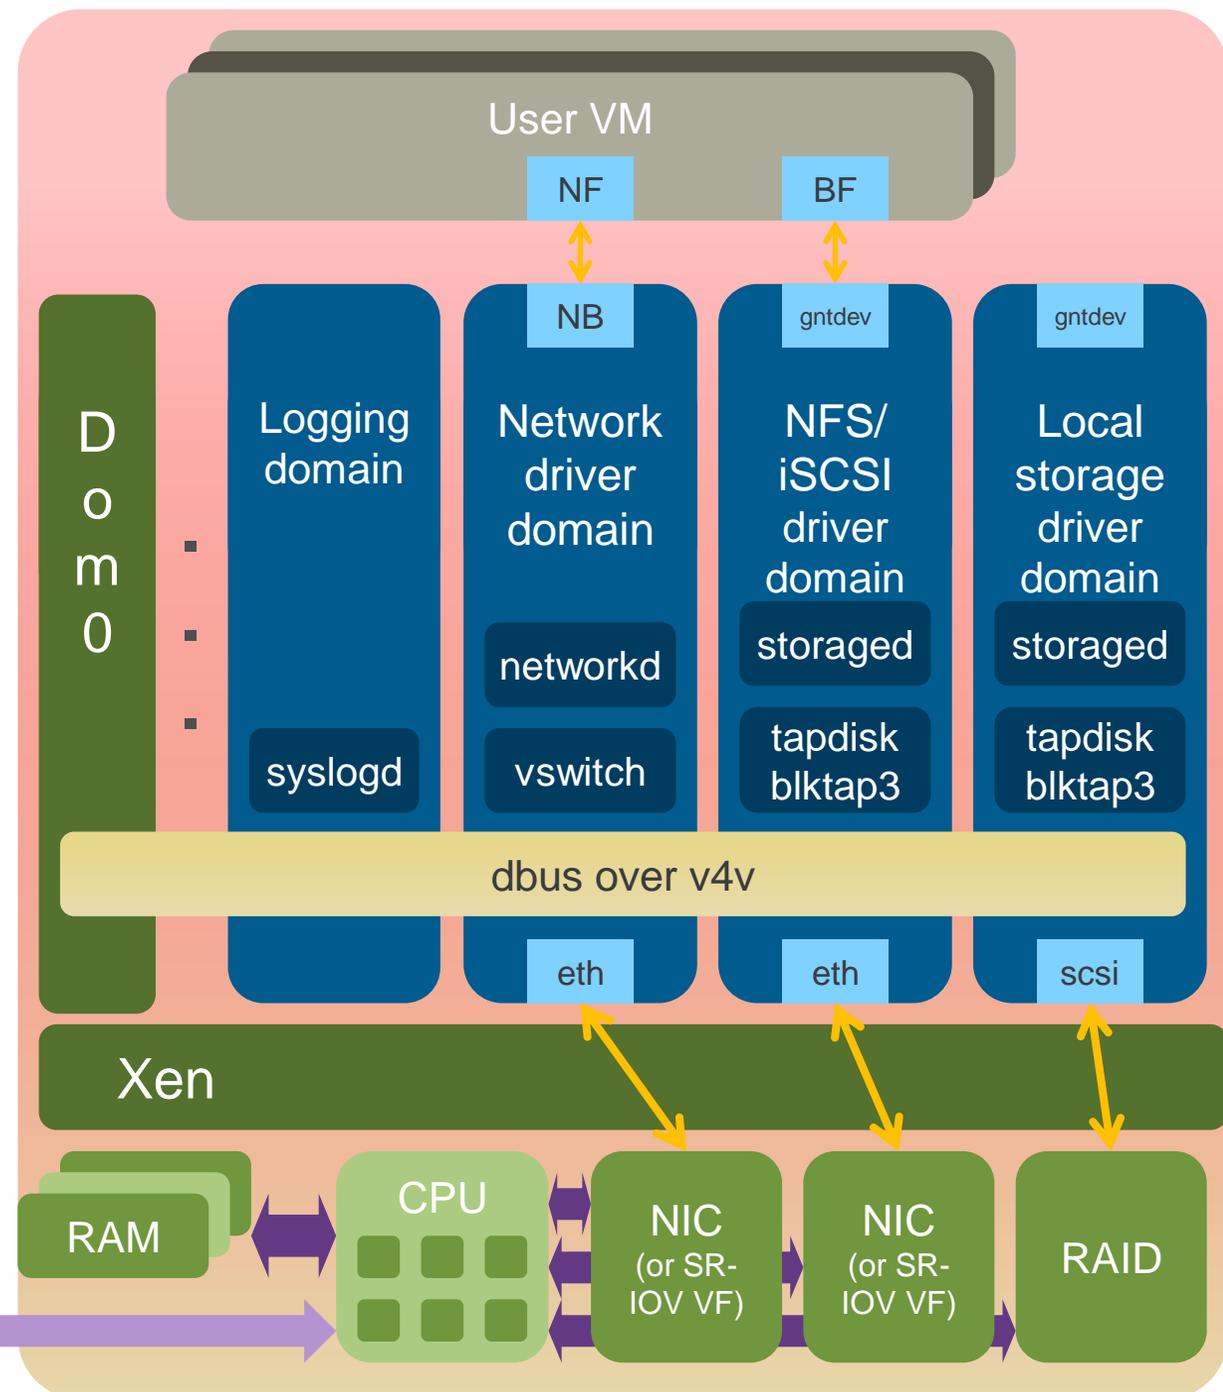
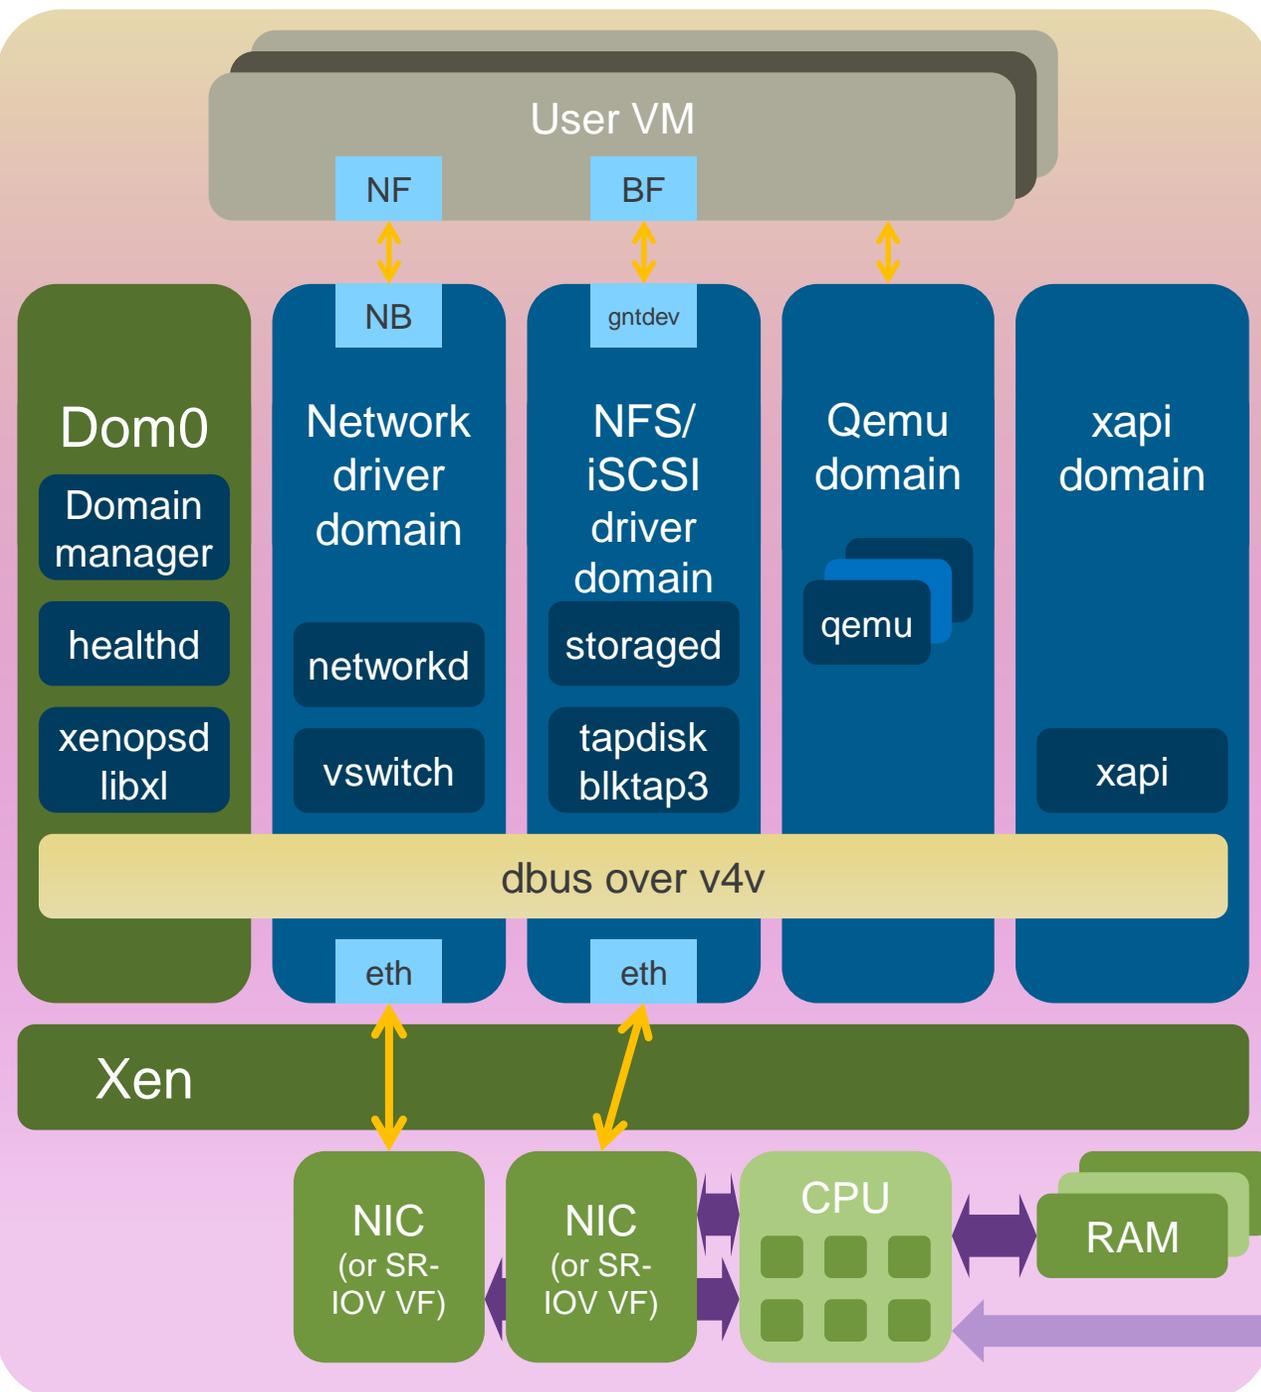
Ability to safely restart parts of the system
(e.g. just 275ms outage from failed Ethernet driver)



**Next: XCP Architecture Diagram
Before and After Disaggregation**





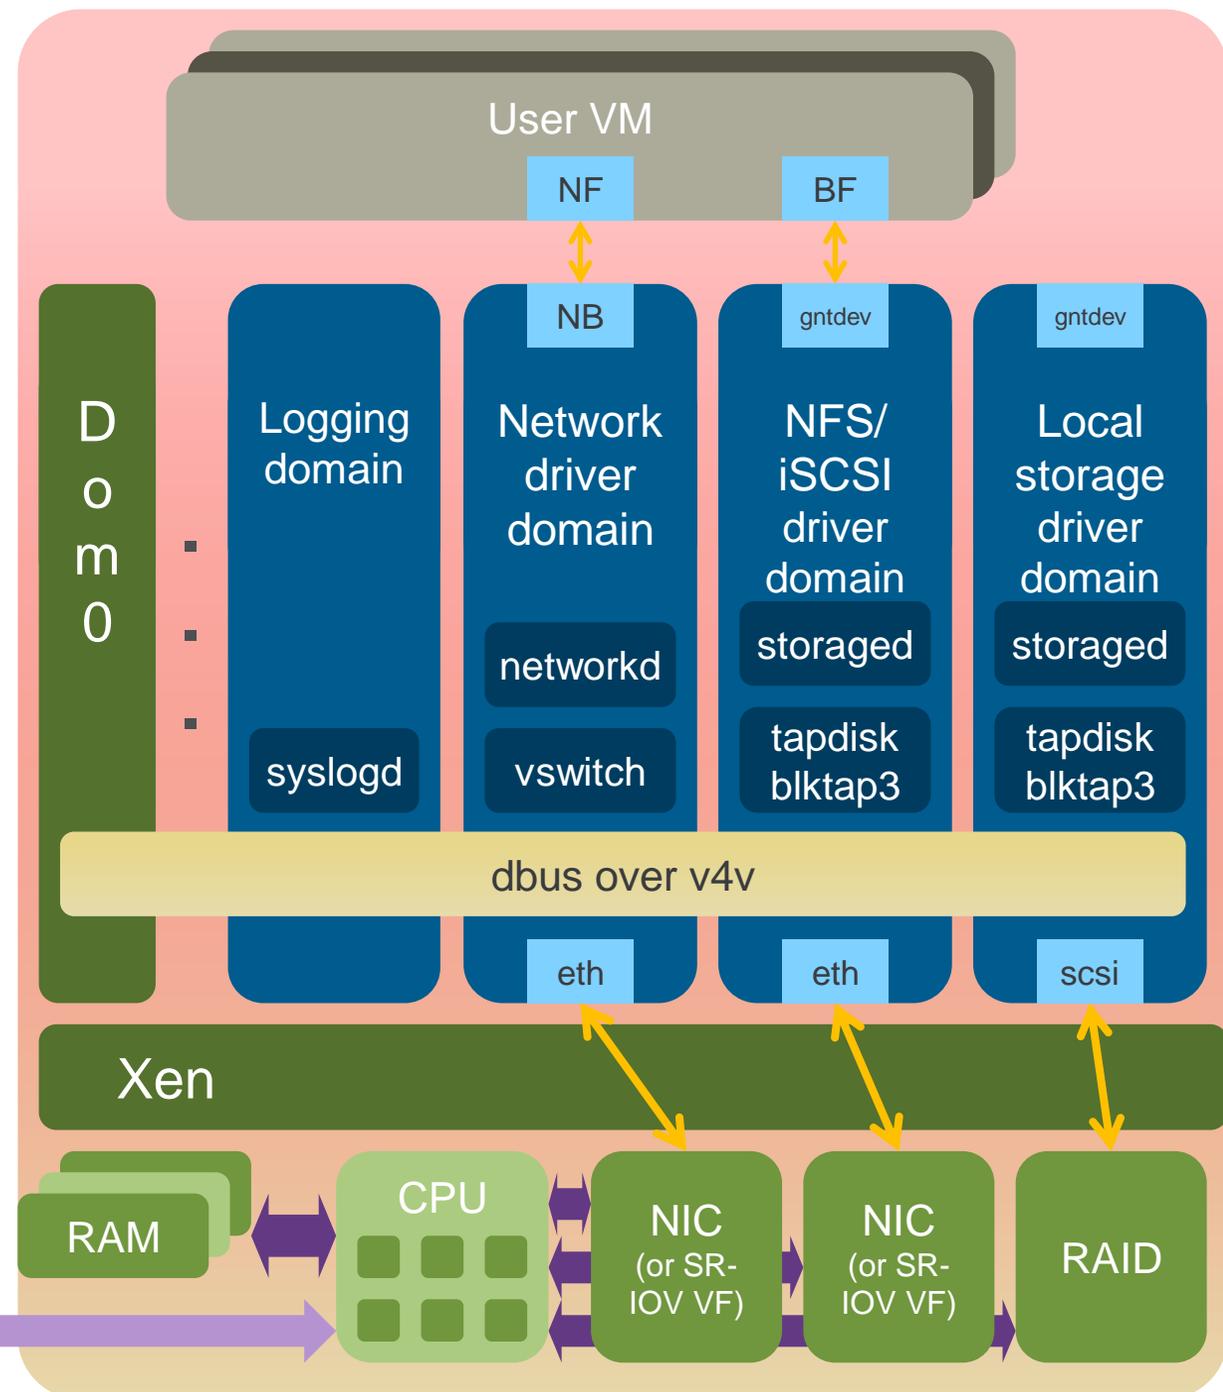
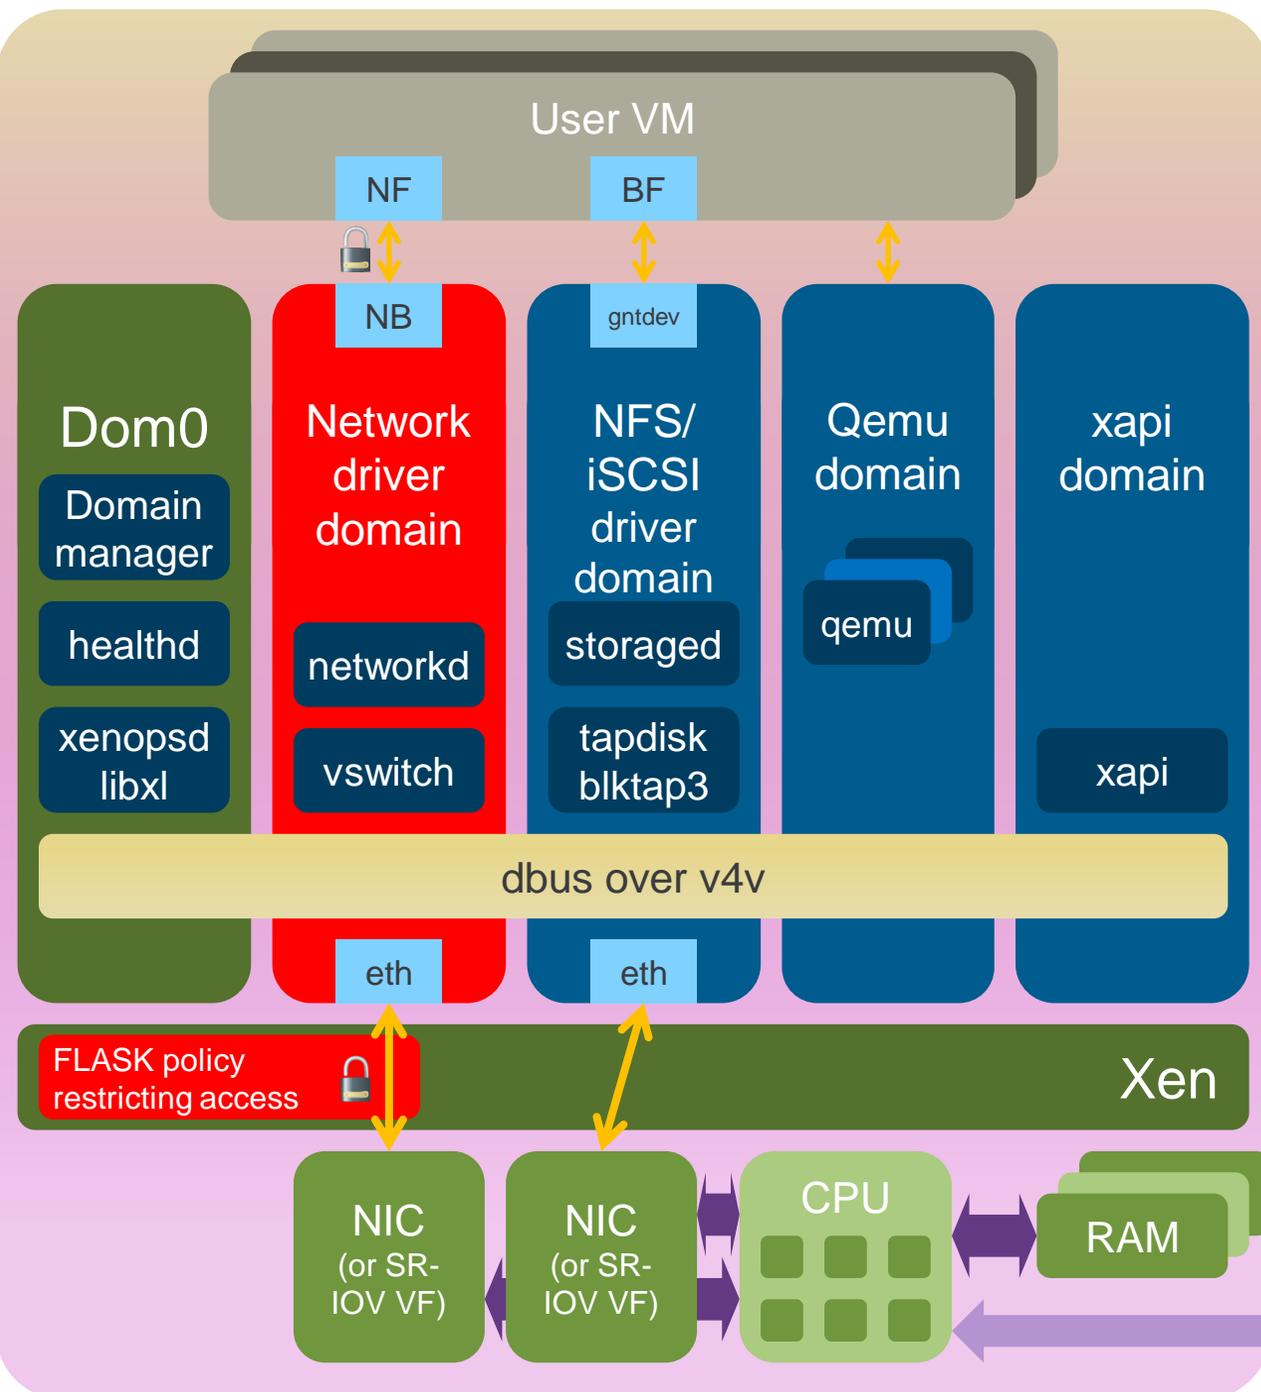


Xen Security Advantages

- Even without Advanced Security Features
 - Well-defined trusted computing base (much smaller than on type-2 HV)
 - Minimal services in hypervisor layer
- Xen Security Modules (or XSM) and FLASK
 - XSM is Xen equivalent of LSM
 - FLASK is Xen equivalent of SELinux
 - Developed, maintained and contributed to Xen by NSA
 - Compatible with [SELinux](#) (tools, architecture)
 - XSM object classes maps onto Xen features

[More info:](http://www.slideshare.net/xen_com_mgr/a-brief-tutorial-on-xens-advanced-security-features) http://www.slideshare.net/xen_com_mgr/a-brief-tutorial-on-xens-advanced-security-features





News from the Xen Community



Coming in Xen 4.3 (Q2 2013)

- PVH virtualization mode
- Extend scope of Xen Security Modules
- qxl Spice support for 3d acceleration
- Updated and improved libvirt drivers for Xen
- Lots of other stuff:
 - scalability, performance, better NUMA support, ...

[More info:](http://blog.xen.org/index.php/2013/02/11/xen-4-3-mid-release-roadmap-update) blog.xen.org/index.php/2013/02/11/xen-4-3-mid-release-roadmap-update



Xen 4.3 for ARM Servers

Fully functional for ARM v7 & v8

ARM v7: Versatile Express, Arndale &
Samsung Chromebook

ARM v8: Fast Model



Xen and ARM : a perfect Match

ARM SOC

Device Tree describes ...



GT

GIC
v2

2 stage
MMU

ARM Architecture Features for Virtualization

User mode : EL0

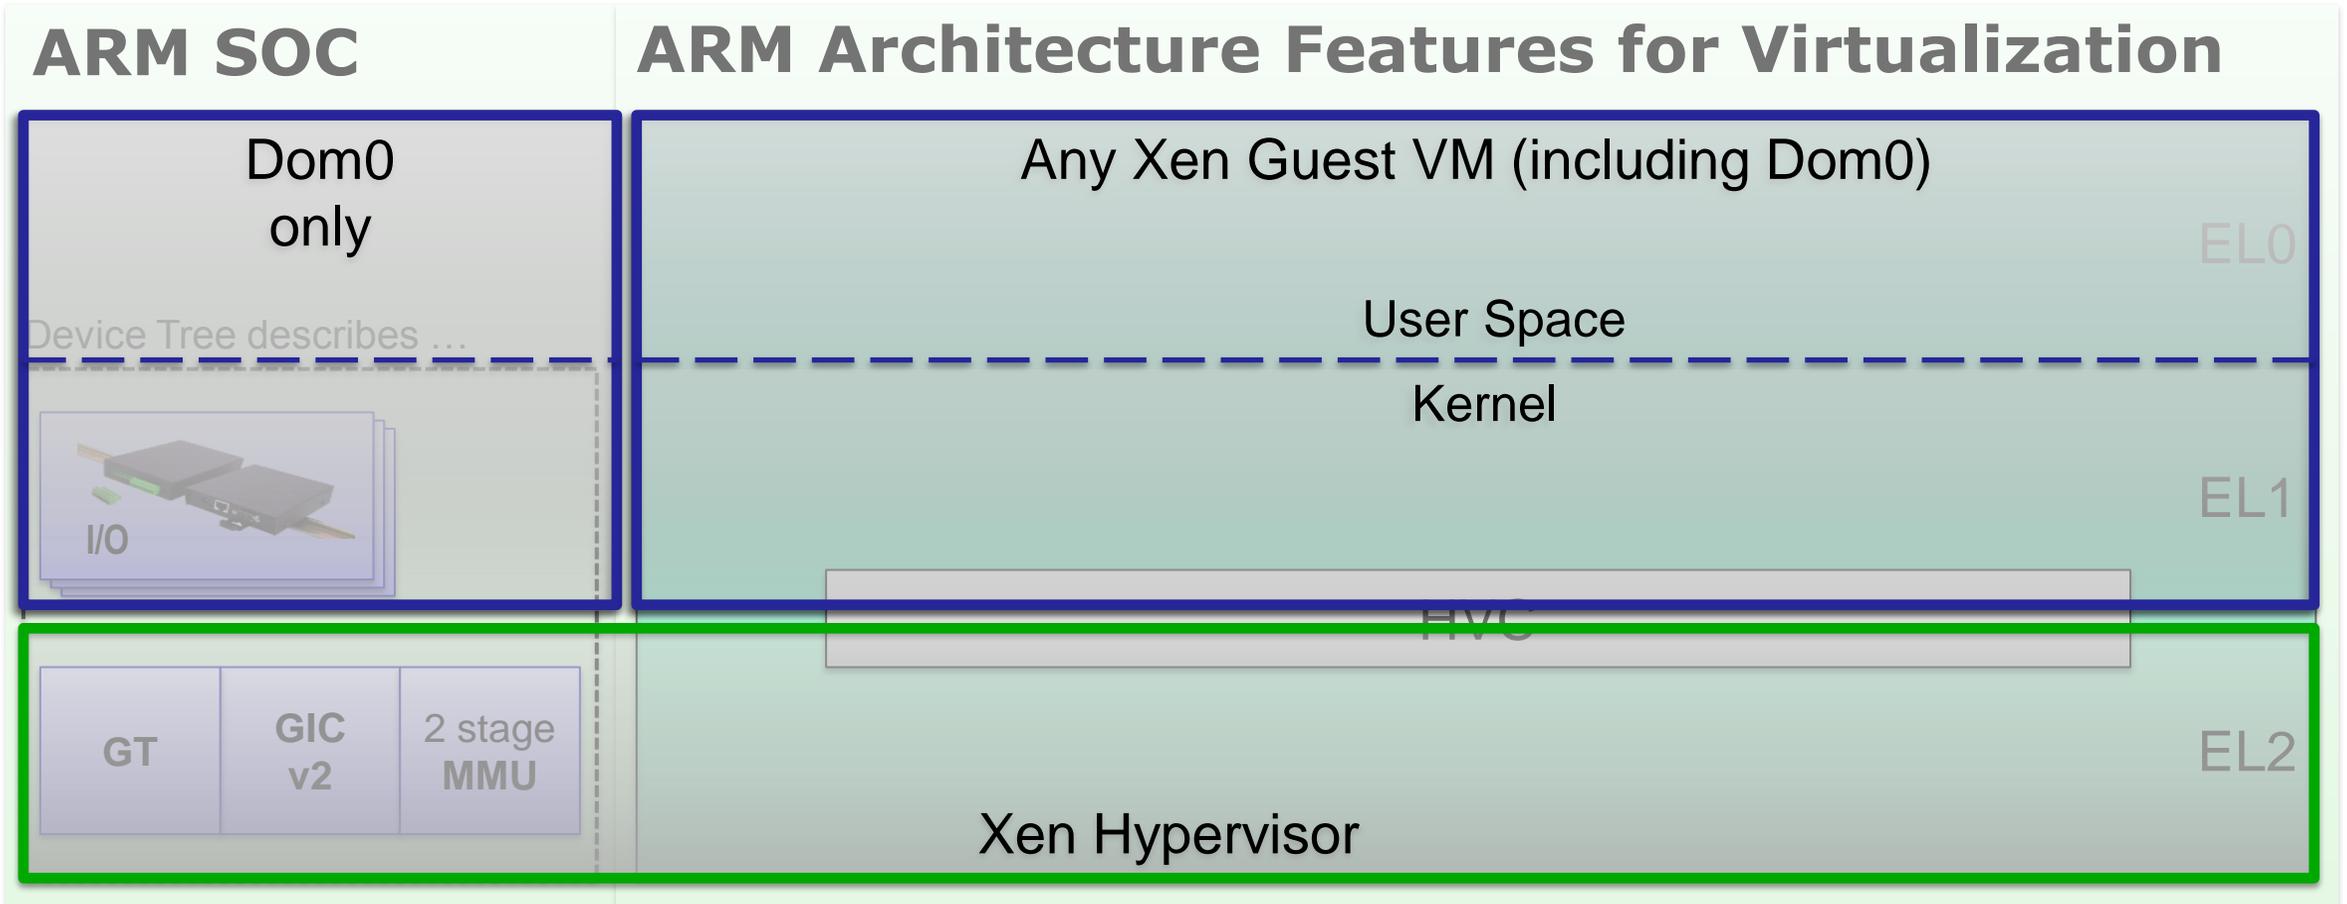
Kernel mode : EL1

Hypercall interface :HVC

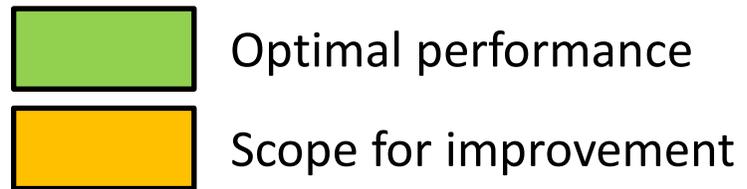
Hypervisor mode : EL2



Xen and ARM : a perfect Match



One mode to rule them all



	<div style="display: flex; justify-content: space-around; text-align: center;"> <div style="border: 1px solid black; padding: 5px; transform: rotate(-45deg);">Disk and Network</div> <div style="border: 1px solid black; padding: 5px; transform: rotate(-45deg);">Interrupts, Timers</div> <div style="border: 1px solid black; padding: 5px; transform: rotate(-45deg);">Emulated Motherboard, Legacy boot</div> <div style="border: 1px solid black; padding: 5px; transform: rotate(-45deg);">Privileged Instructions and page tables</div> </div>				
x86: PVHVM	P	P	VS	VH	HVM mode/domain PV mode/domain
x86: PVH	P	P	P	VH	
ARM v7 & v8	P	VH	VH	VH	



Xen in CentOS 6.4+

Xen is coming back to CentOS

In semi-private beta

Planned release in CentOS 6.4

Include XAPI packages – aka XCP in CentOS



Xen Library Operating Systems

Application stacks only running on Xen APIs

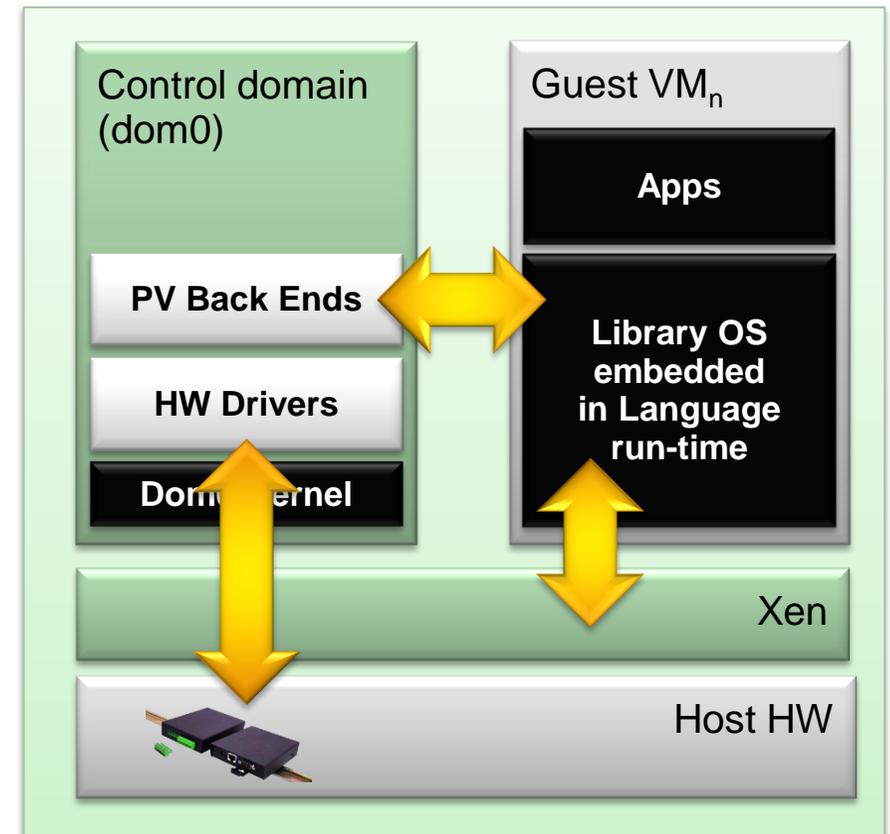
Works on any Xen based cloud or hosting service

Examples

- ErlangOnXen.org : Erlang
- HalVM : Haskell
- OpenMirage : Ocaml

Benefits:

- Small footprint
- Low startup latency
- Extremely fast migration of VMs



Summary: Why Xen?



- Designed for the Cloud : many advantages for cloud use!
 - Resilience, Robustness & Scalability
 - Security: Small surface of attack, Isolation & Advanced Security Features
- Widely used by Cloud Providers and Vendors
- XCP
 - Ready for use with cloud orchestration stacks
- Open Source with a large community and eco-system
 - Xen is still on top of the game
 - Exciting new developments and features in the pipeline





- IRC: [##xen @ FREENODE](#)
- Mailing List: [xen-users](#) & [xen-api](#) ([lists.xen.org](#))
- Wiki: [wiki.xen.org](#)
- Ecosystem pages:
[xen.org/community/ecosystem.html](#)
- Presentations & Videos:
[xen.org/community/presentations.html](#)

Thank You!



@lars_kurth

FREENODE: lars_kurth



Slides available under CC-BY-SA 3.0
From [www.slideshare.net/xen_com_mgr](#)

