

# 10 Years of Xen and beyond ...



 **LINUX FOUNDATION**  
COLLABORATIVE PROJECTS

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

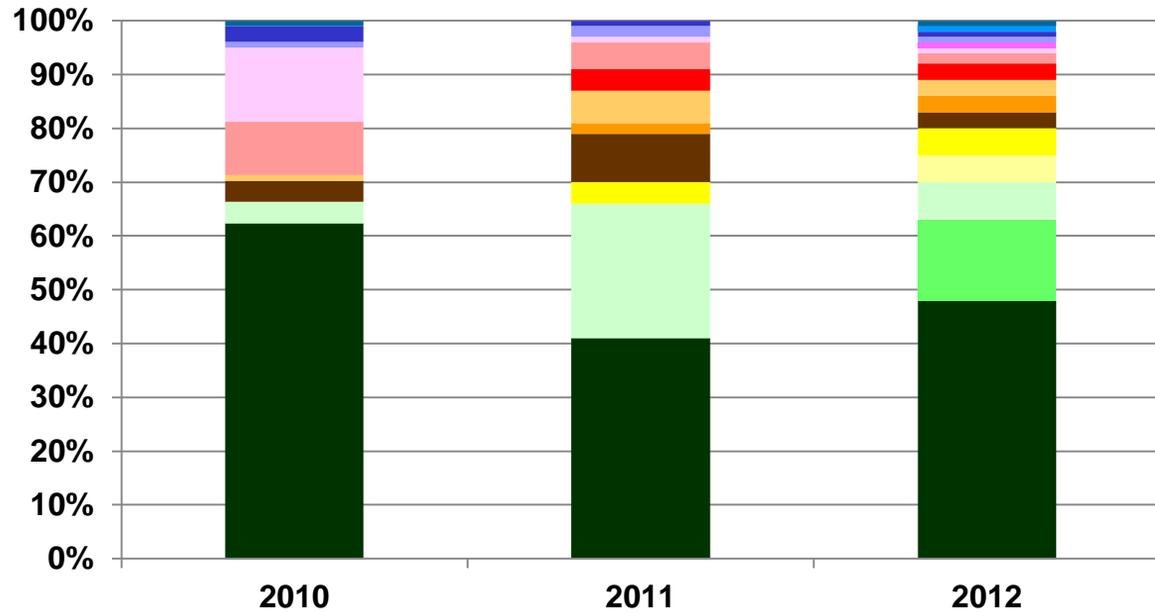
 [@lars\\_kurth](https://twitter.com/lars_kurth)  
FREEMODE: lars\_kurth

# Xen.org becomes XenProject.org

- Teams aka sub-projects
  - Hypervisor
  - XAPI
  - ARM Hypervisor (for Servers as well as Mobile Devices)
  - Mirage OS
- Governance : mixture between Linux Kernel and Apache
  - Consensus decision making
  - Sub-project life-cycle (aka incubator)
  - PMC style structure for team leadership
  - Funded by member companies (Advisory Board)



# Xen contributor community is diversifying



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



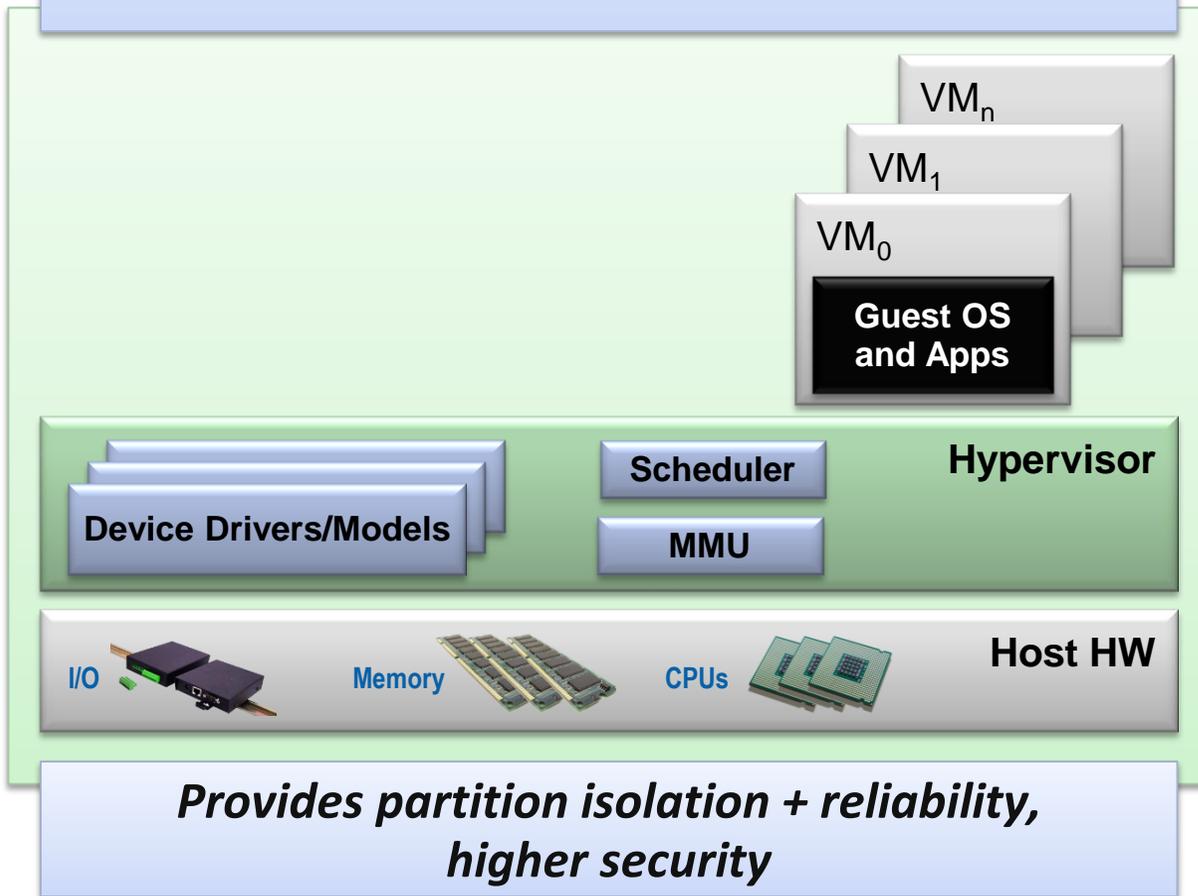
# Hypervisor Architecture



# Hypervisor Architectures

## Type 1: Bare metal Hypervisor

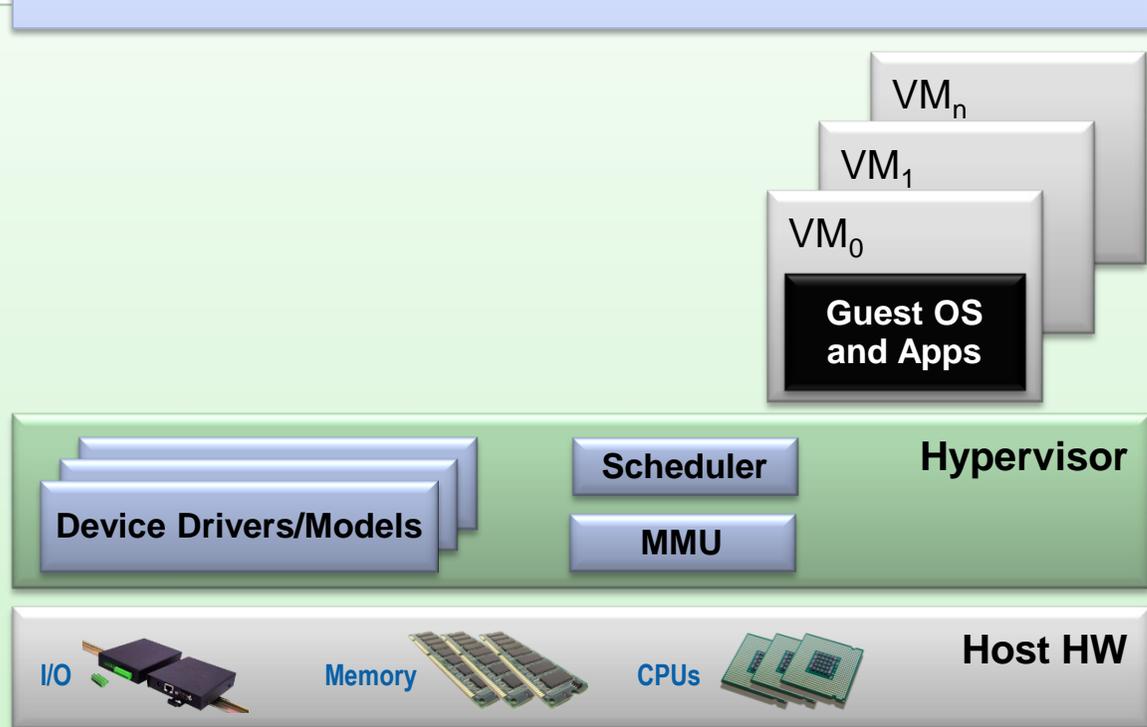
A pure Hypervisor that runs directly on the hardware and hosts Guest OS's.



# Hypervisor Architectures

## Type 1: Bare metal Hypervisor

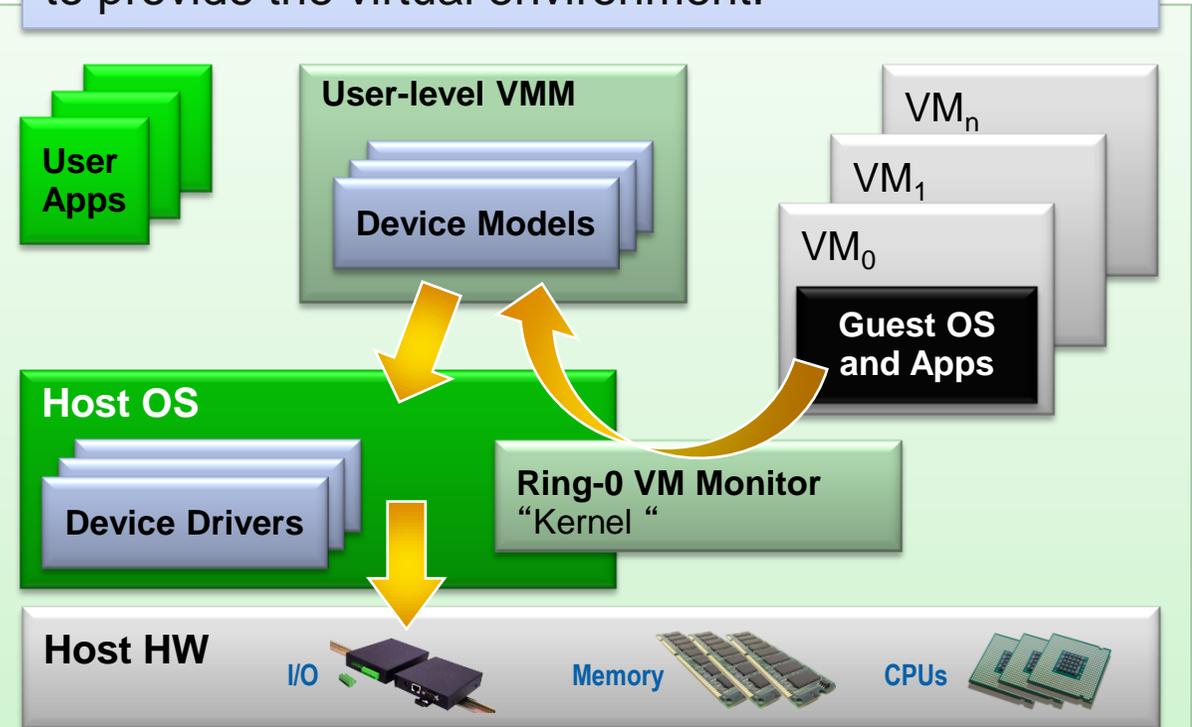
A pure Hypervisor that runs directly on the hardware and hosts Guest OS's.



*Provides partition isolation + reliability,  
higher security*

## Type 2: OS 'Hosted'

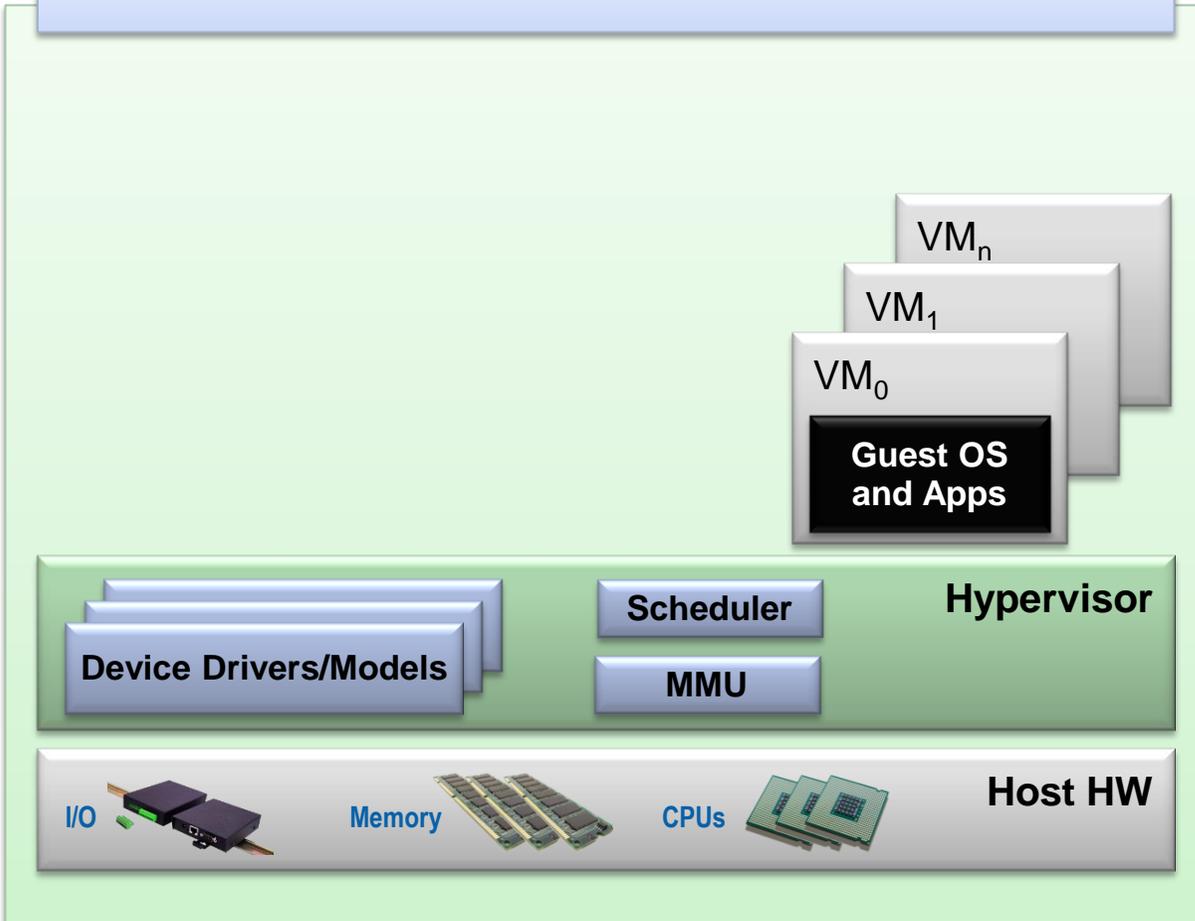
A Hypervisor that runs within a Host OS and hosts Guest OS's inside of it, using the host OS services to provide the virtual environment.



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

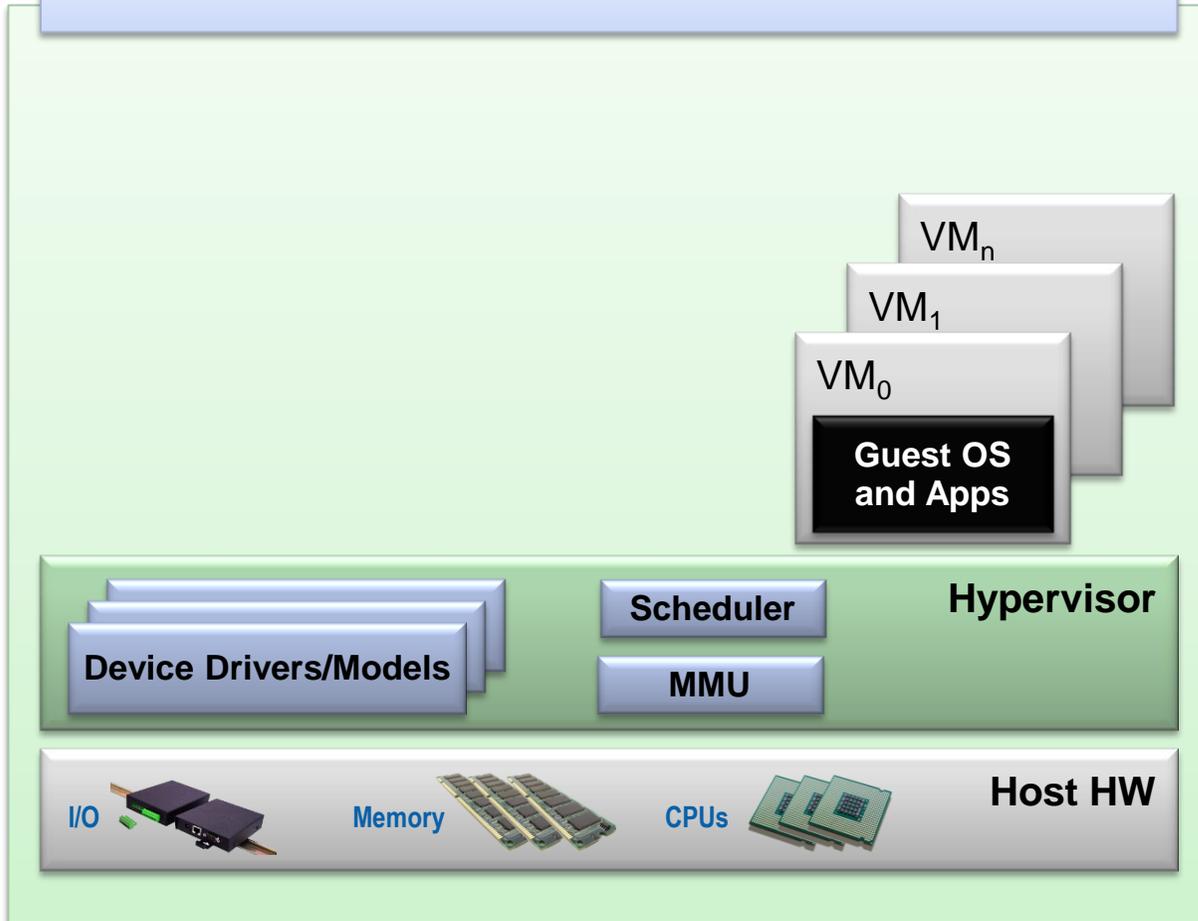
# Xen: Type 1 with a Twist

## Type 1: Bare metal Hypervisor

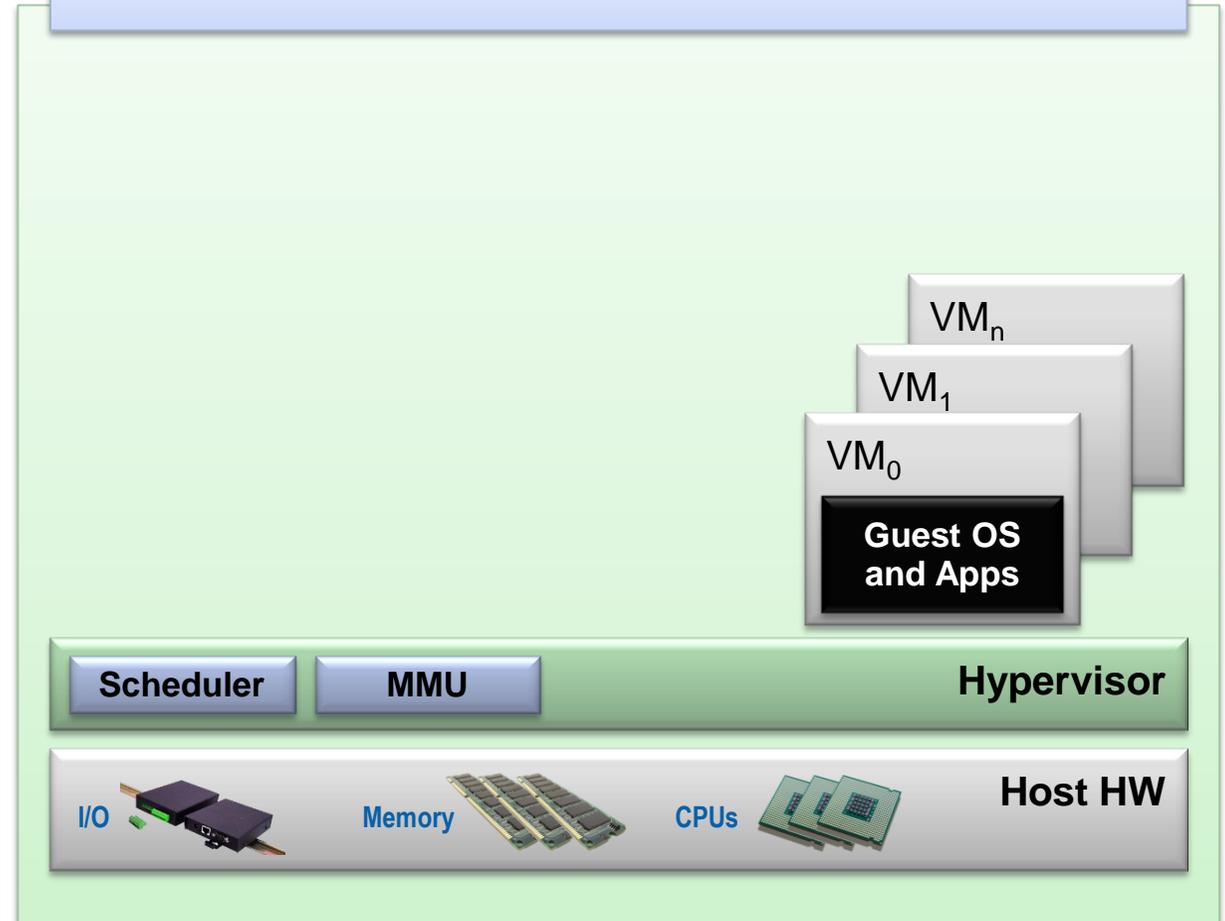


# Xen: Type 1 with a Twist

## Type 1: Bare metal Hypervisor

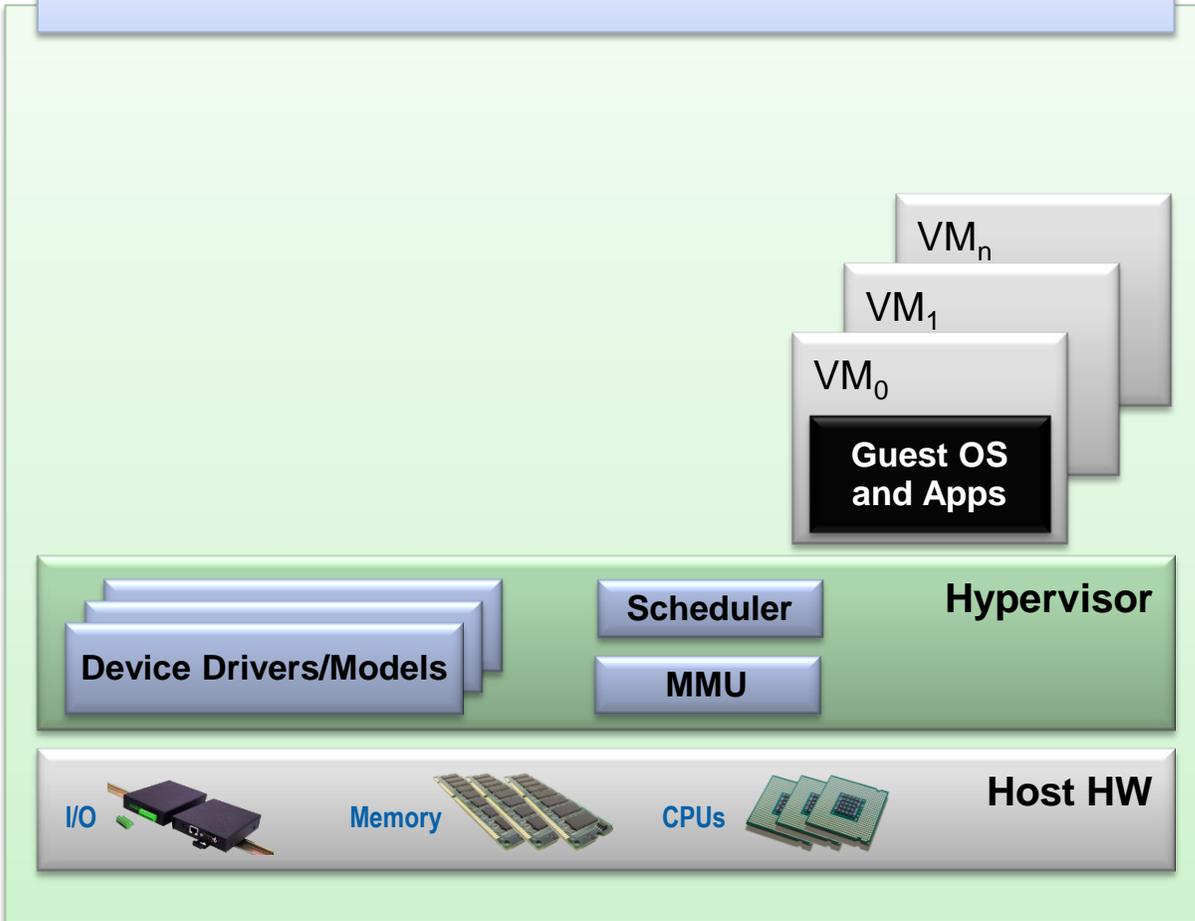


## Xen Architecture

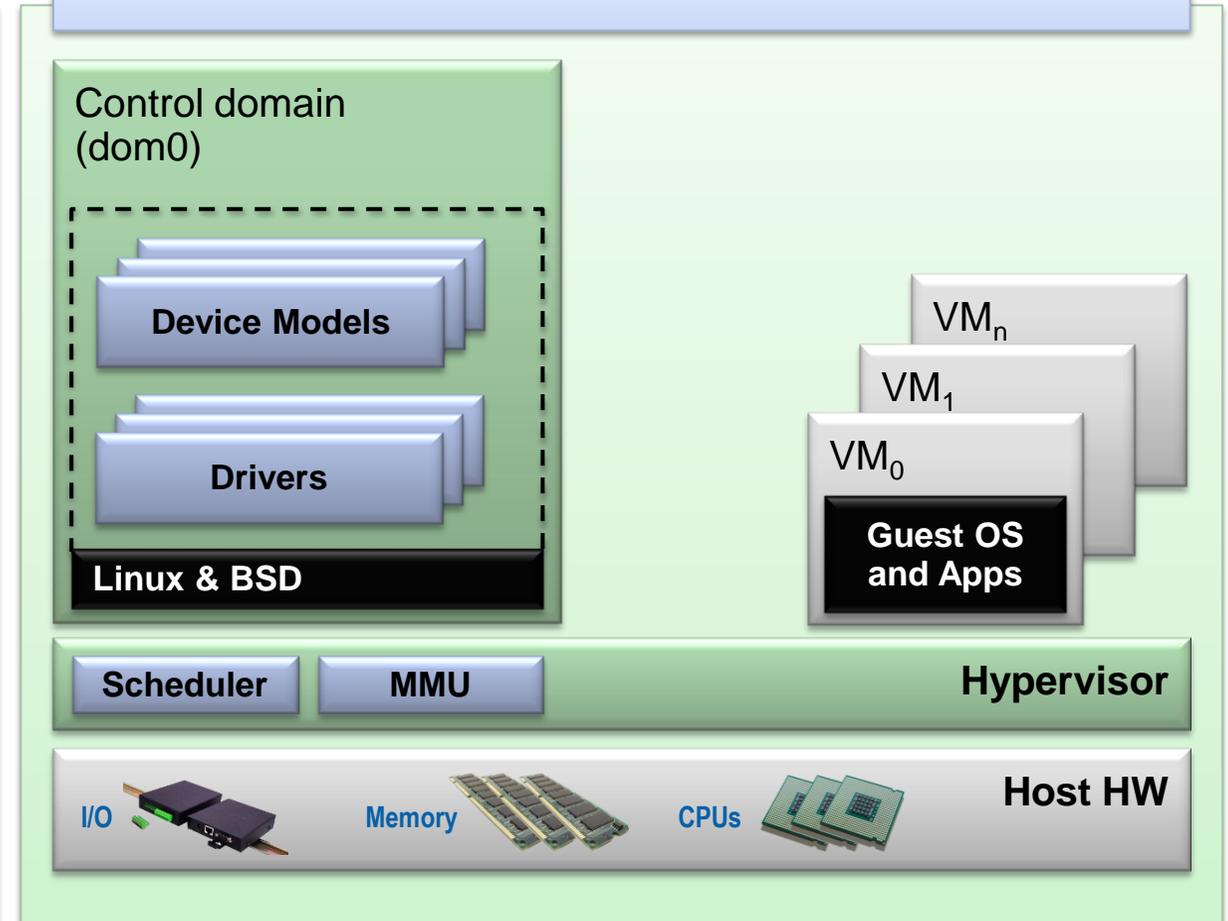


# Xen: Type 1 with a Twist

## Type 1: Bare metal Hypervisor



## Xen Architecture



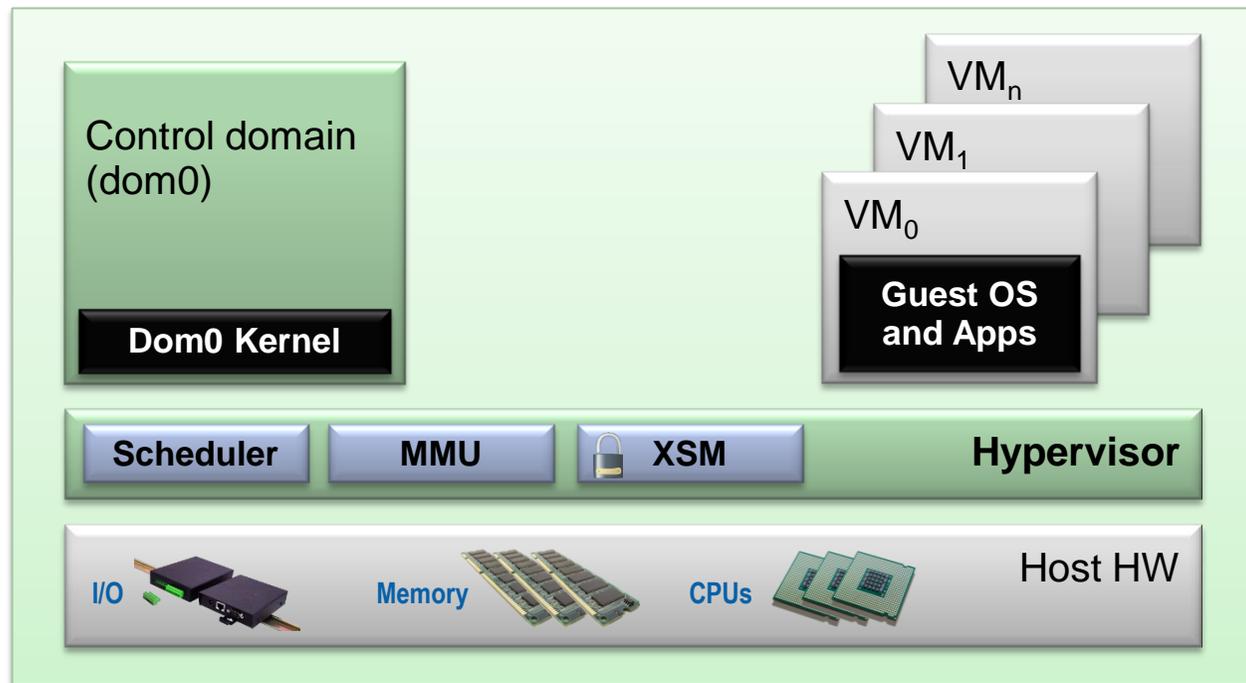
# Xen Project and Linux

- 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



## Control Domain aka **Dom0**

- Dom0 kernel with drivers

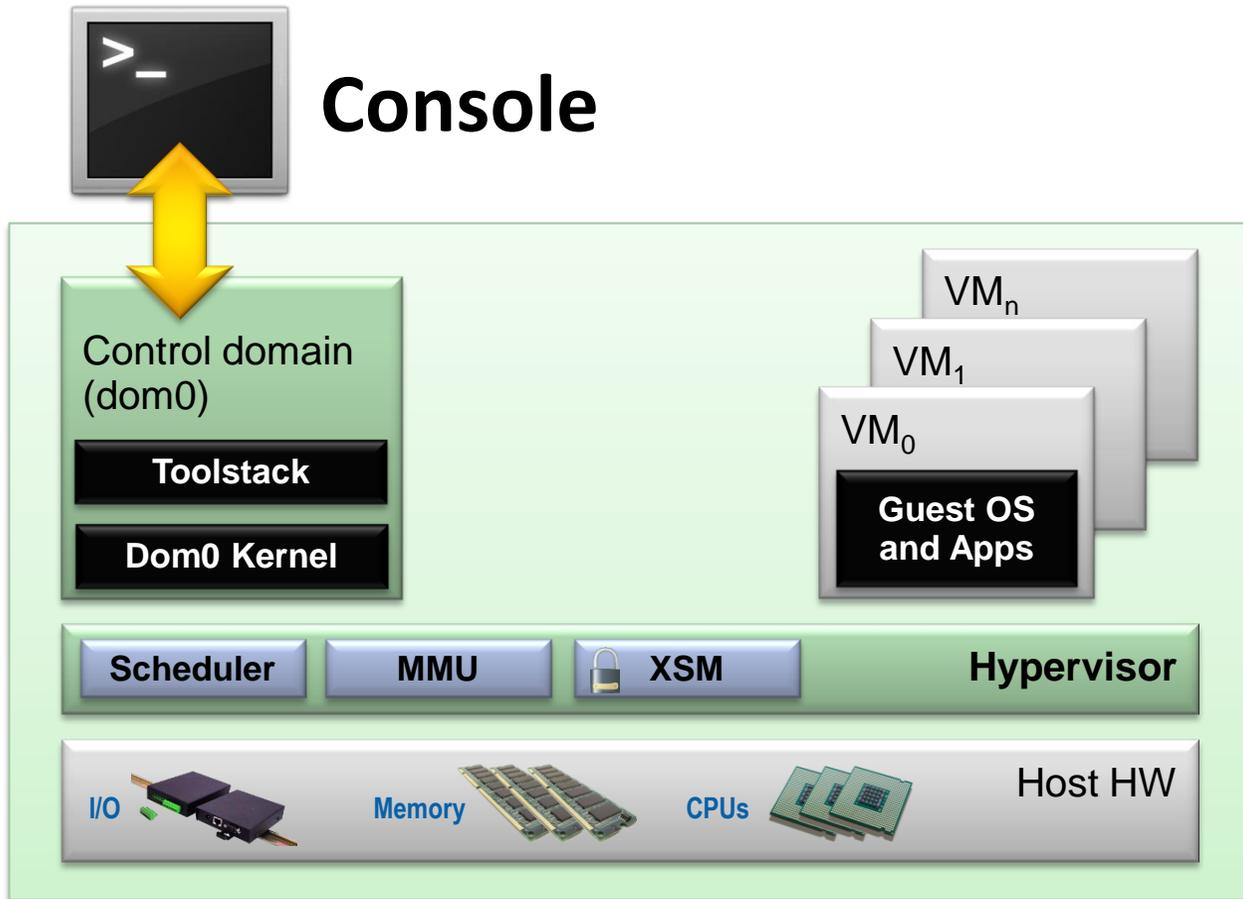
## Guest Domains

- Your apps

■ Trusted Computing Base



# Basic Xen Concepts



■ Trusted Computing Base

## Console

- Interface to the outside world

## Control Domain aka **Dom0**

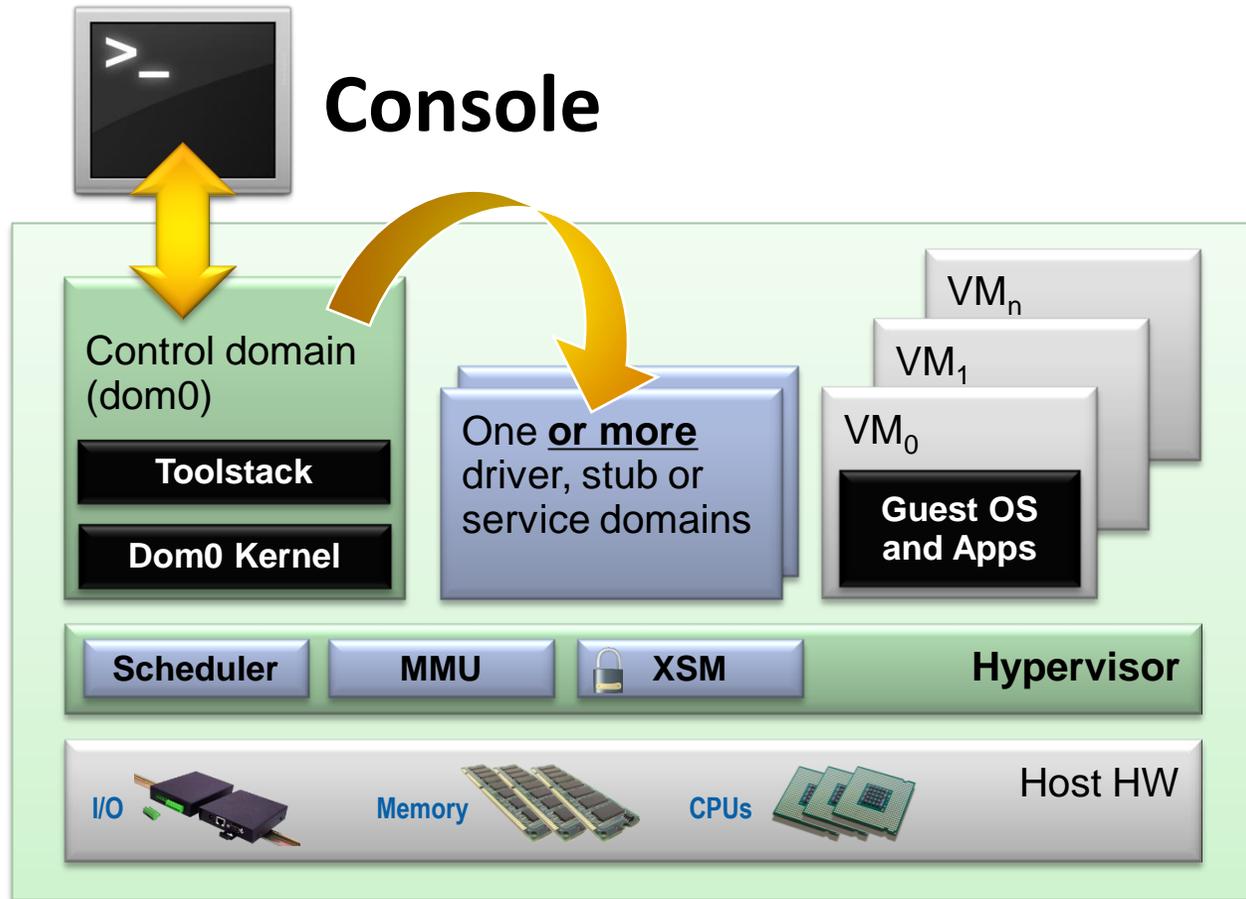
- Dom0 kernel with drivers
- Xen Management Toolstack

## 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

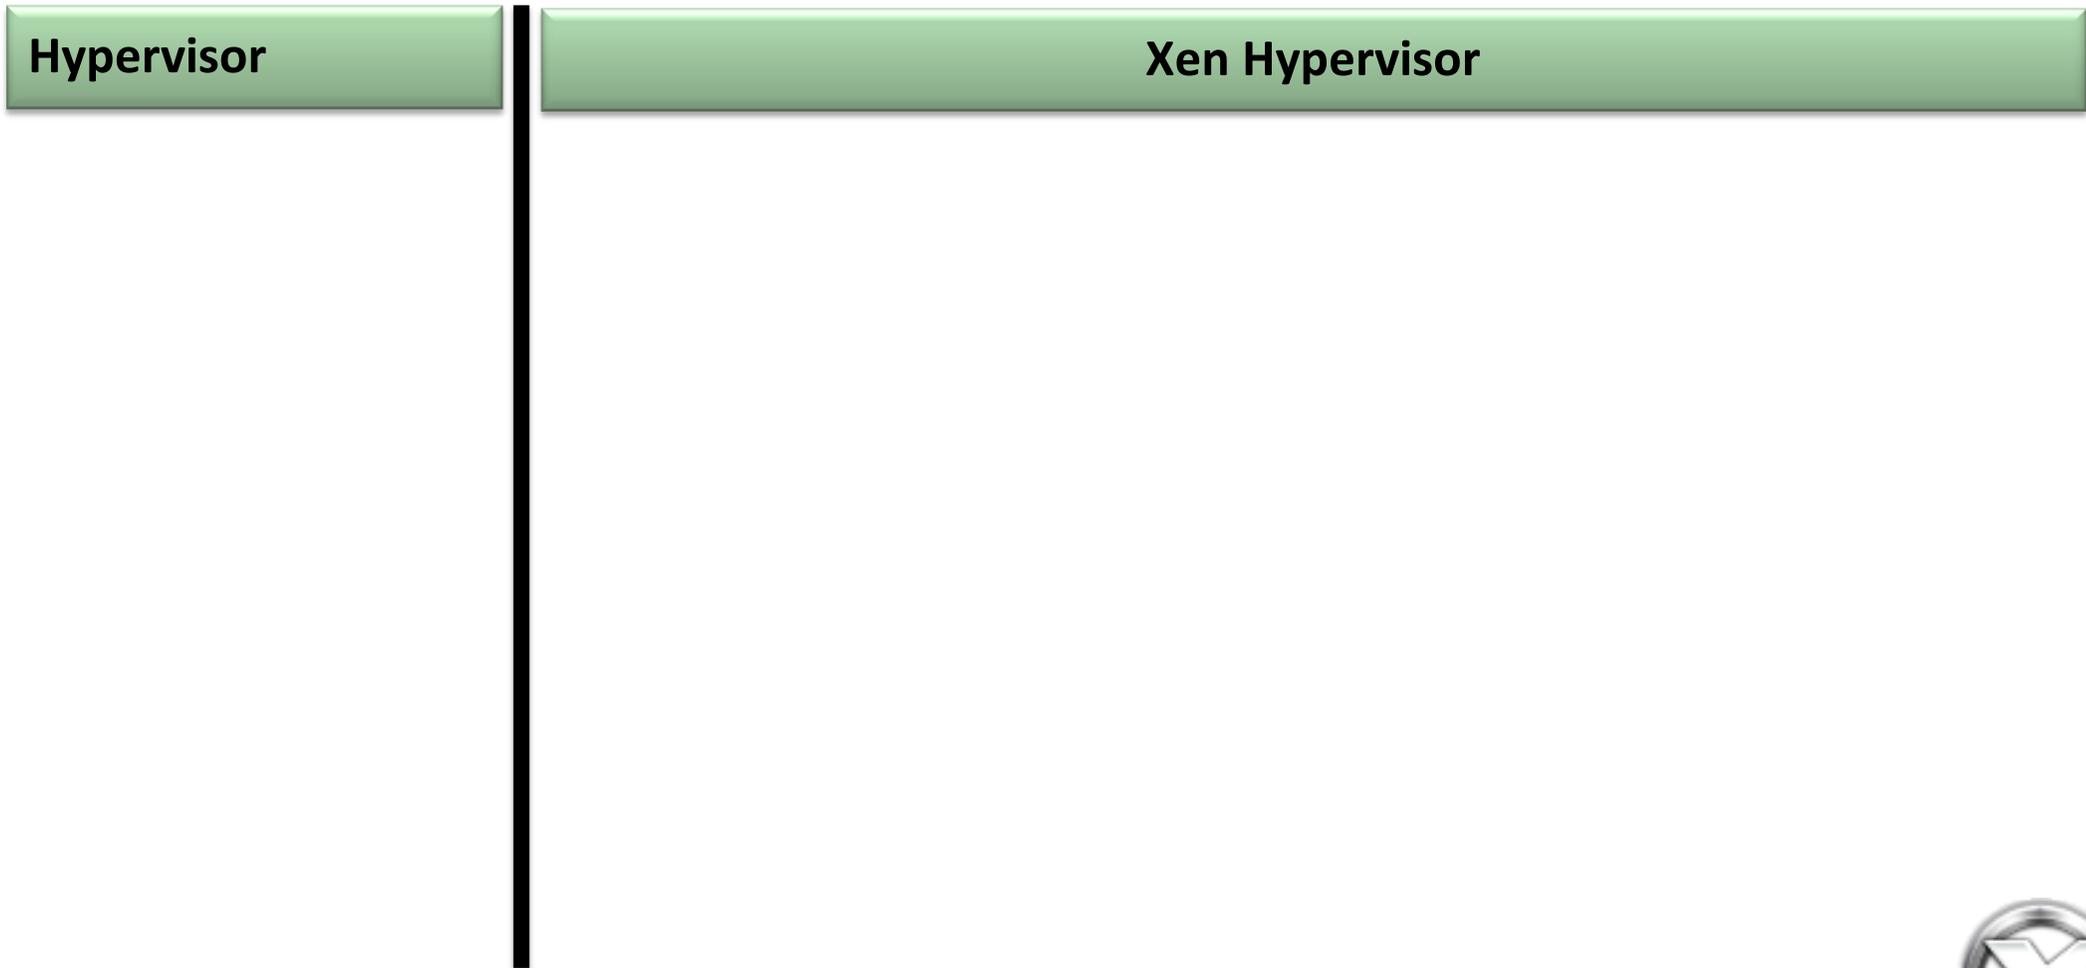
- Your apps

## Driver/Stub/Service Domain(s)

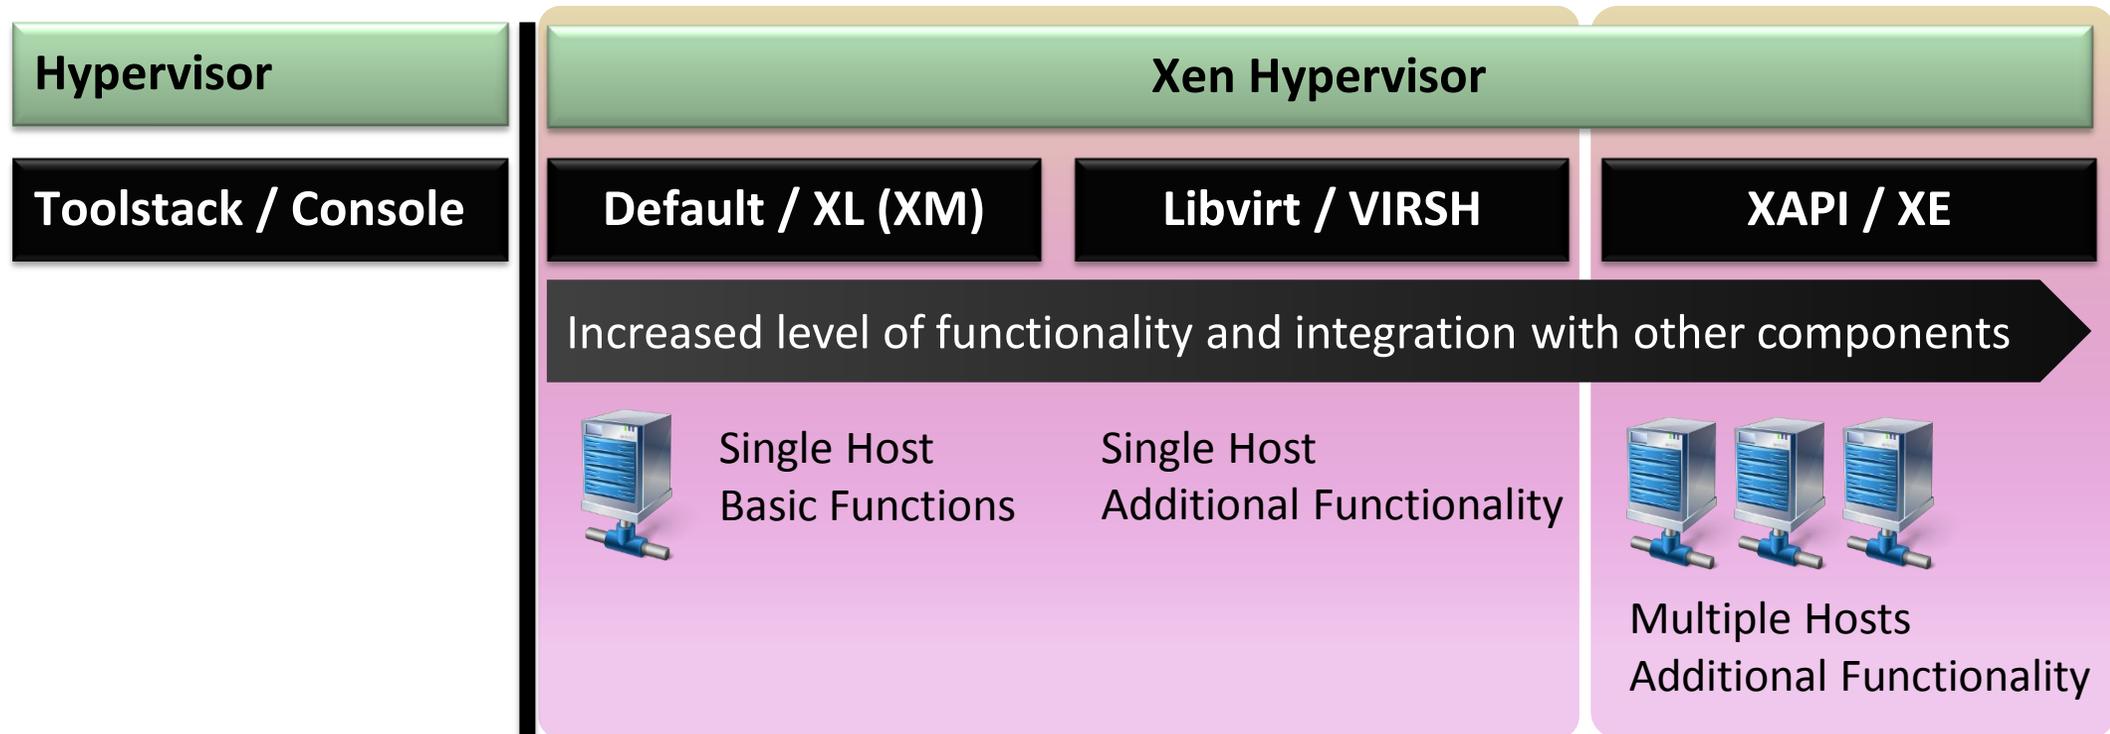
- A “driver, device model or control service in a box”
- De-privileged and isolated
- Lifetime: start, stop, kill



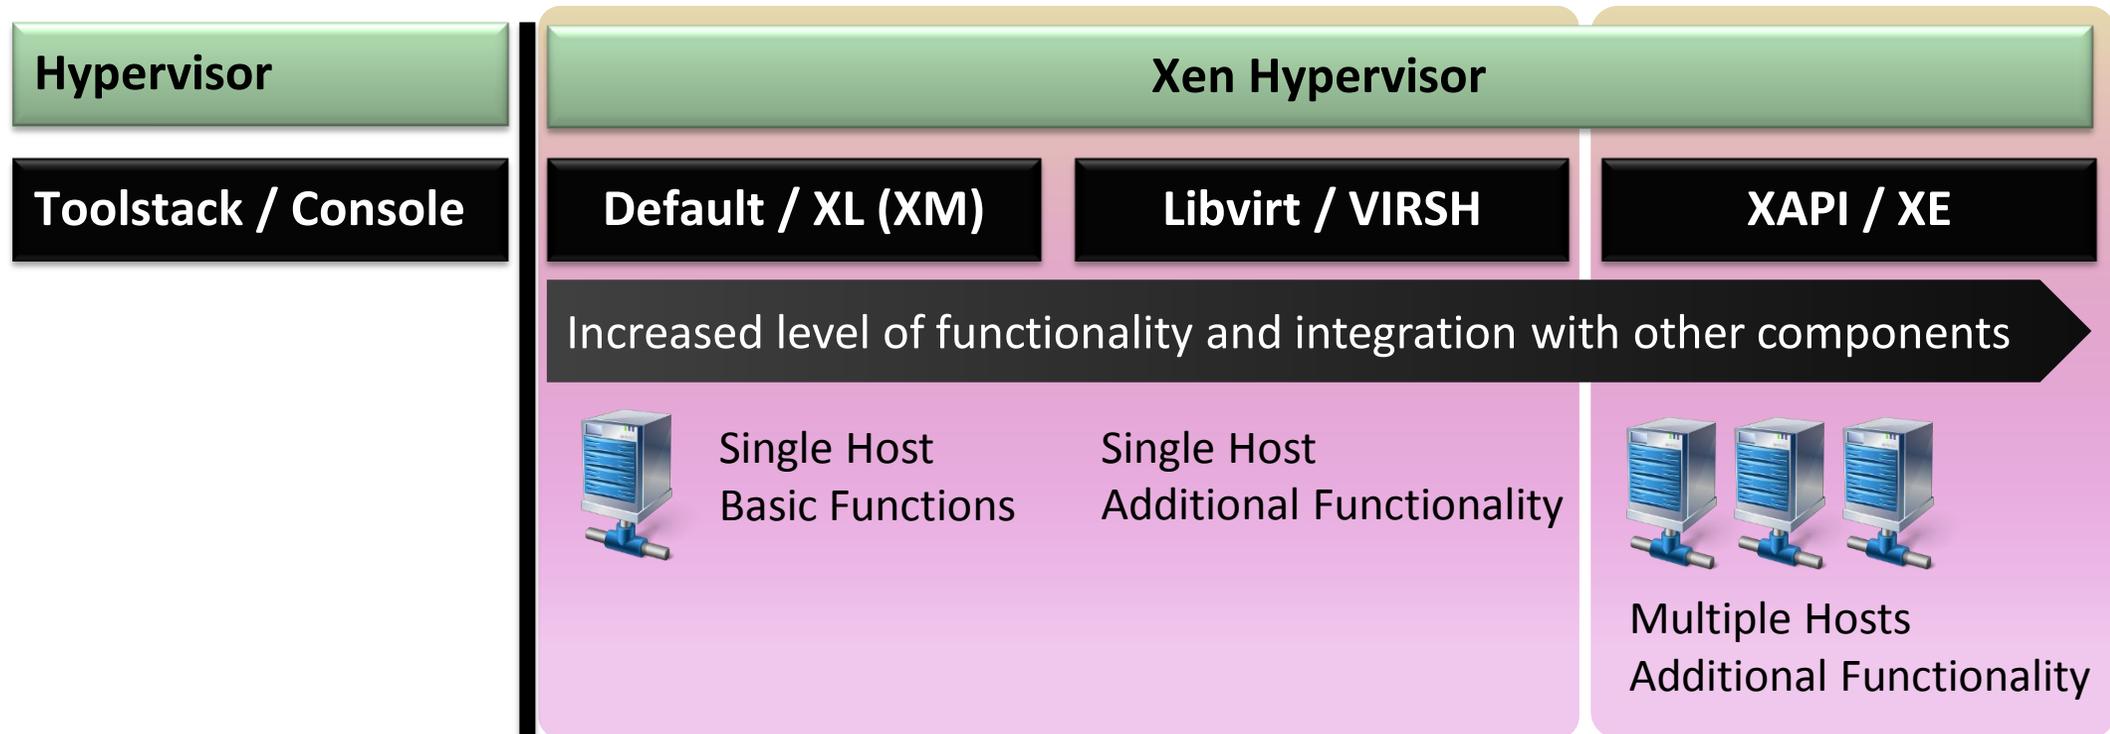
# Xen Variants for Server & Cloud



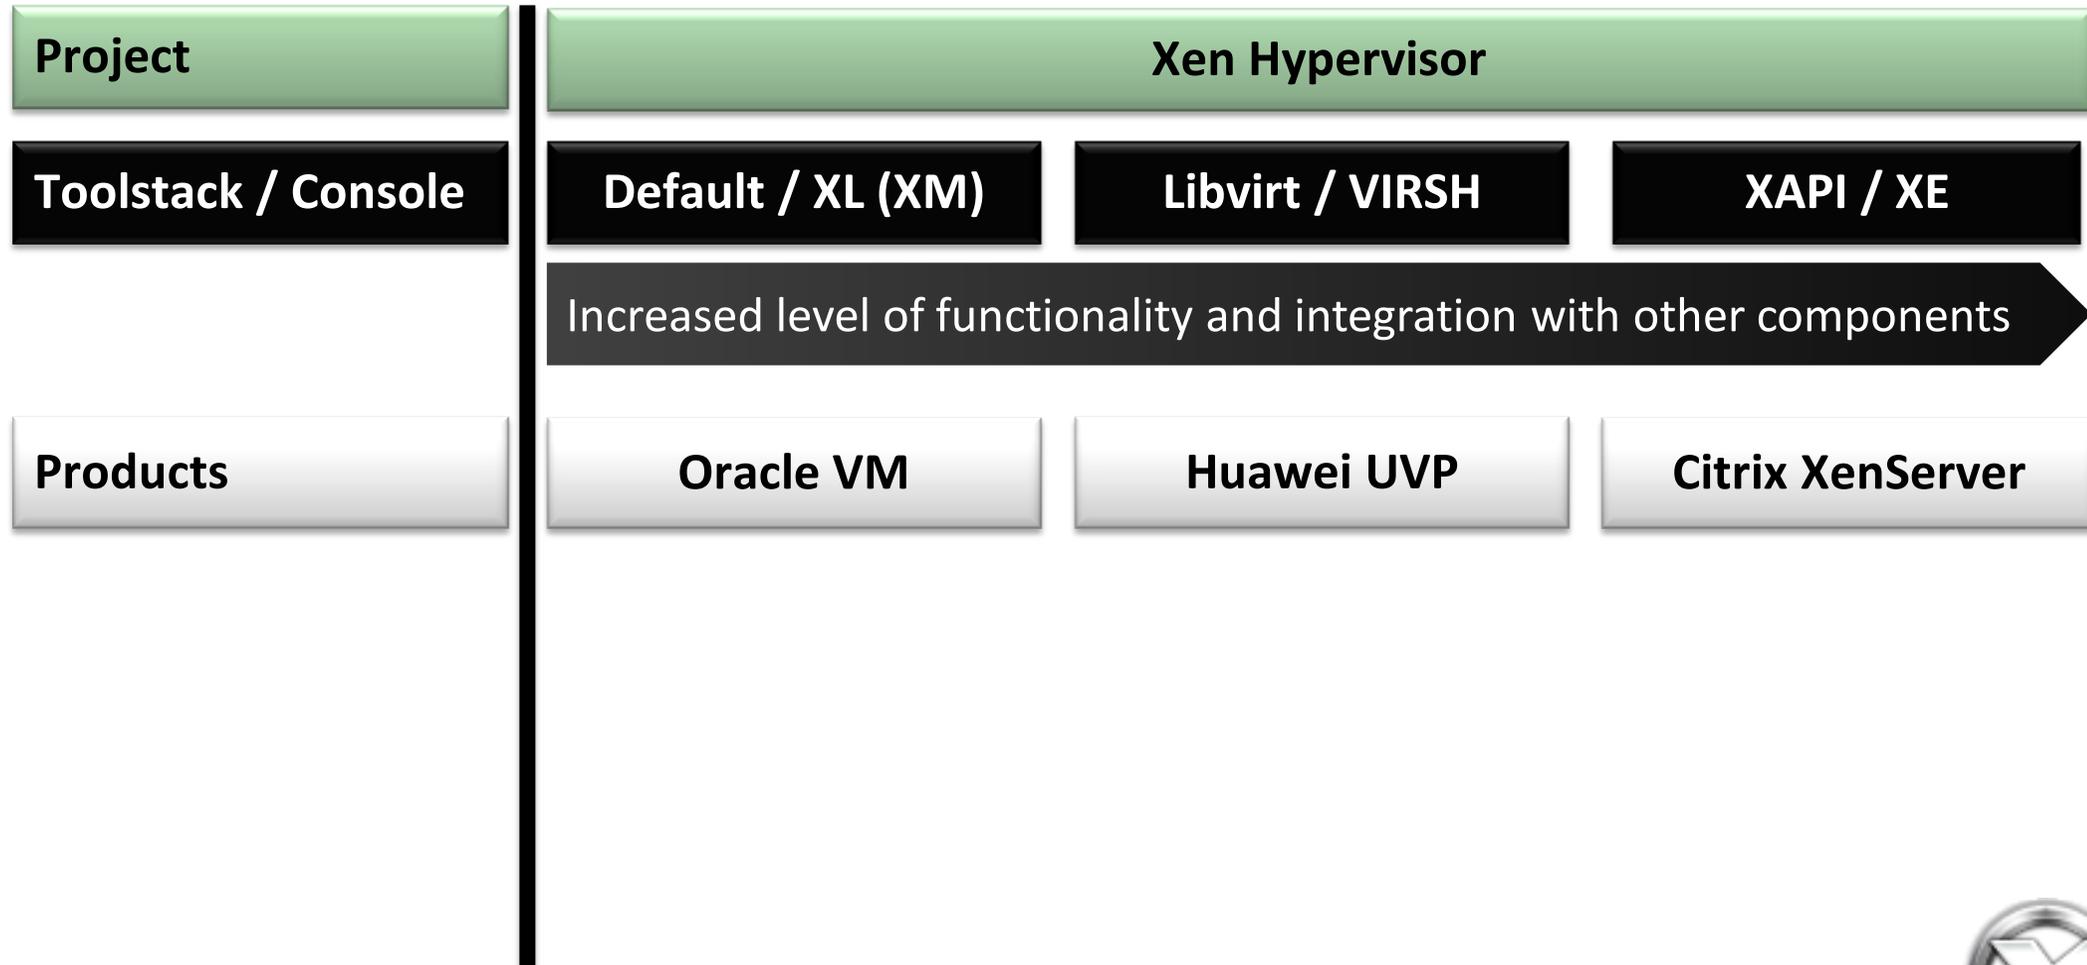
# Xen Variants for Server & Cloud



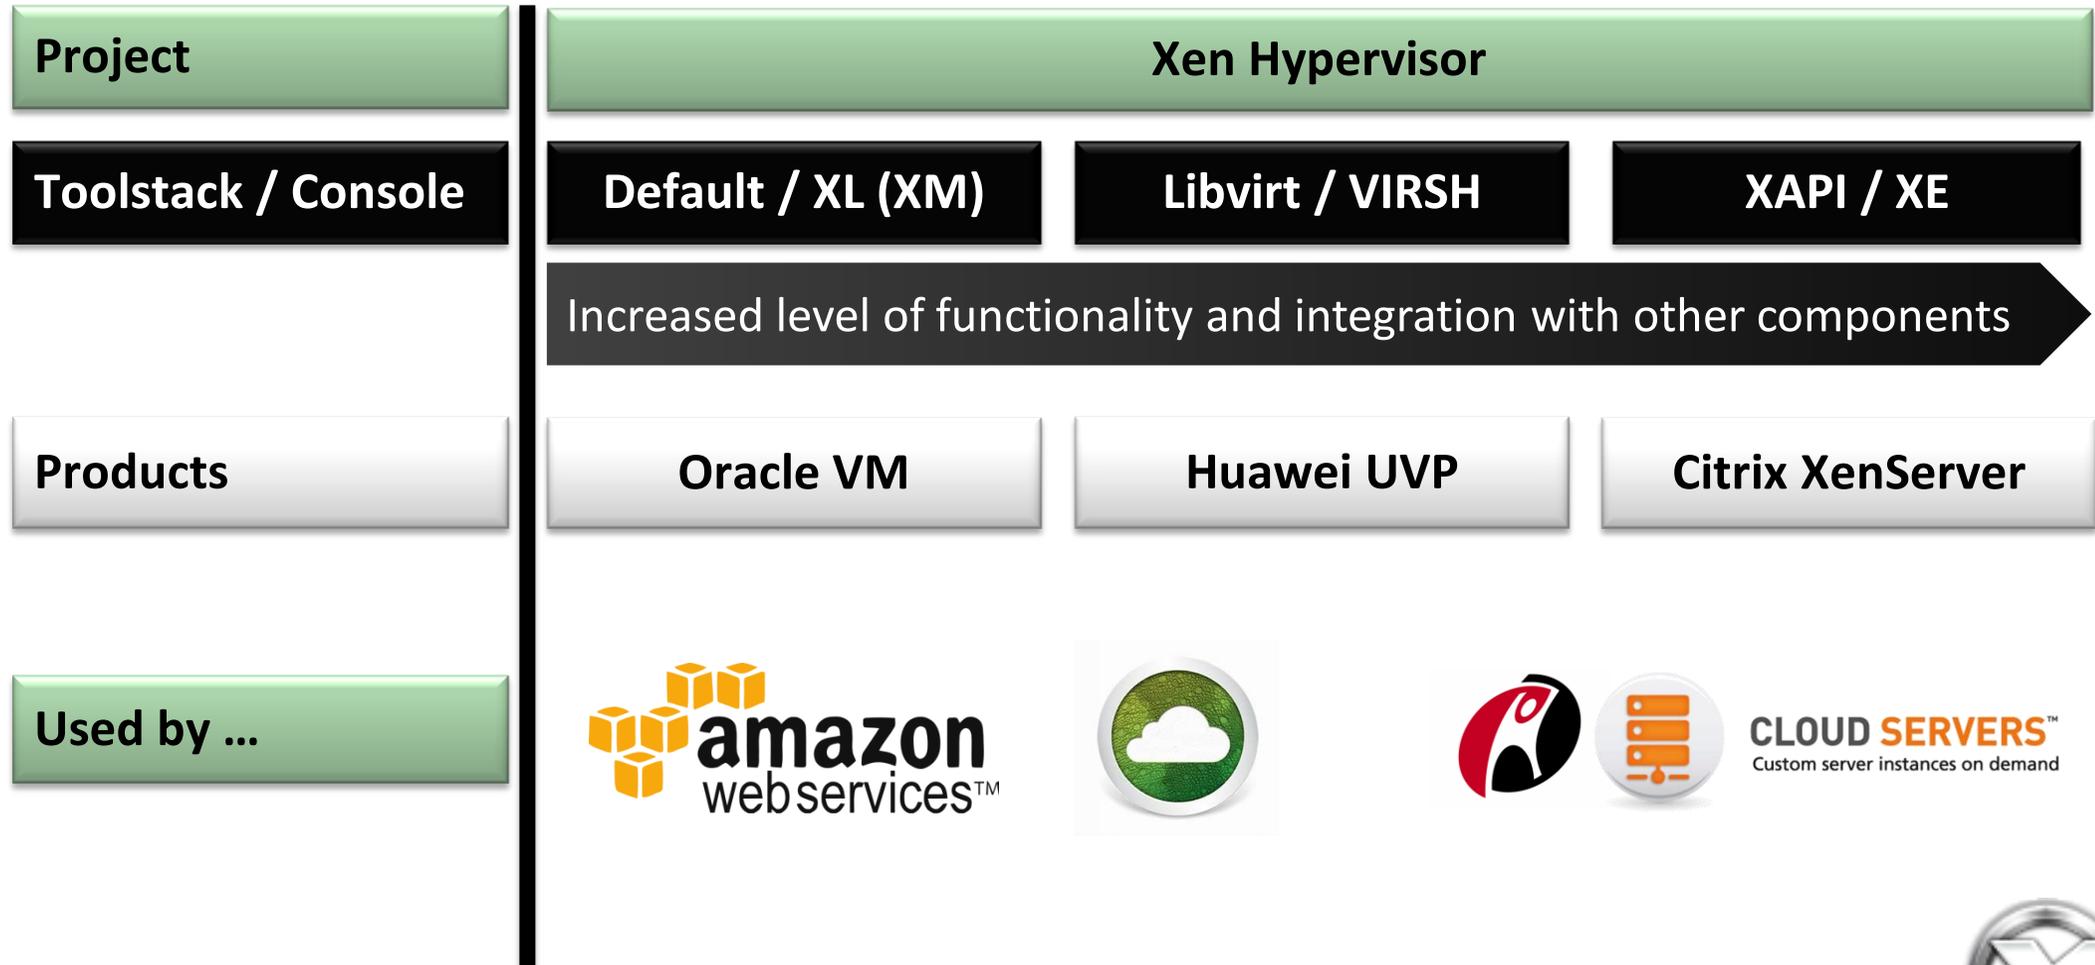
# Xen Variants for Server & Cloud



# Xen Variants for Server & Cloud



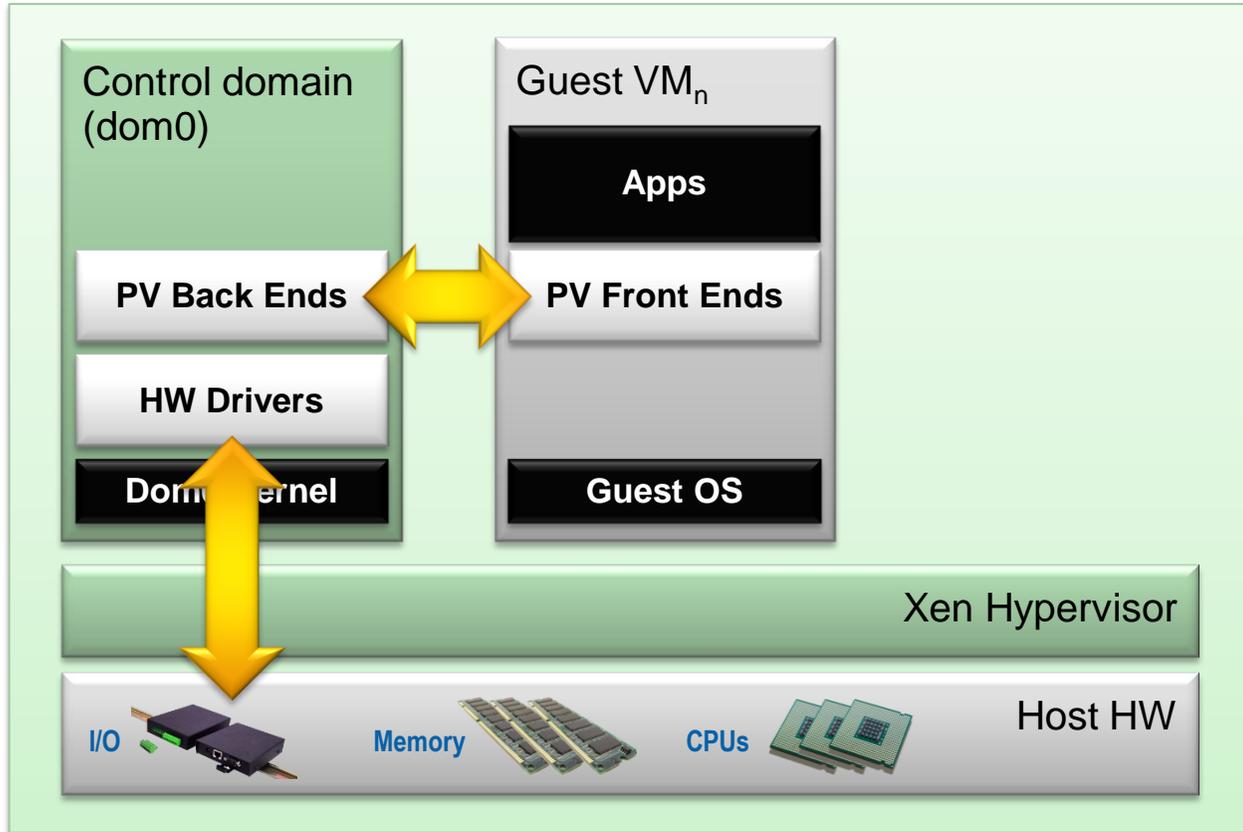
# Xen Variants for Server & Cloud



# Xen : Types of Virtualization



# PV Domains



## Technology:

- Paravirtualization

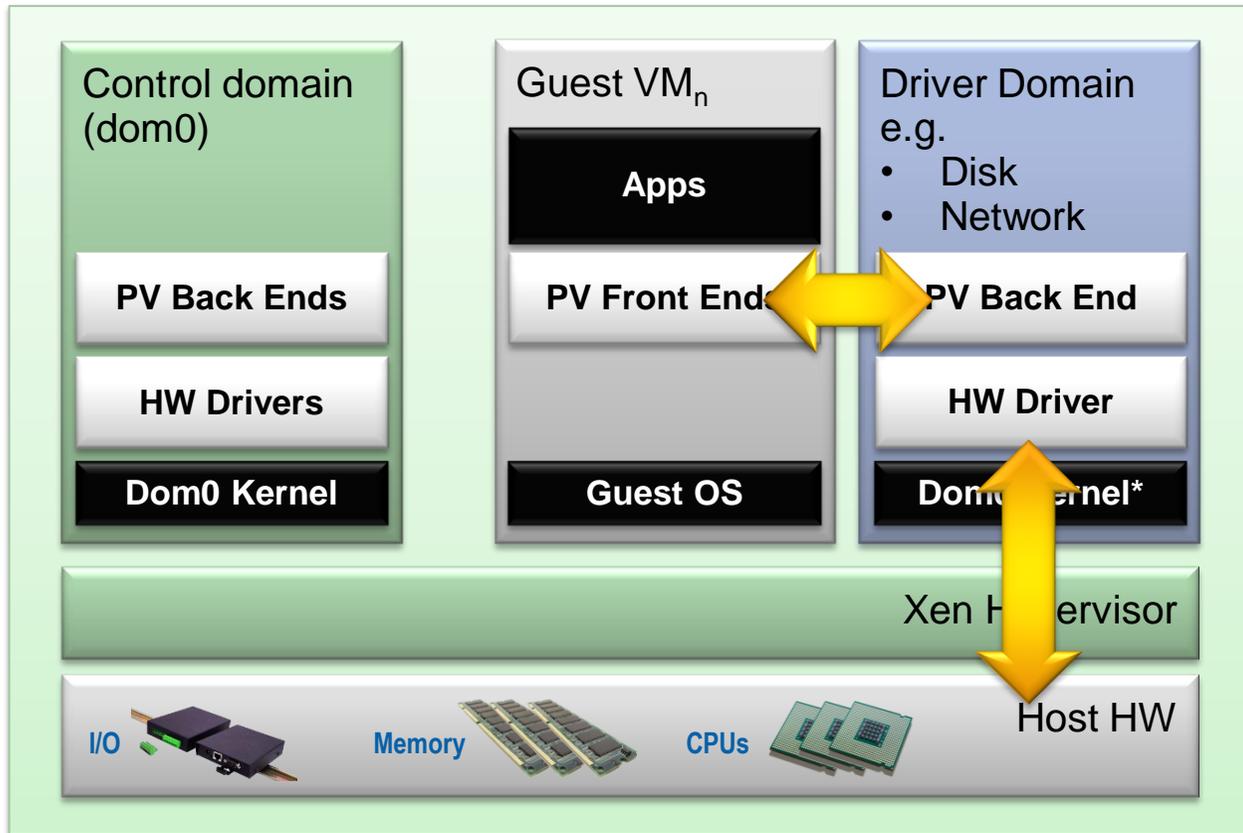
## Linux PV guests have limitations:

- limited to a subset of set of virtual HW

## 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 to a subset of virtual HW

## Advantages

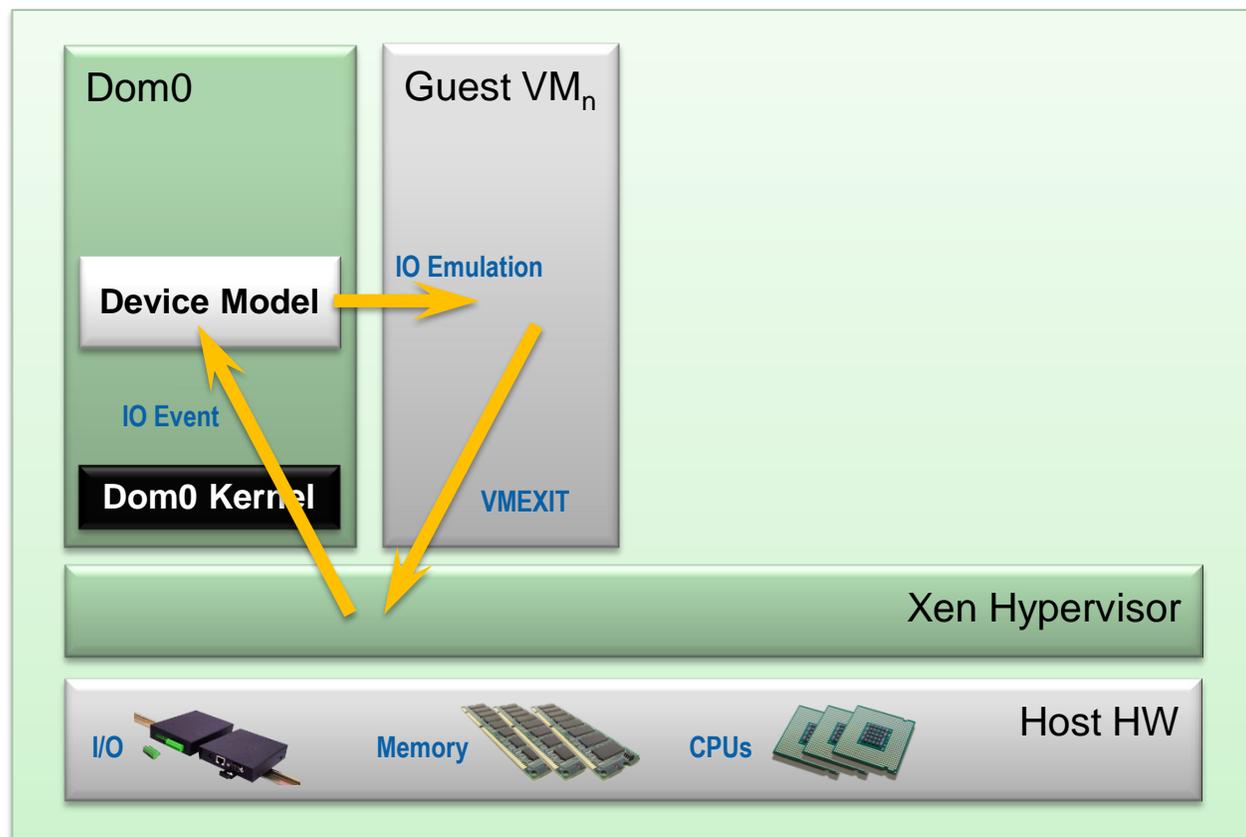
- Fast
- Works on any system  
(even without virt extensions)

## 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

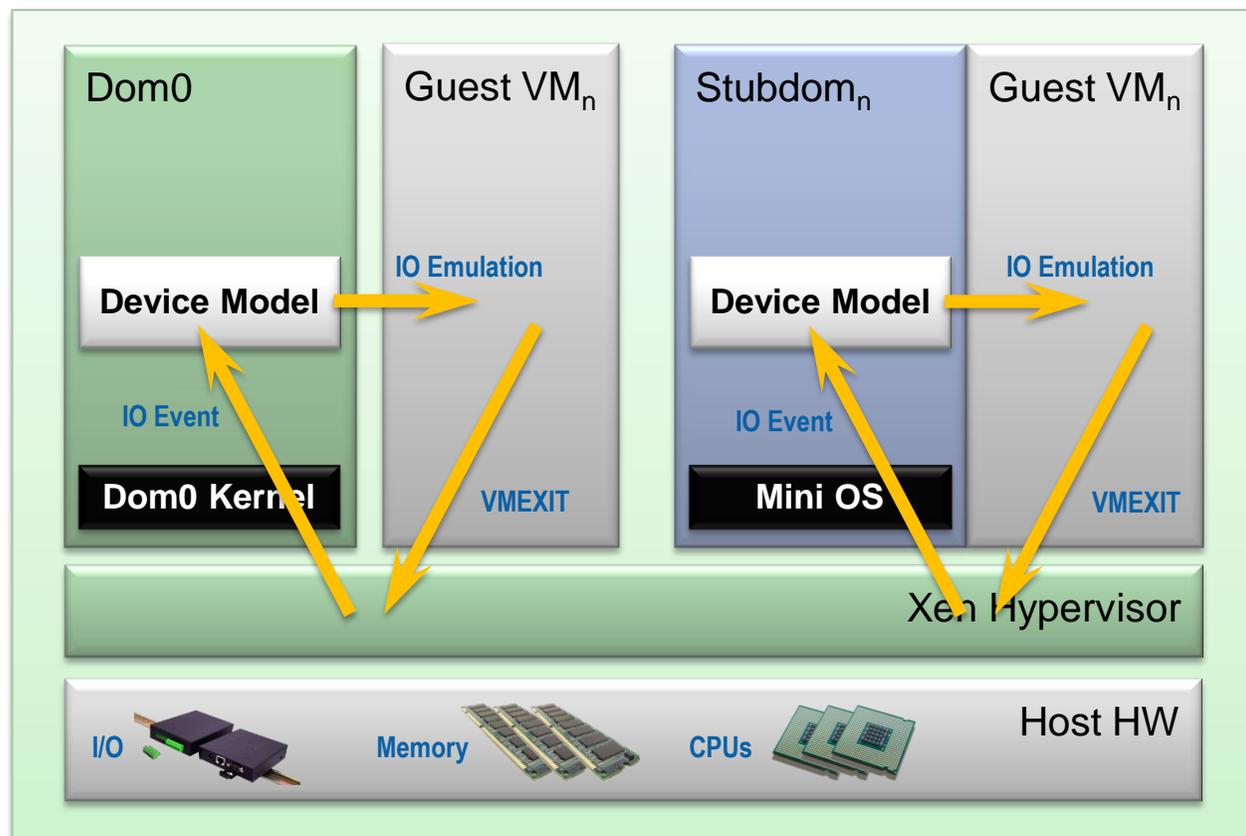
## Disadvantages

- 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

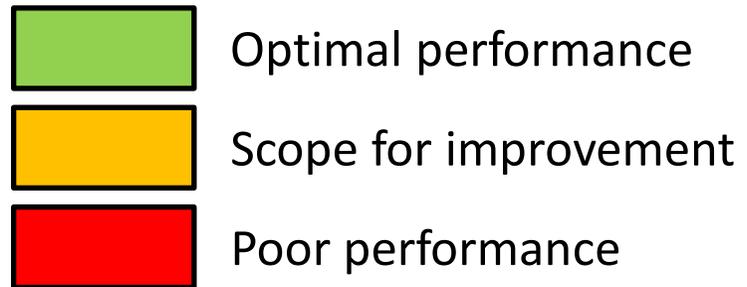


# The Virtualization Spectrum

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

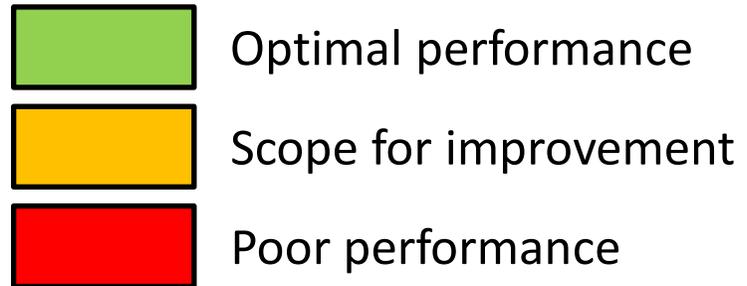
	Disk and Network	Interrupts, Timers	Emulated Motherboard, Legacy boot	Privileged Instructions and page tables	
Fully Virtualized (FV)	VS	VS	VS	VH	} HVM mode/domain
FV with PV for disk & network	P	VS	VS	VH	
PVHVM	P	P	VS	VH	
PVH <span style="color: blue; font-weight: bold;">NEW Xen 4.4</span>	P	P	P	VH	} 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 <span style="float: right; font-size: small;">NEW Xen 4.4</span>	P	P	P	VH	} PV mode/domain
Fully Paravirtualized (PV)	P	P	P	P	

# The Virtualization Spectrum

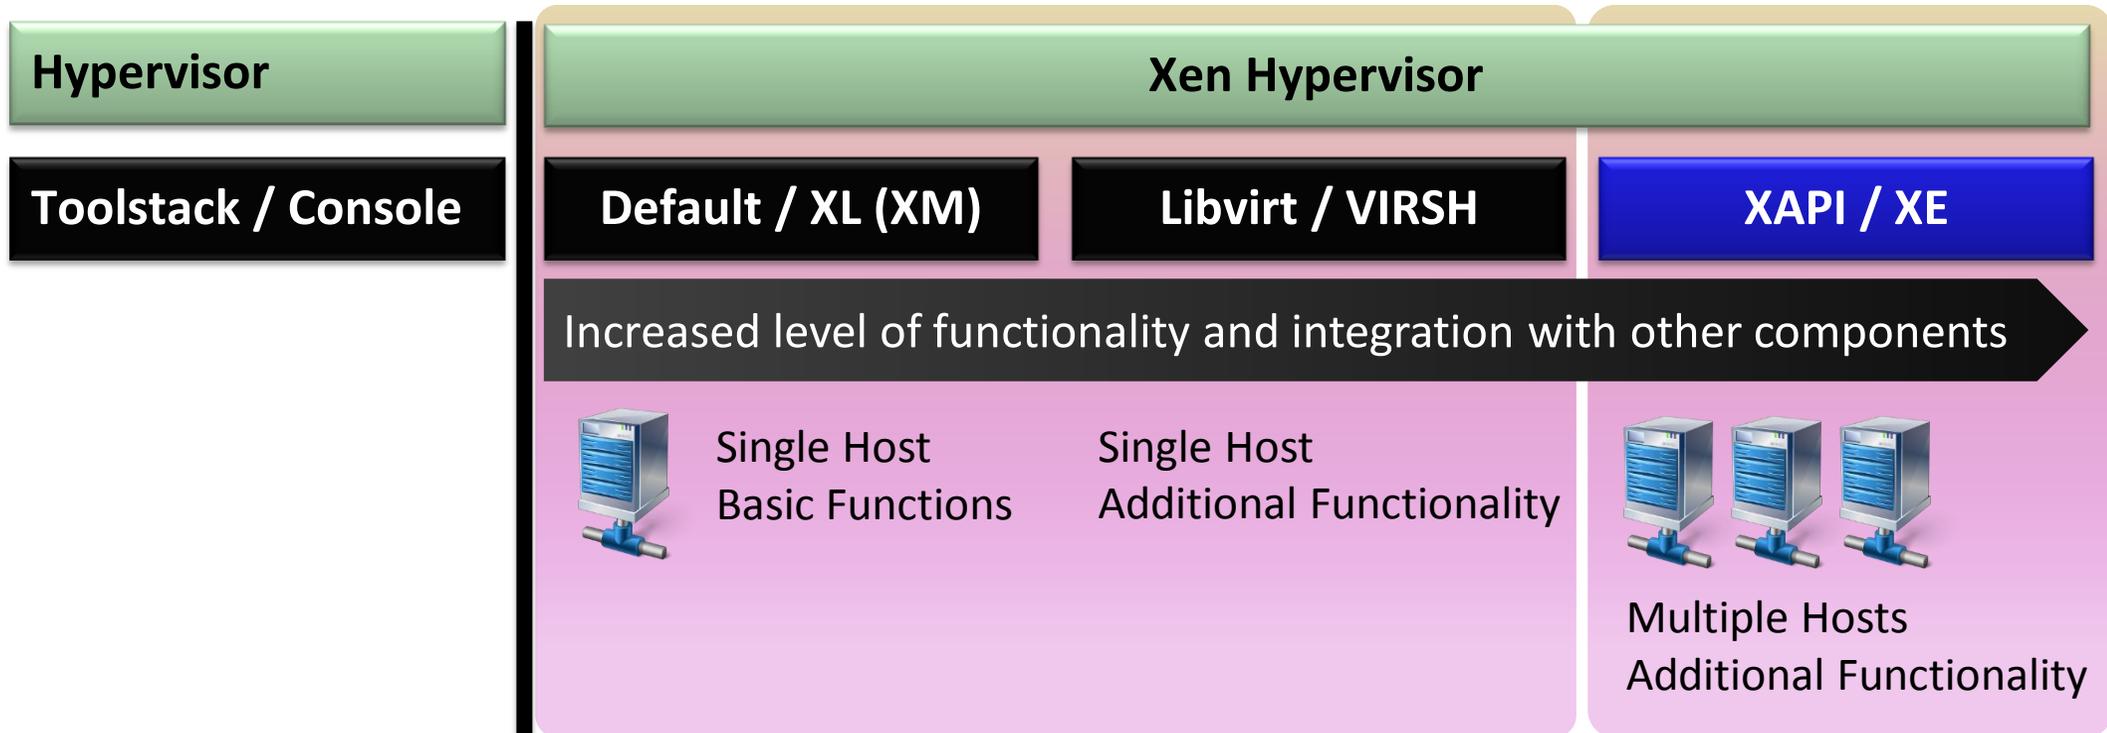


**Important:** Xen automatically picks the best option based on HW & OS capabilities and available drivers.  
 As a Xen user I chose a HVM or PV domain.

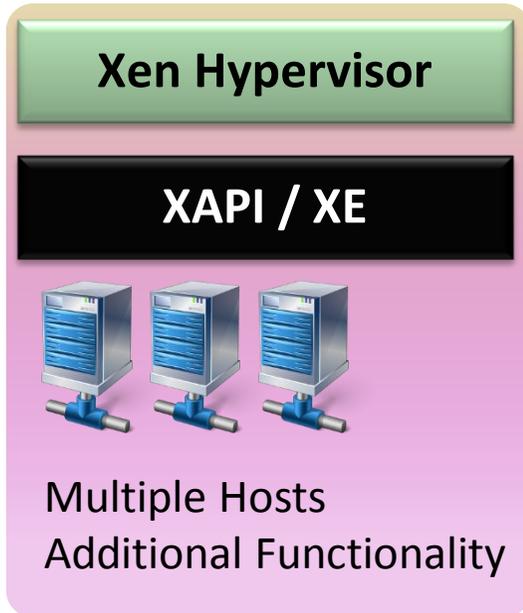
	Disk and Network	Interrupts, Tim	Emulated Mo	Legacy boot	Privileged Ins and page tar	
Fully Virtualized (FV)	VS	VS	VS	VH	} HVM mode/domain	
FV with PV for disk & network	P	VS	VS	VH		
PVHVM	P	P	VS	VH		
PVH <span style="float: right; font-size: small;">NEW Xen 4.4</span>	P	P	P	VH	} PV mode/domain	
Fully Paravirtualized (PV)	P	P	P	P		



# XAPI, XCP and XCP-XAPI : What is it?



# XAPI : What do I get?



- VM lifecycle: live snapshots, checkpoint, migration
- Storage XenMotion: Migrate VMs between hosts or pools without shared storage (while the VM is running)
- Resource pools: flexible storage and networking
- Event tracking: progress, notification
- Upgrade and patching capabilities
- Real-time performance monitoring and alerting
- 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](http://wiki.xen.org/wiki/XCP_Release_Features)

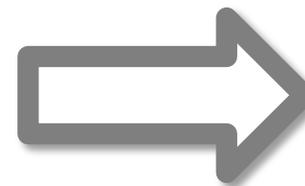
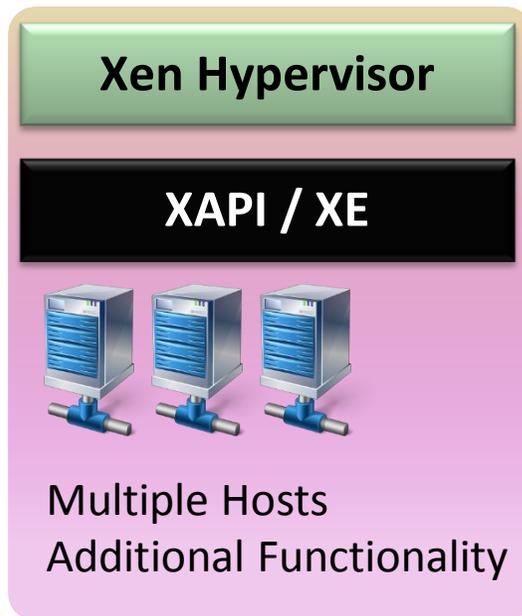
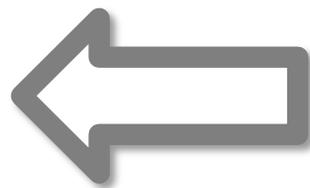


# XAPI : two variants!



## XCP ISO (at v1.6)

Xen 4.1.3 + XAPI  
CentOS 5.3  
Kernel (v2.6.32.43)  
OVS 1.4.2



## XCP-XAPI packages

Debian Wheezy  
Ubuntu 12.04 LTS  
Others in progress ...

# XAPI : Orchestration and UIs

**Xen Hypervisor**

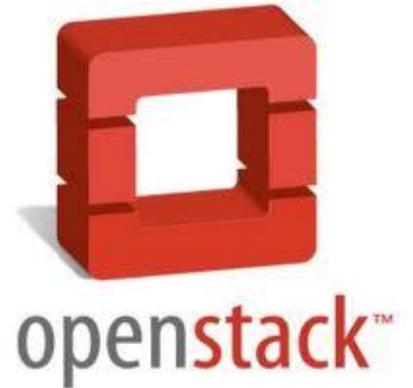
**XAPI / XE**



Multiple Hosts  
Additional Functionality

**apache cloudstack**  
open source cloud computing

**OpenNebula.org**



# Other XEN Cloud Orchestration

OpenNebula.org  
EUCALYPTUS

Xen Hypervisor

XM or XL



Single Host  
Basic Functions

Xen Hypervisor

Libvirt



Single Host  
Additional Functionality



# Challenges for FOSS hypervisors



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

[www.colt.net/cio-research](http://www.colt.net/cio-research)



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

Results XCP User Survey 2013 – 90% of users quoted these as most important attributes



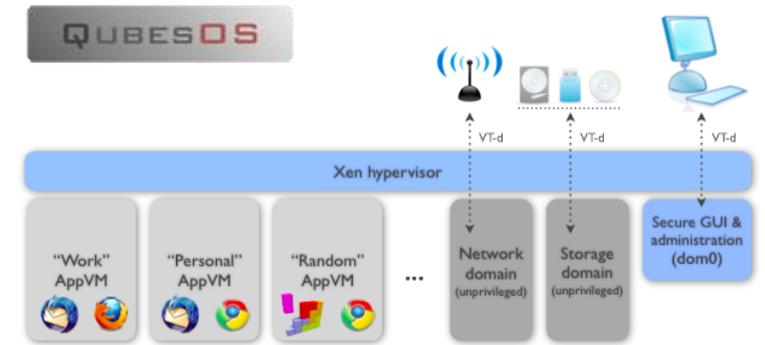
# Disaggregation

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](#)"

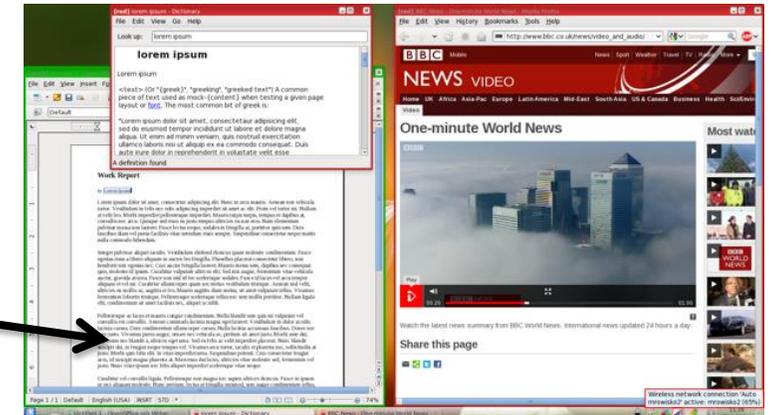
Used today by [Qubes OS](#) and Citrix XenClient XT

Prototypes for XAPI



See [qubes-os.org](http://qubes-os.org)

Different windows run in different VMs



# Benefits of Disaggregation

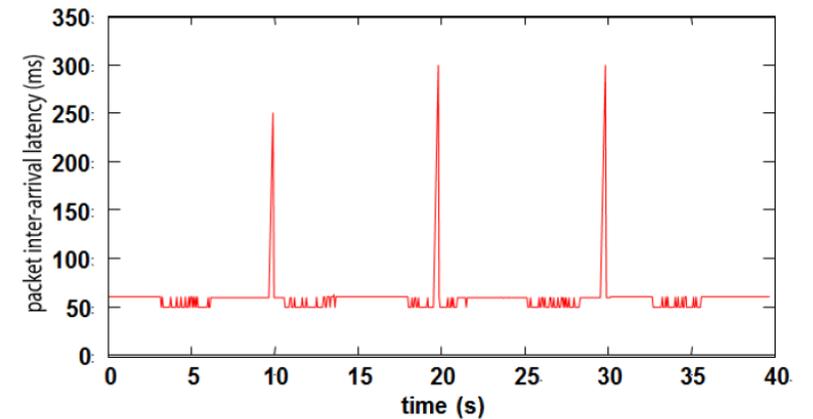
More Security

Increased serviceability and flexibility

Better Robustness

Better Performance

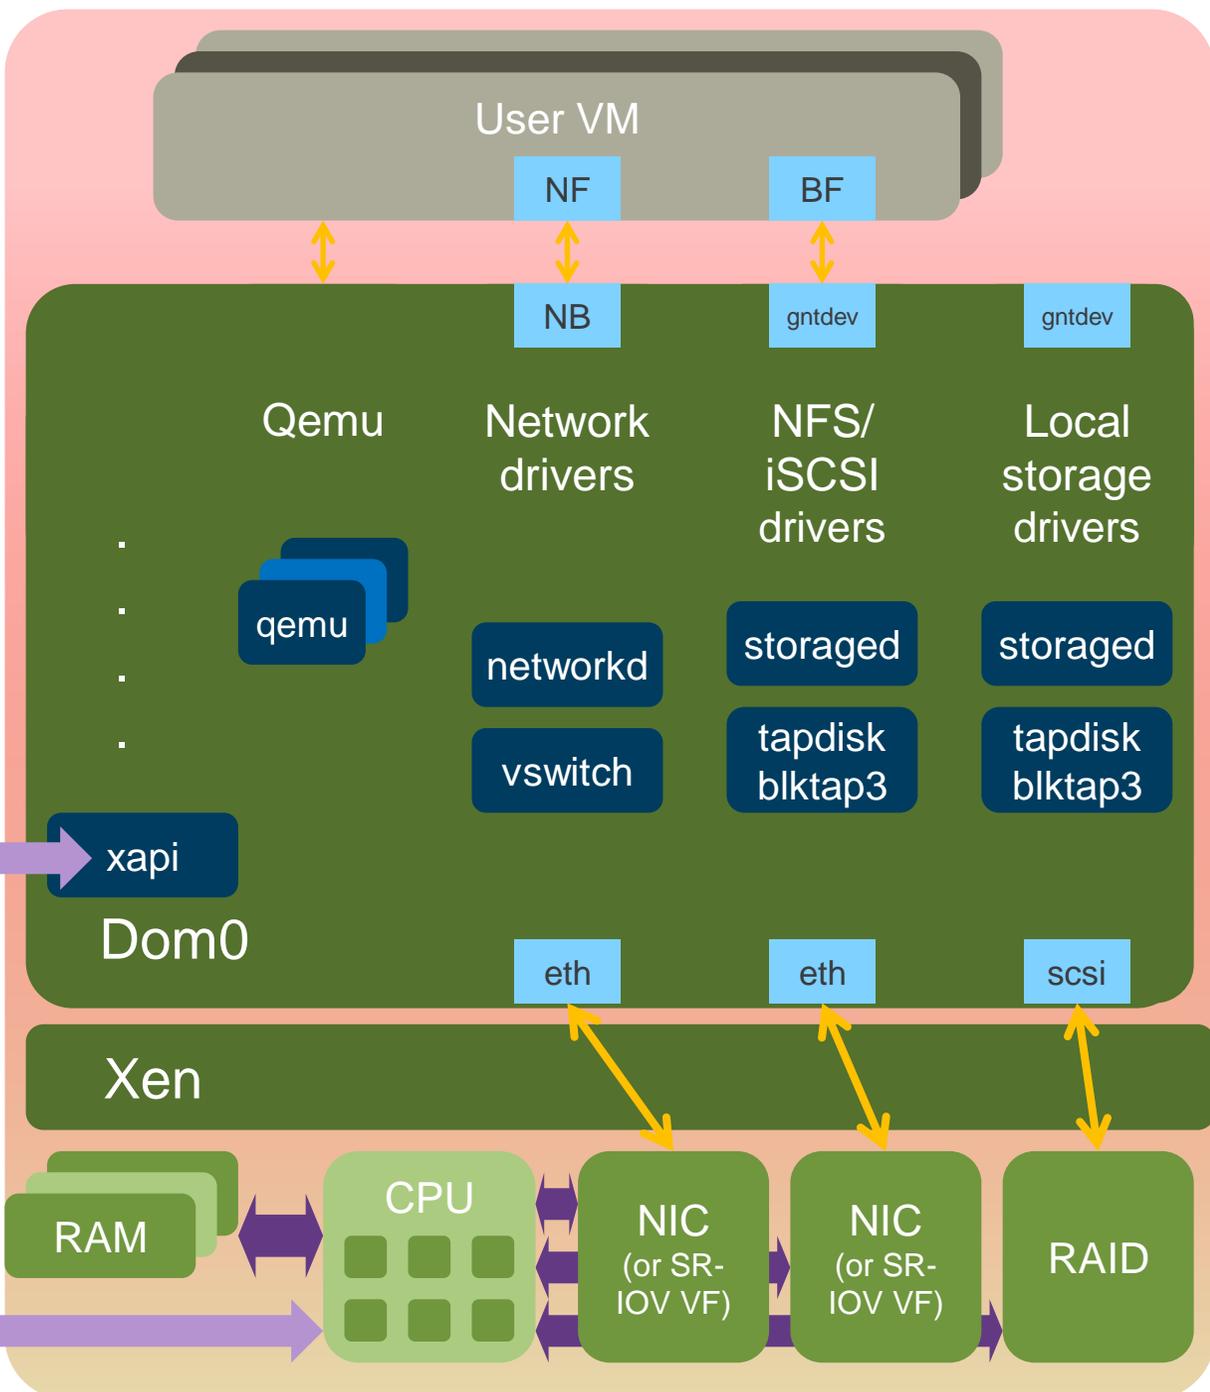
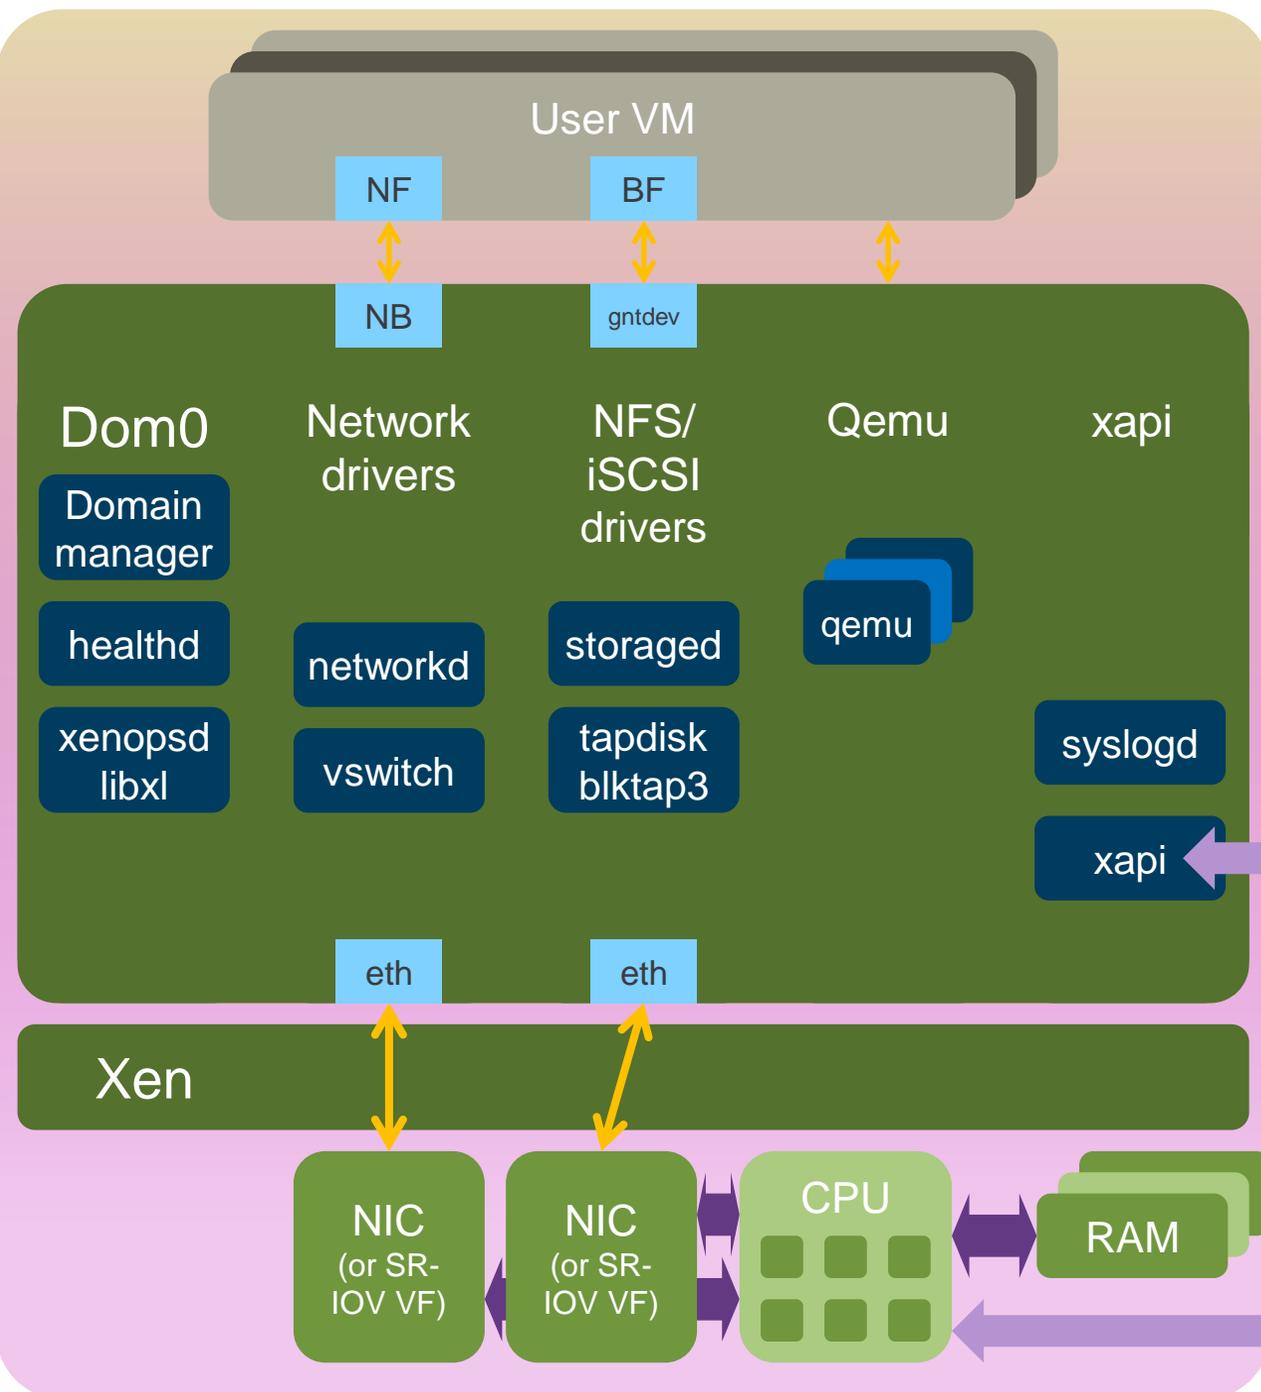
Better Scalability

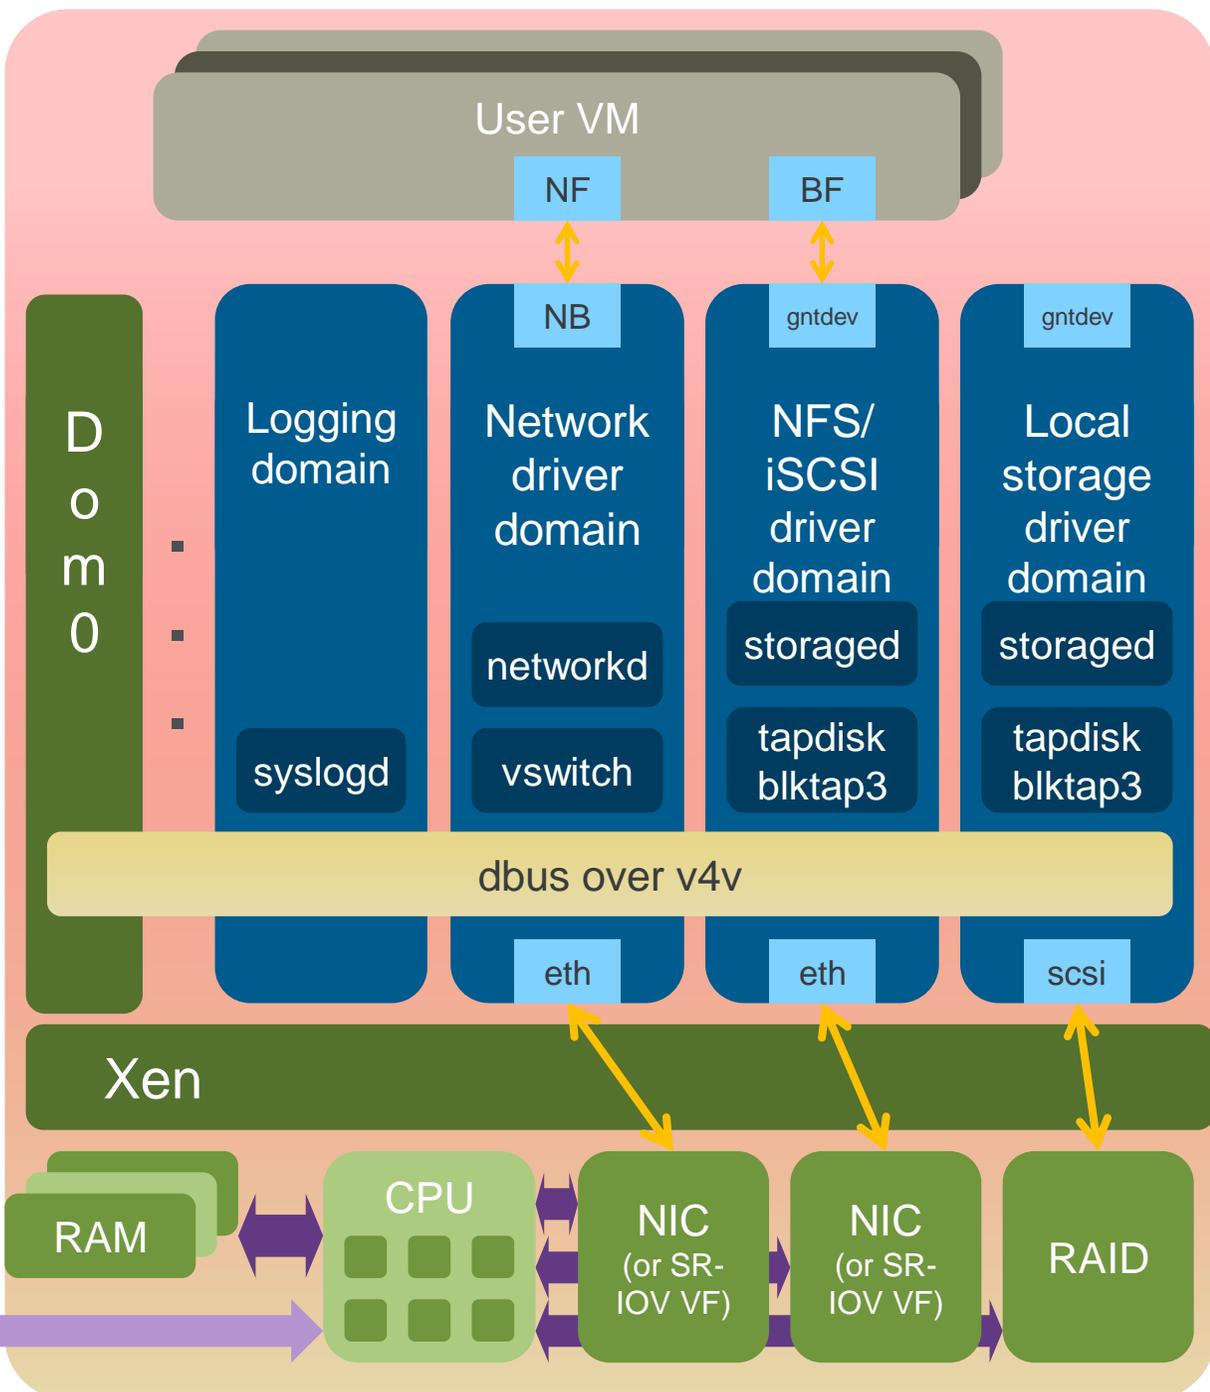
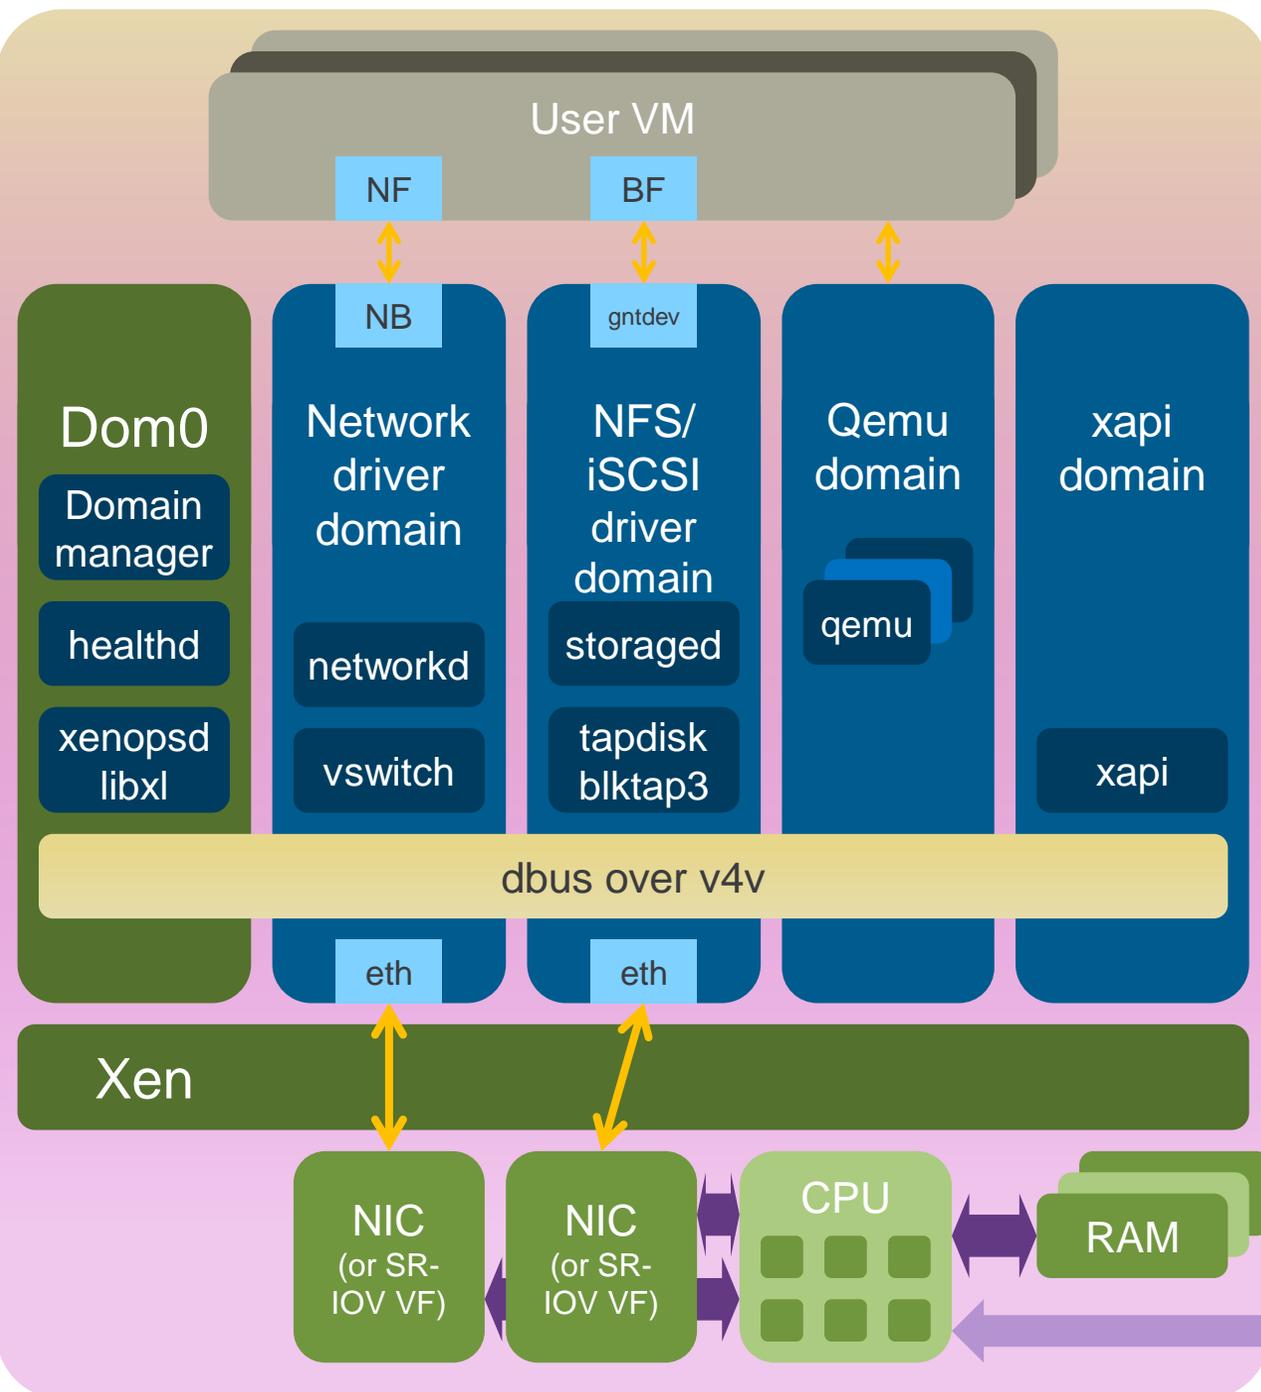


Ability to safely restart parts of the system  
(e.g. just 275ms outage from failed Ethernet driver)



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



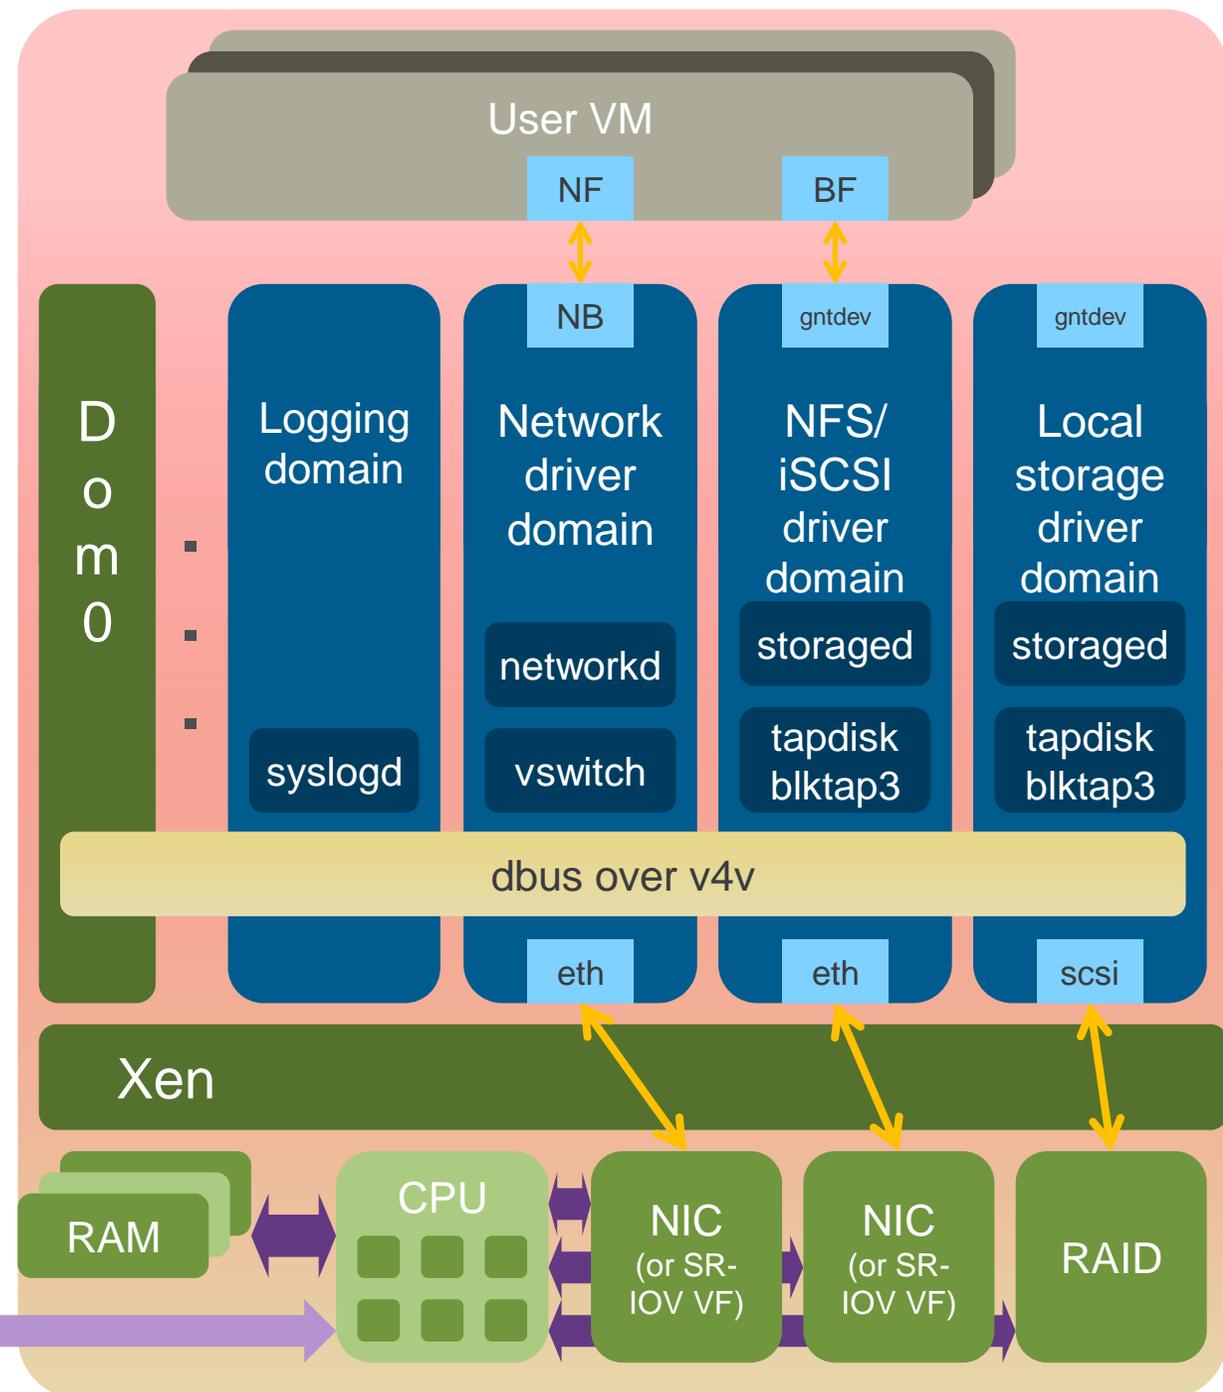
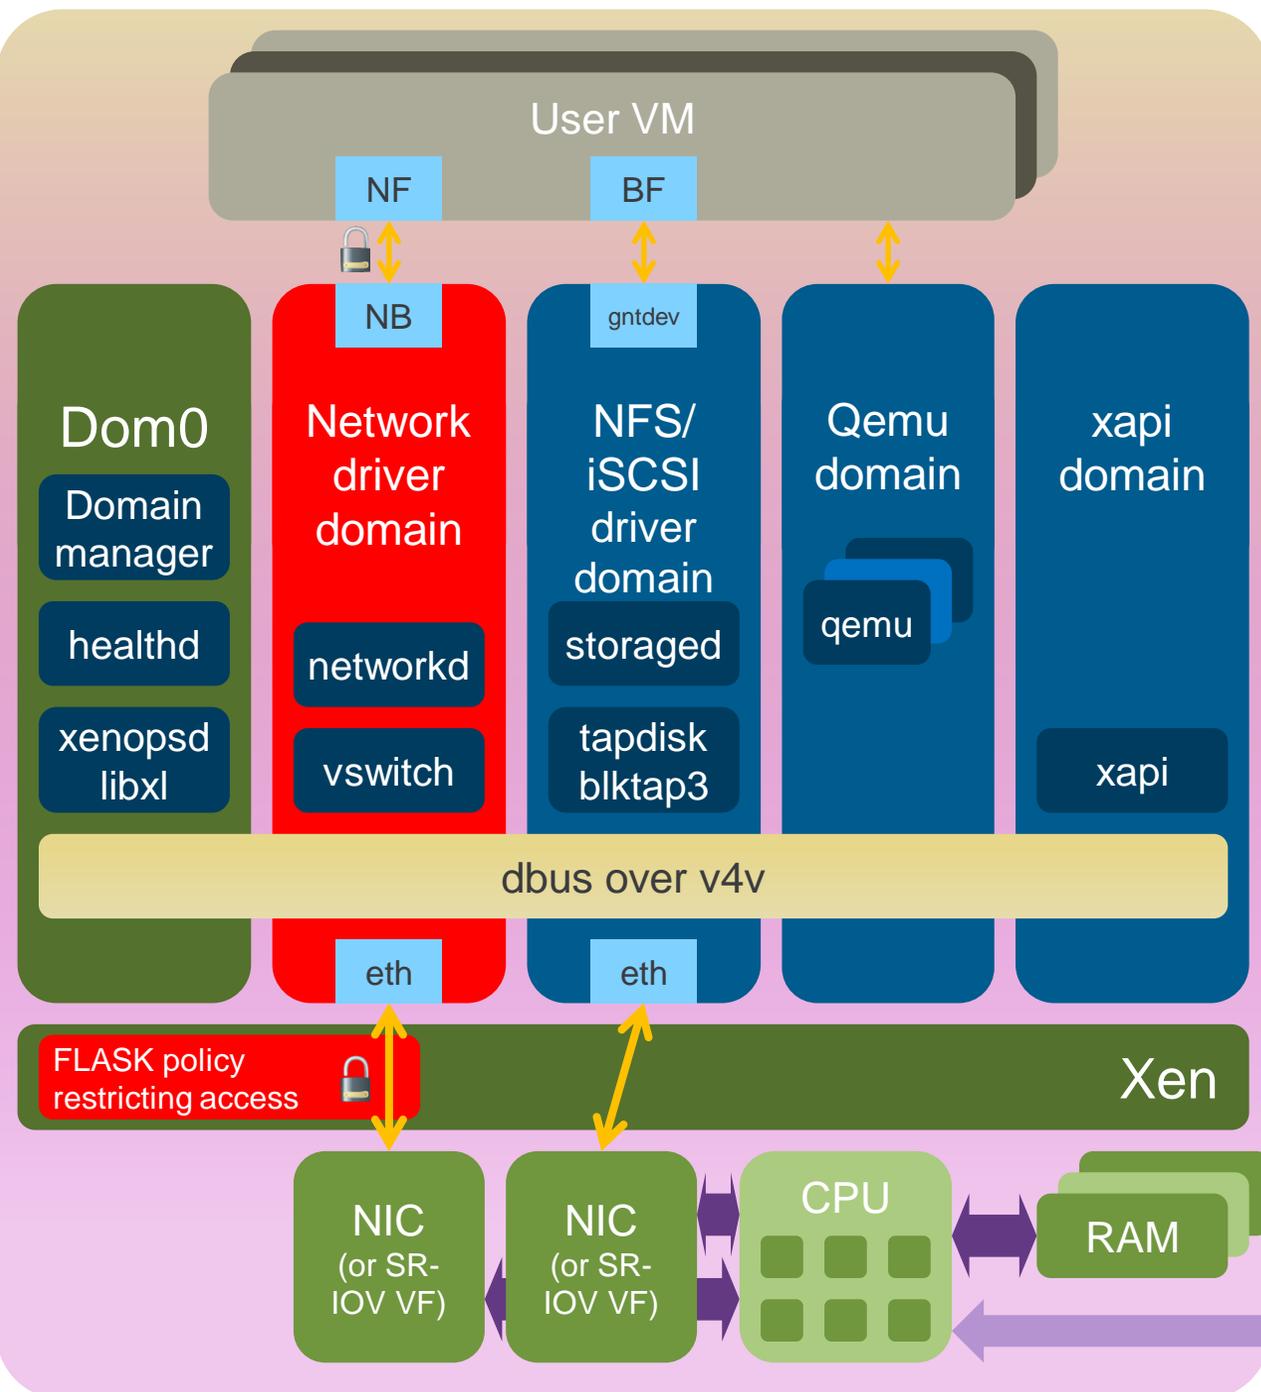


# Xen Security Advantages

- 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)





# Xen Advanced Security Recipes

- Xen has many Security Features (besides the ones I covered)
- Most are **not** switched on by default
- Although most are simple to use, some seem complicated

See: [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)

And more are be coming!



# ARM Hypervisor



# In depth presentation tomorrow ...

## Xen on ARM

- by Stefano Stabellini –  
Xen and Xen-ARM Linux maintainer
- New York III from 10:45 to 11:30
- Including a demo



*One mode to  
rule them all*

# Mirage OS



# Library Operating Systems

Application stacks only running on Xen APIs

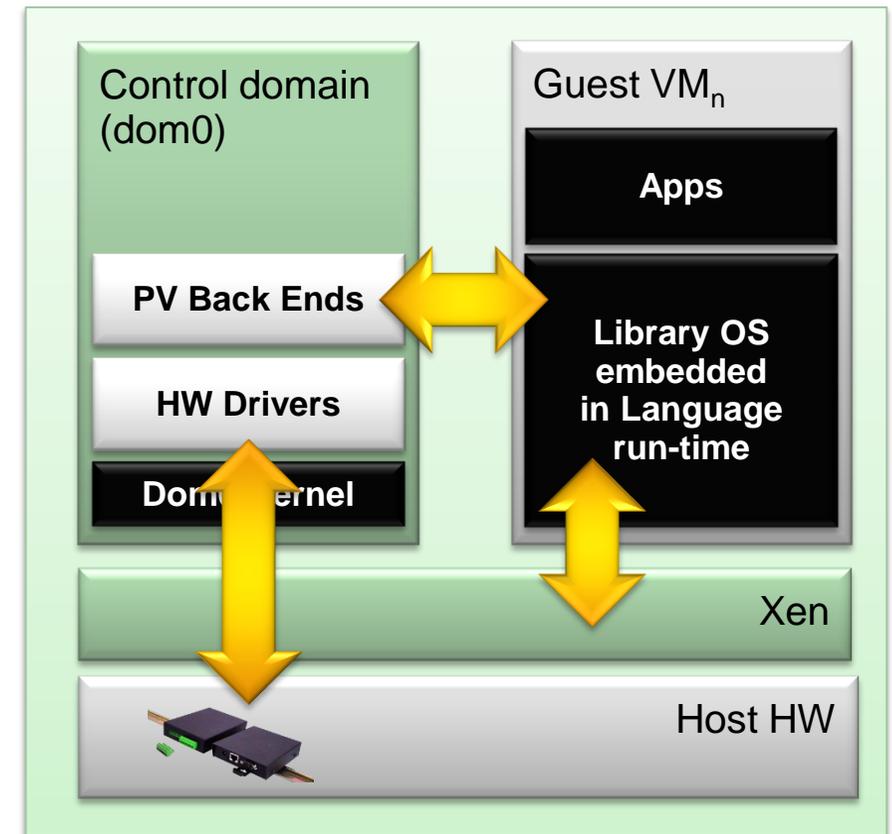
Works on any Xen based cloud or hosting service

## Examples

- ErlangOnXen.org : Erlang
- HalVM : Haskell
- Mirage OS : Ocaml

## Benefits:

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



# Mirage OS

- Recently added to Xen Project incubator
- In beta stage : first release on its way (July 2013)
- Clean-slate protocols implementations, e.g.
  - TCP/IP, DNS, SSH, Openflow (switch/controller), HTTP, XMPP, ...
  - New applications using next generation XAPI (disaggregated XAPI architecture)

More info: [http://www.slideshare.net/xen\\_com\\_mgr/mirage-extreme-specialisation-of-virtual-appliances](http://www.slideshare.net/xen_com_mgr/mirage-extreme-specialisation-of-virtual-appliances)



# Hot Topics and Projects



# Xen 4.3 Release (June 2013)

- Release candidates & Xen Test Days (today, June 5<sup>th</sup>)

---
- Xen ARM for Servers
- Extend scope of Xen Security Modules
- Default to QEMU upstream
- Updated and improved libvirt drivers for Xen
- Lots of other stuff:
  - scalability, performance, better NUMA support, ...



[More info:](http://lists.xen.org/archives/html/xen-devel/2013-05/msg01134.html) <http://lists.xen.org/archives/html/xen-devel/2013-05/msg01134.html>



# Xen 4 + XAPI in CentOS 6 (June 2013)

*“For about a year members of the Xen Project, the CentOS community and large Xen Users have worked on bringing Xen and XAPI to CentOS 6”*

---

Driven by demand from the community :

- Can run Xen on CentOS 6 today, but non-trivial
- We wanted “YUM INSTALL XEN”
- Mostly a packaging problem
- Teams from CentOS, Citrix, Go Daddy & Rackspace
- QA and usability sanity checks

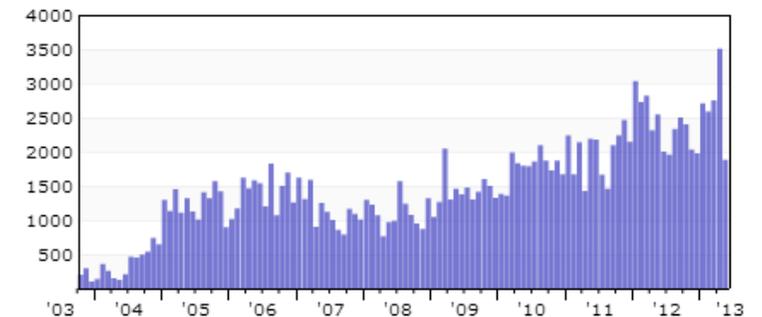


# The Xen Community is Changing

*“Growth is leading to more structure & more collaboration & more openness!”*

---

- Establishing a shared and open test infrastructure
  - Goal: Increase development velocity
- Improved usability and better distro-integration
  - Xen + XAPI in CentOS 6.4
- More focus on downstreams
  - OpenStack and Xen Orchestra
  - Better libvirt and virt-manager integration
- Changing the XAPI / XCP release model
- Xen on ARM and collaboration with Linaro



# Getting Started with Xen Projects

## Online : [xenproject.org](http://xenproject.org) > User & Help menus

- Mailing Lists and IRC
- Q&A System
- Find me and I can get you hooked up!

## Events : [xenproject.org/about/events.html](http://xenproject.org/about/events.html)

- Test Days (IRC at #xentest) ... today, June 5th
- Document Days (IRC at #xendocs) ... the next one is May 28<sup>th</sup>
- User (Sept 18, New Orleans) and Developer Summits (Oct 24-25, Edinburgh)
- Hackathons - last week search for “[Hackathon Minutes]” on xen-devel





- **News:** [blog.xenproject.org](http://blog.xenproject.org)
- **Web:** [xenproject.org](http://xenproject.org) > **Help**
  - Help for IRC, Lists, ...
  - Stackoverflow like Q&A
- **Wiki:** [wiki.xenproject.org](http://wiki.xenproject.org)
- **Presentations:** [slideshare.net/xen\\_com\\_mgr](http://slideshare.net/xen_com_mgr)
- **Videos:** [vimeo.com/channels/xen](http://vimeo.com/channels/xen)

# Thank You!



@lars\_kurth

FREENODE: lars\_kurth



Slides available under CC-BY-SA 3.0  
From [www.slideshare.net/xen\\_com\\_mgr](http://www.slideshare.net/xen_com_mgr)

