Virtualization in the Cloud

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

@lars_kurth
@xen_com_mgr

A Brief History of Xen in the Cloud

Late 90s

XenoServer Project
(Cambridge Univ.)

The **XenoServer project** is building

> *public infrastructure for wide-area distributed computing*.

We envisage a world in which **XenoServer** execution platforms will be scattered across the globe and available for any member of the public to submit code for execution.

Global Public Computing

"*This dissertation proposes a new **distributed computing paradigm**, termed global public computing, which allows*

*any user to run any code anywhere. Such platforms **price***

***computing resources**, and ultimately **charge users for resources** consumed.*"

Evangelos Kotsovinos, PhD dissertation, 2004

A Brief History of Xen in the Cloud

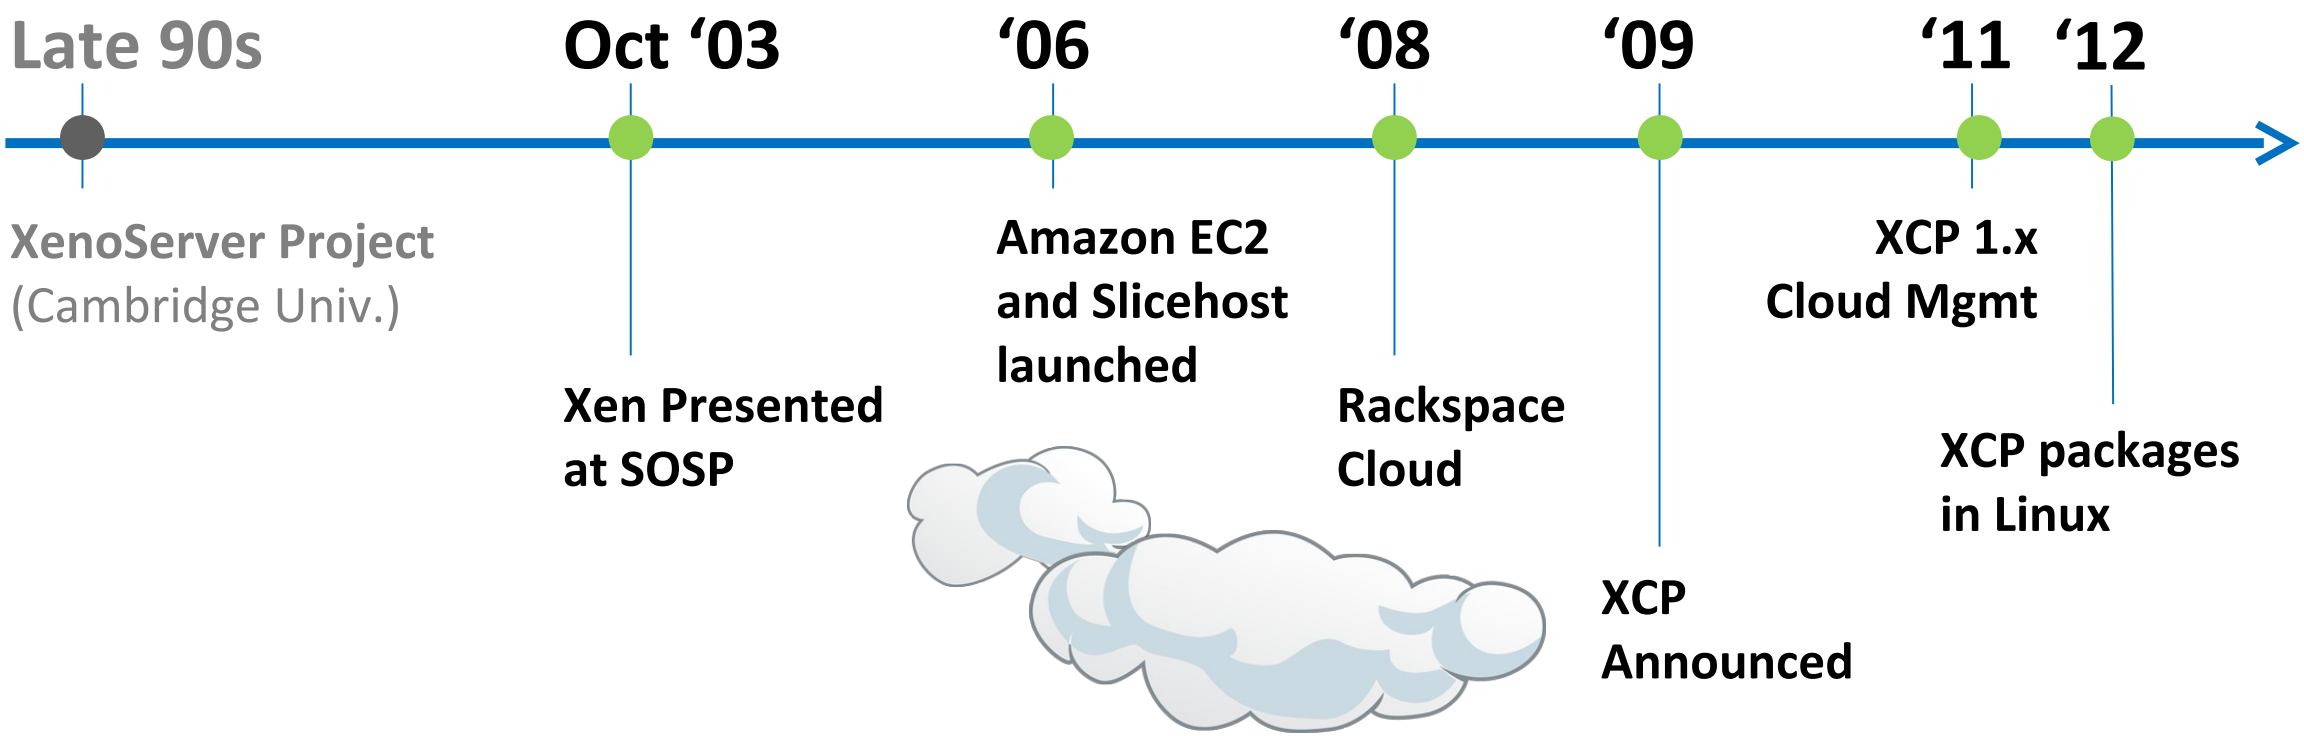

The Xen Hypervisor was designed for the Cloud straight from the outset!

Xen.org

- Guardian of Xen Hypervisor and related OSS Projects
- Xen project Governance similar to Linux Kernel
- Projects
  - Xen Hypervisor (led by Citrix)
  - Xen Cloud Platform aka XCP (led by Citrix)
  - Xen ARM : Xen for mobile devices (led by Samsung)
  - PVOPS : Xen components and support in Linux Kernel (led by Oracle)
- 10+ vendors contributing more than 1% to the project
  (AMD, Citrix, Fujitsu, Huawei, Intel, Novell, Oracle, Samsung, Suse, …)

Xen Overview

Basic Xen Concepts

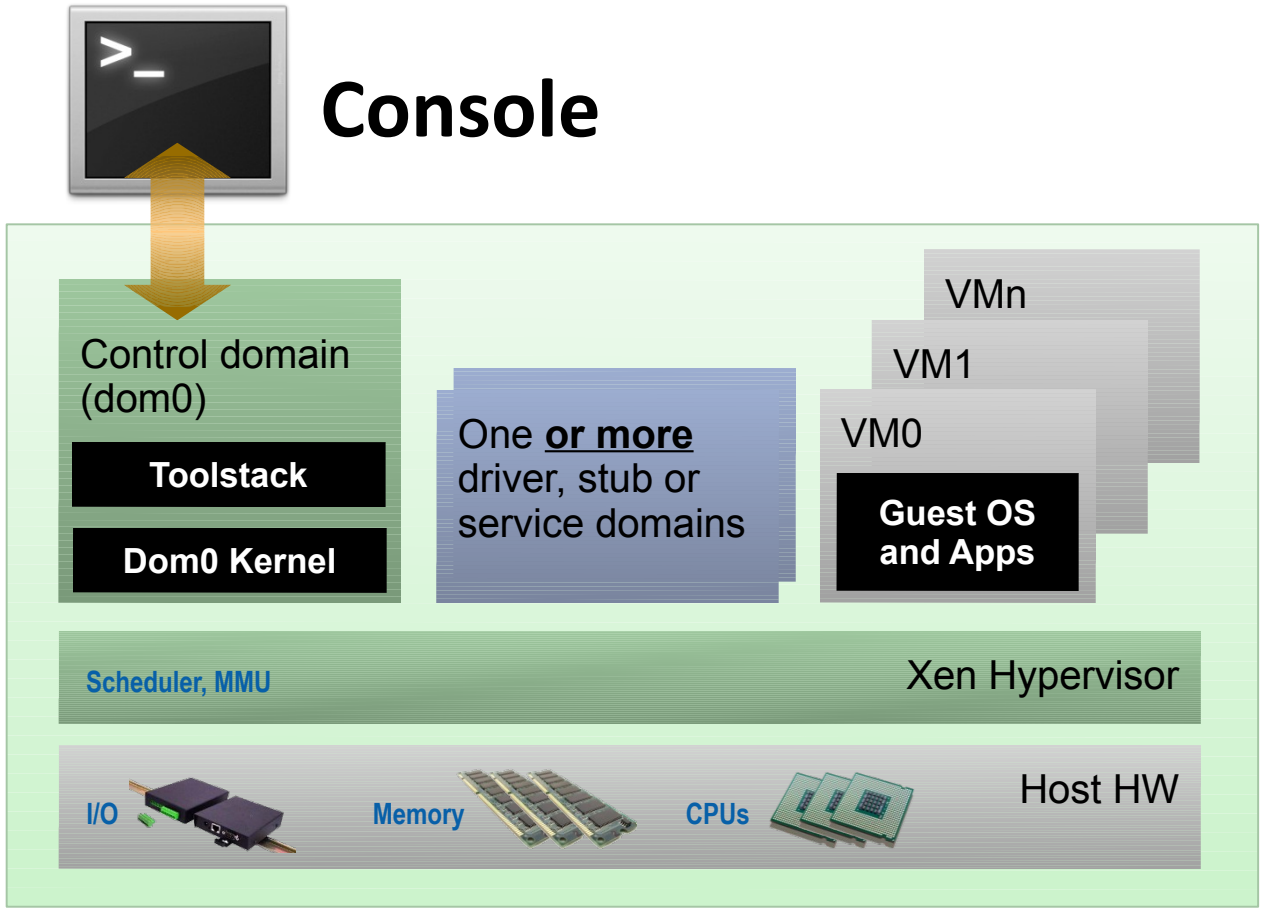

- Control Domain aka Dom0
  - Dom0 kernel with drivers
  - Xen Management Toolstack
  - Trusted Computing Base
- Guest Domains
  - Your apps
  - E.g. your cloud management stack
- Driver/Stub/Service Domain(s)
  - A “driver, device model or control service in a box”
  - De-privileged and isolated
  - Lifetime: start, stop, kill

Xen Variants for Server & Cloud

	Xen 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** **XCP from Xen.org**
Products	**Oracle VM**	**Huawei UVP**	**Citrix XenServer**
Used by …	amazon web services™	**Many Others**	the rackspace cloud

Xen

Xen : Types of Virtualization

PV Domains & Driver Domains

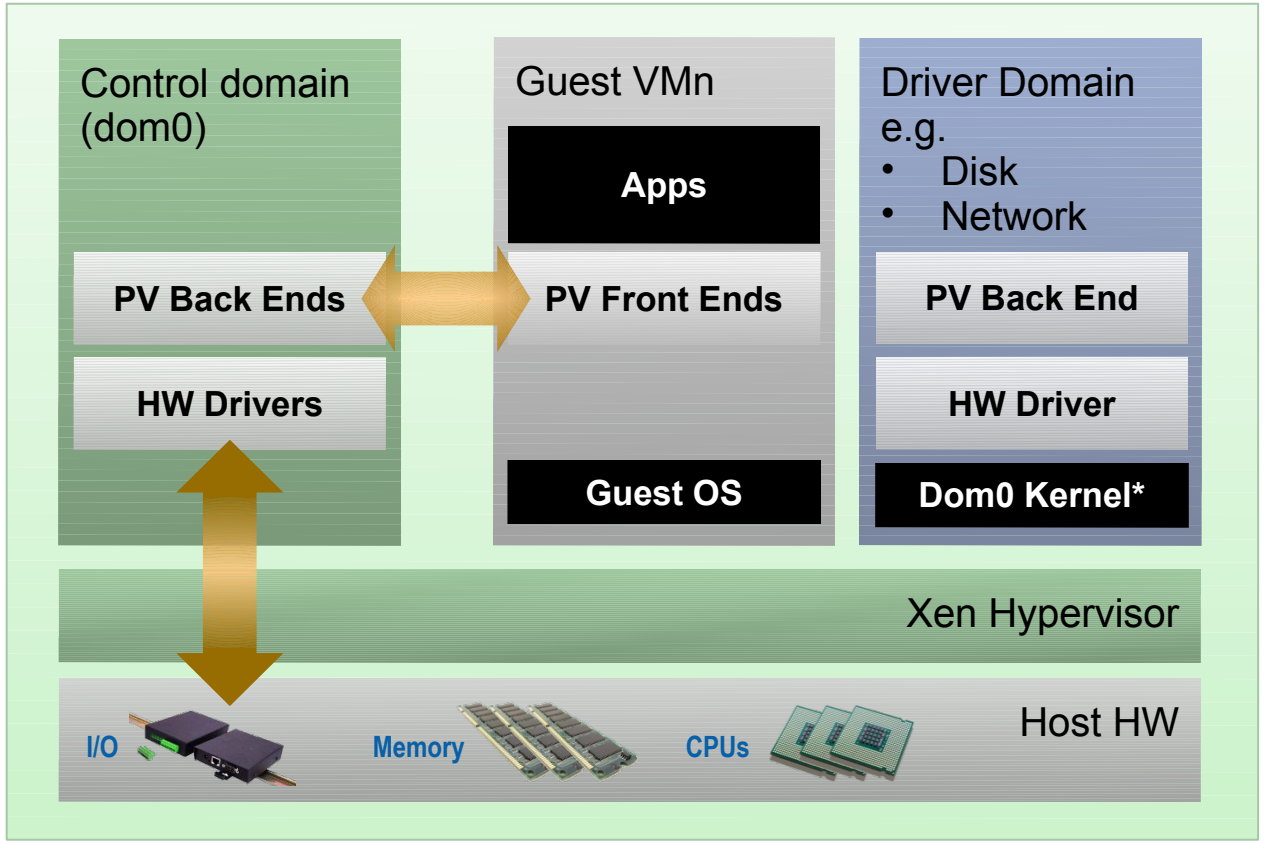

*) Can be MiniOS

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

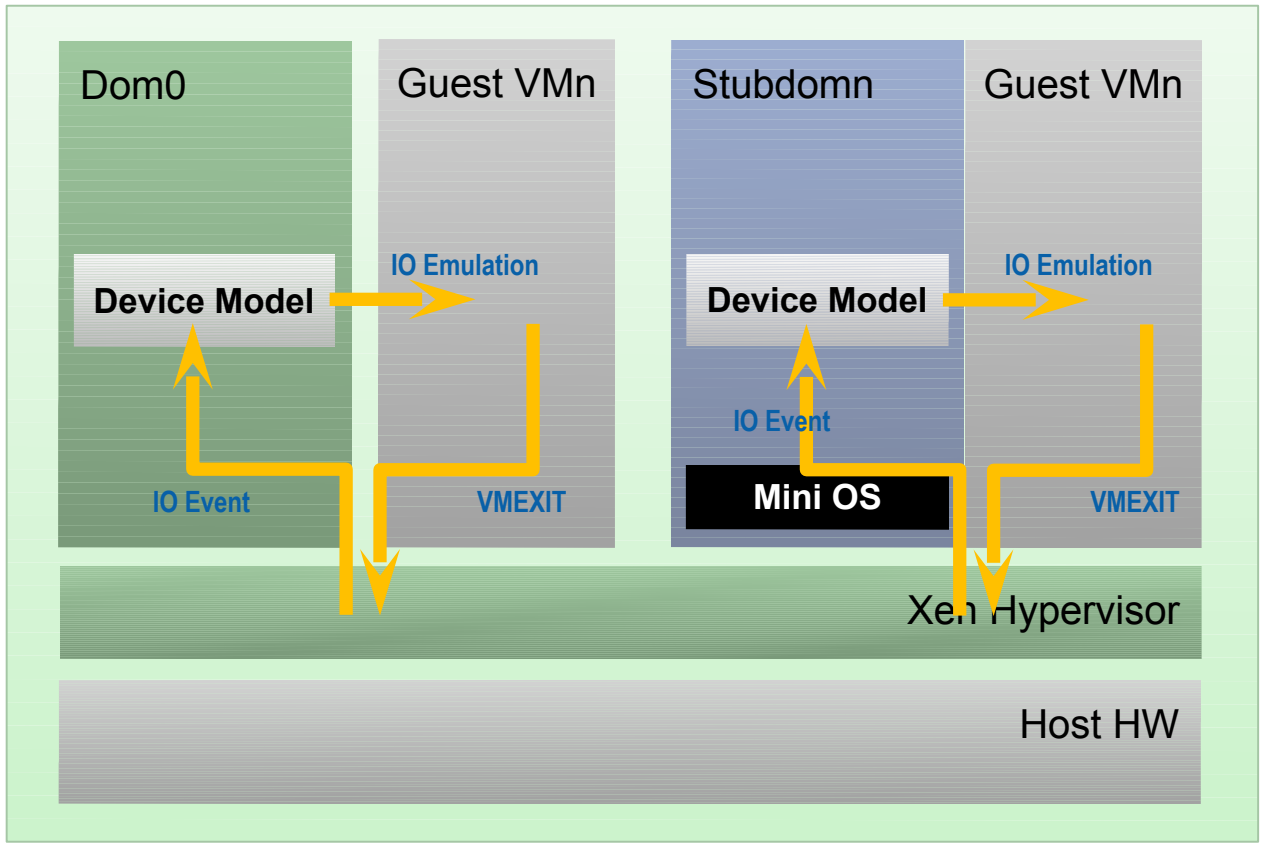

Disadvantages

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

Advantages

- Install the same way as native Linux

Stub Domains

- Security
- Isolation
- Reliability and Robustness

PV on HVM

- A mixture of PV and HVM
- Linux enables as many PV interfaces as possible
- This has advantages
  - install the same way as native
  - PC-like hardware
  - access to fast PV devices
  - exploit nested paging
  - Good performance trade-offs
- Drivers in Linux 3.x

	HVM	PV on HVM	PV
Boot Sequence	Emulated	Emulated	PV
Memory	HW	HW	PV
Interrupts, Timers & Spinlocks	Emulated	PV*	PV
Disk & Network	Emulated	PV	PV
Privileged Operations	HW	HW	PV

*) Emulated for Windows

Xen and Linux

Xen and the Linux Kernel

Xen was initially a University research project

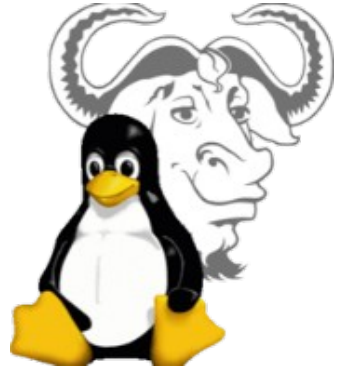

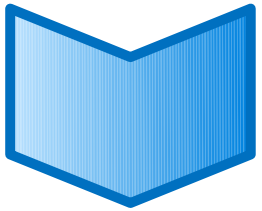

Invasive changes to the kernel to run Linux as a PV guest and Dom0

Current State

PVOPS Project

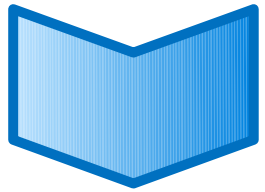

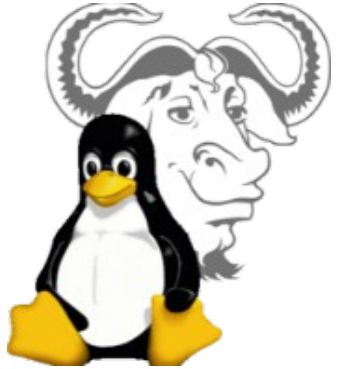

Xen support in Linux 3.0+

(it is functional but not yet fully optimized)

On-going optimization work in Linux 3.4 +

What does this mean?

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


“Security and QoS/Reliability are amongst the top 3 blockers for cloud adoption”
www.colt.net/cio-research
Xen

Security and the Next Wave of Virtualization

- Security is a key requirement for Cloud
- Security is the primary goal of virtualization on the Client
  - Desktop, Laptops, Tablets & Smart Phones
- Maintaining isolation between VMs is critical

Xen Security & Robustness Advantages

- Even without Advanced Security Features
  - Well-defined trusted computing base (much smaller than on type-2 HV)
  - No extra services in hypervisor layer
- **More Robustness:** Mature, Tried & Tested, Architecture
- Xen Security Modules (or XSM)
  - Developed and contributed to Xen by NSA
  - Generalized Security Framework for Xen
- Can also run with SELinux

Advanced Security: Disaggregation

- Split Control Domain into Driver, Stub and Service Domains
  - Each contains a specific set of control logic
  - See: ”Breaking up is hard to do” @ Xen Papers
- Unique benefit of the Xen architecture
  - **Security**: Minimum privilege; Narrow interfaces
  - **Robustness:** ability to safely restart parts of the system
  - **Performance:** lightweight, e.g. Mini OS directly on hypervisor
  - **Scalability:** more distributed system (less reliable on Dom0)
- Used by Qubes OS and Citrix XenClient XT

Advanced XenClient Architecture

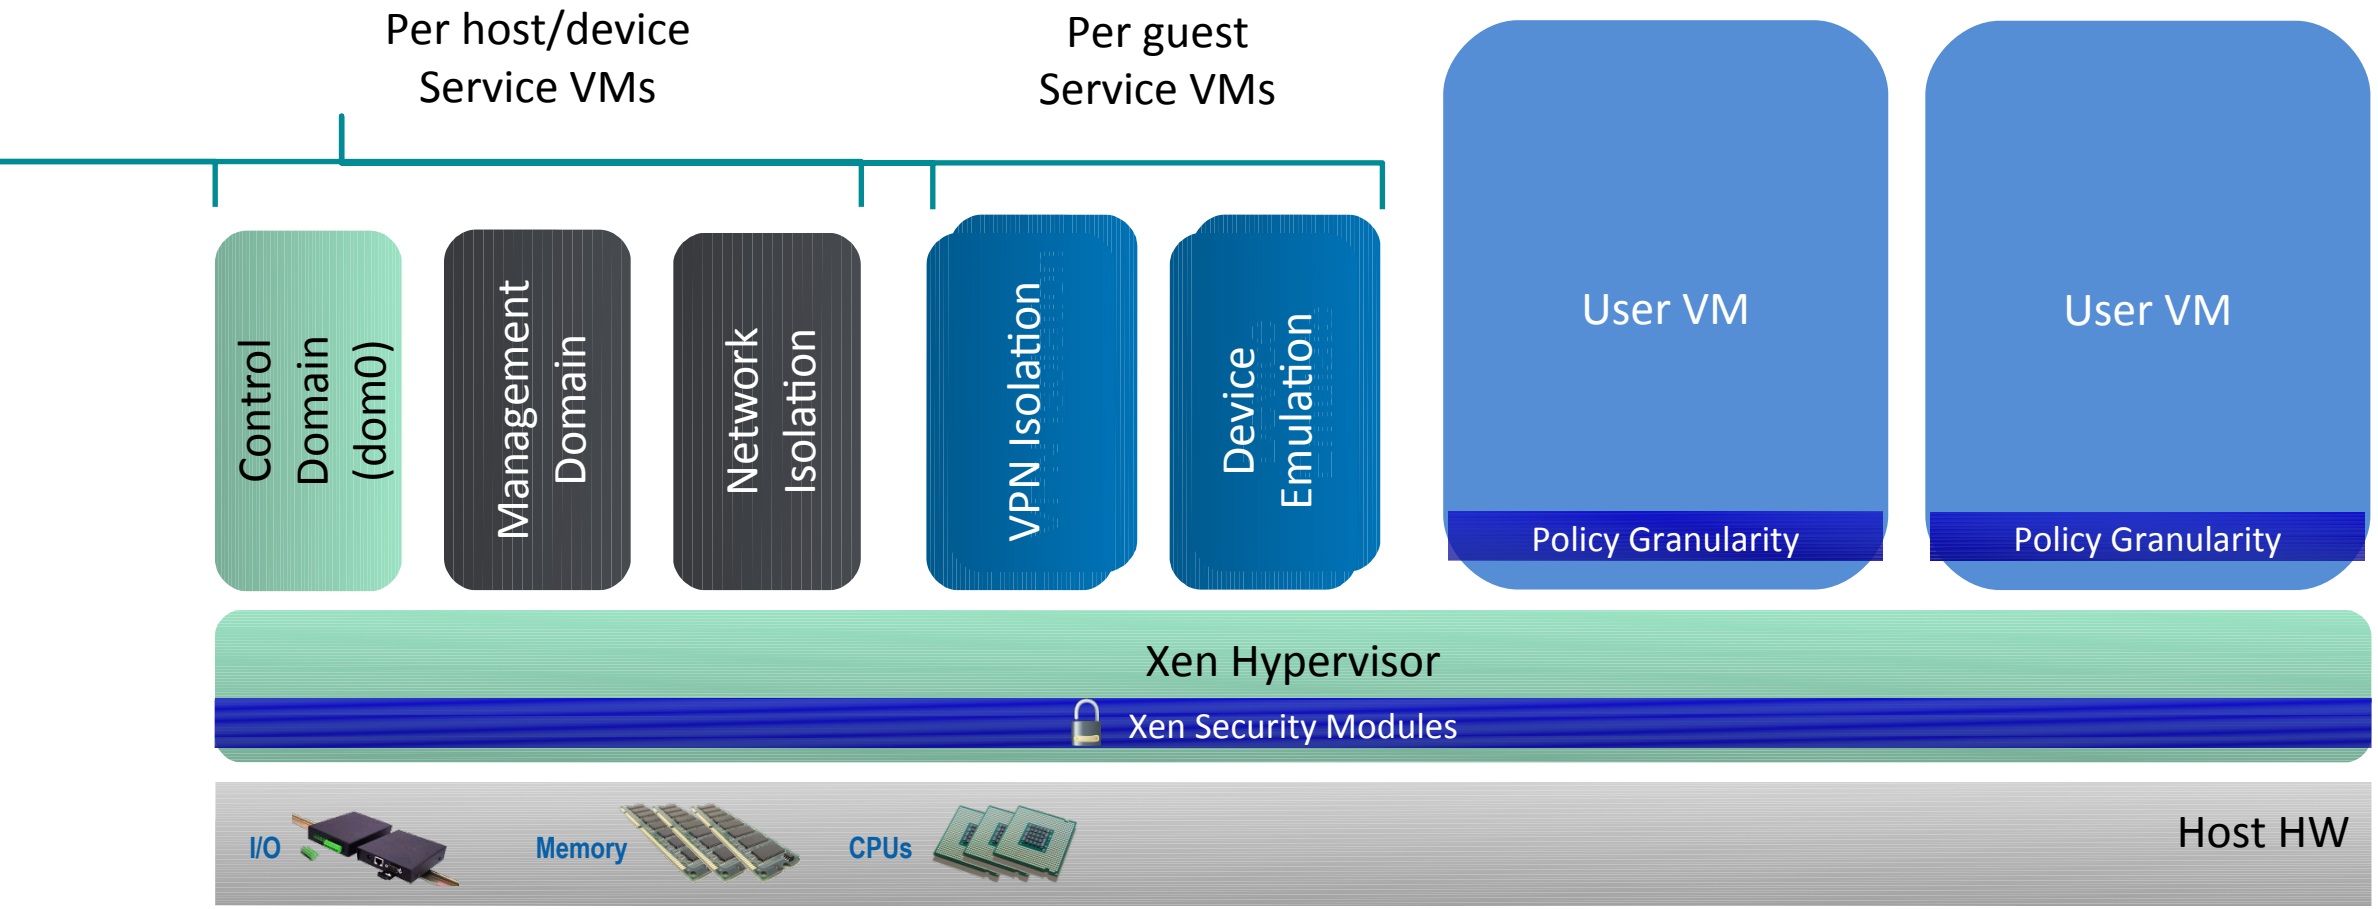

Example: Network Driver Domain for HA

- Detect failure e.g.
  - Illegal access
  - Timeout
- Kill domain, restart
  - E.g. Just 275ms outage from failed Ethernet driver
- Auto-restarts to enhance security

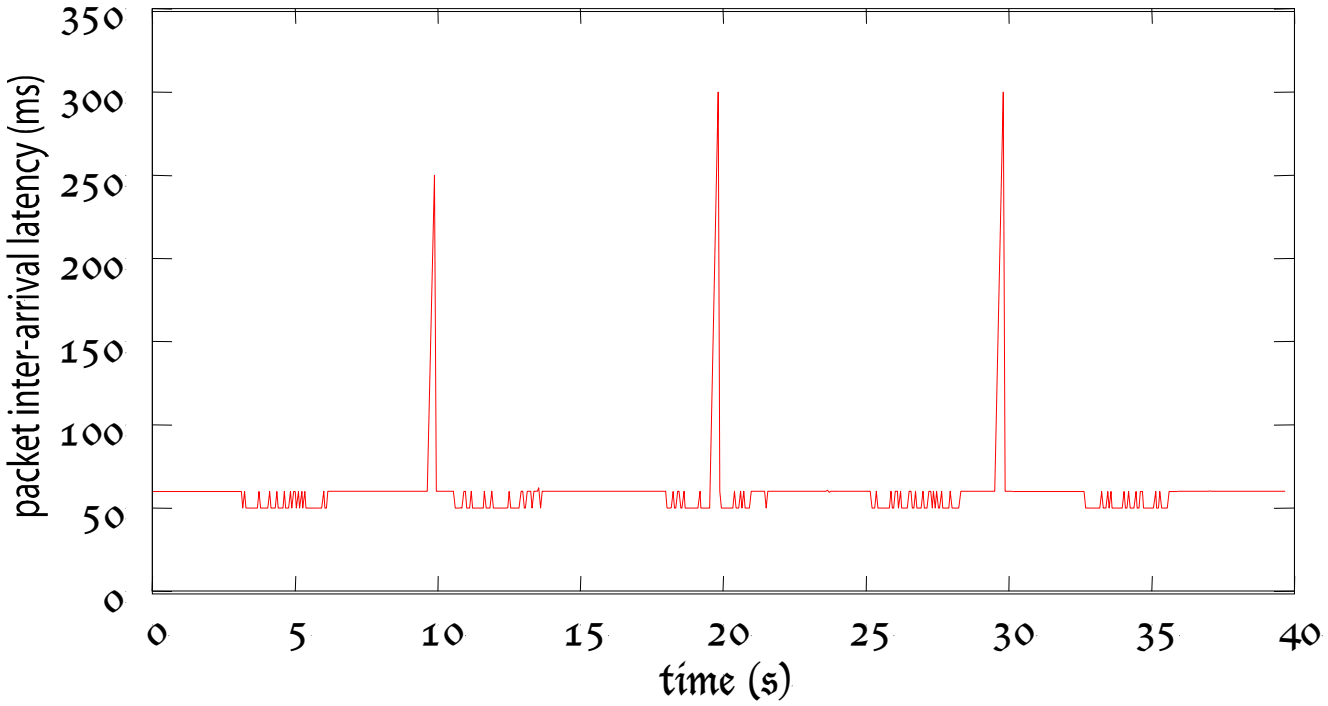

BUT…

- Today, XCP and commercial Xen based Server products
  - Do not yet make use of XSM
  - Do not yet make use of Advanced Security Features (Disaggregation)
- In XCP, work has started to add these features
  - Various articles of how this may be done on the xen wiki
  - More information soon (likely at XenSummit)
  - Commitment on improving docs for Security, Reliability & Tuning

XCP Project

XCP

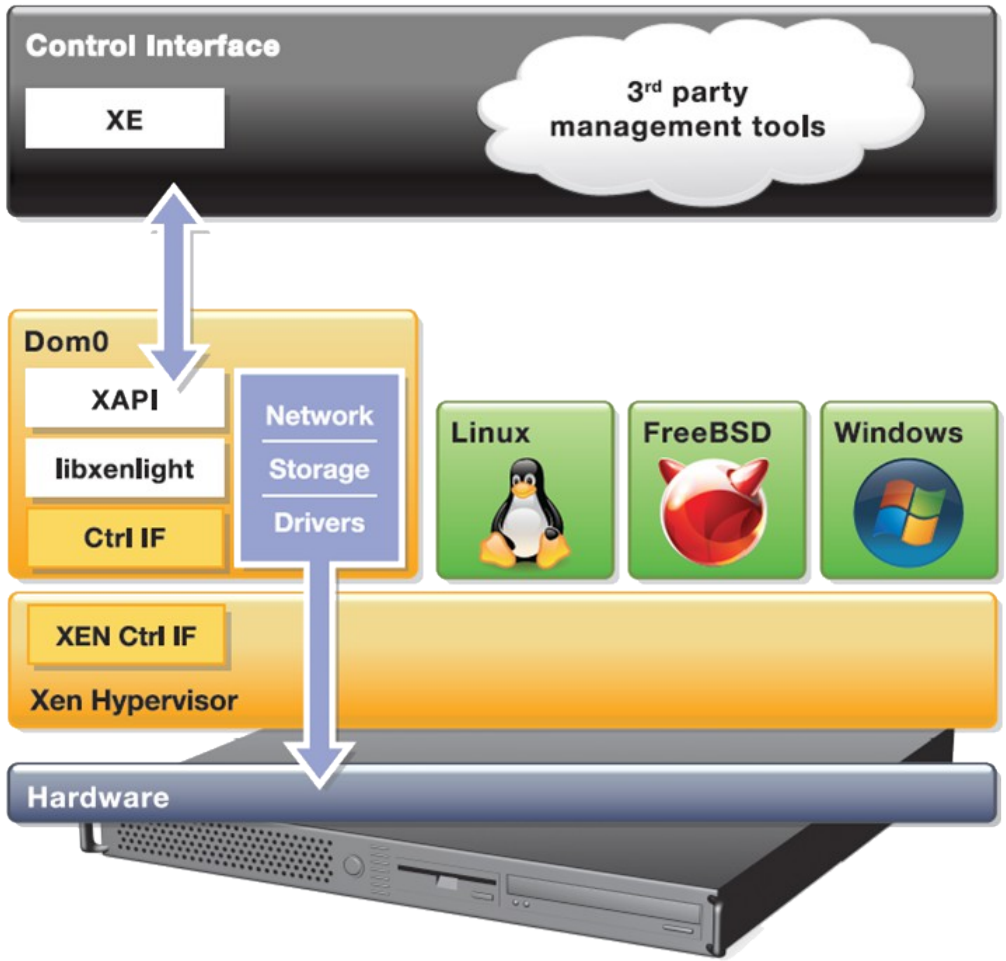

- Complete vertical stack for server virtualization
- Distributed as a closed appliance (ISO) with CentOS 5.5 Dom0, misc DomU's, network & storage support and Xen API
- Open source distribution of Citrix XenServer

XCP Overview

- Open source version of Citrix XenServer
  - wiki.xen.org/wiki/XCP/XenServer_Feature_Matrix
- Enterprise-ready server virtualization and cloud platform
  - Extends Xen beyond one physical machine and other functionality
  - Lots of other additional functionality compared to Xen
- Datacenter and cloud-ready management API
  - XenAPI (XAPI) is fully open source
  - CloudStack and OpenStack integration

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

XCP 1.5 (in beta)

- **Architectural Improvements:** Xen 4.1, GPT, smaller Dom0
- **GPU pass through:** for VMs serving high end graphics
- **Performance and Scalability:**
  - 1 TB mem/host
  - 16 VCPUs/VM, 128 GB/VM
- **Networking:** Open vSwitch (default), Active-Backup NIC Bonding
- **Virtual Appliance:** multi-VM and boot sequenced, OVF support
- More guest OS templates

XAPI: What is it?

- XAPI is the backbone of XCP
  - Provides the glue between all components
  - Is the backend for all management applications
  - Also called XenAPI
- It's a XML-RPC style API, served via HTTPS
  - Provided by a service on every XCP dom0 host
  - Designed to by highly programmable
  - API bindings for many languages: .NET, Java, C, Powershell, Python
- XAPI is Extensible via plugins
  - E.g. used by OpenStack

XCP-XAPI Packages in Linux

- Make the XAPI toolstack independent of CentOS 5.5
- Extend the delivery model
  - Deliver Xen, XAPI and everything in between (storage manager, network support, OCaml libs, etc.) via your favorite Linux distro
    “apt-get install xcp-xapi” or “yum install xcp-xapi”

- Debian 7.0 “Wheezy"

- Ubuntu 12.04 LTS
- Next: Fedora & CentOS

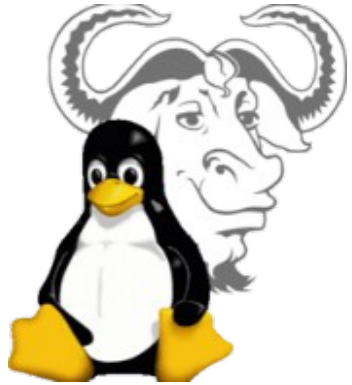

XAPI Management Options

- XAPI frontend command line tool: XE (tab-completable)
- Desktop GUIs
  - Citrix XenCenter (Windows-only)
  - OpenXenManager (open source cross-platform XenCenter clone)
- Web interfaces
  - Xen VNC Proxy (XVP)
  - XenWebManager (web-basedclone of OpenXenManager)
- XCP Ecosystem:
  - xen.org/community/vendors/XCPProjectsPage.html
  - xen.org/community/vendors/XCPProductsPage.html

XCP and Cloud Orchestration Stacks

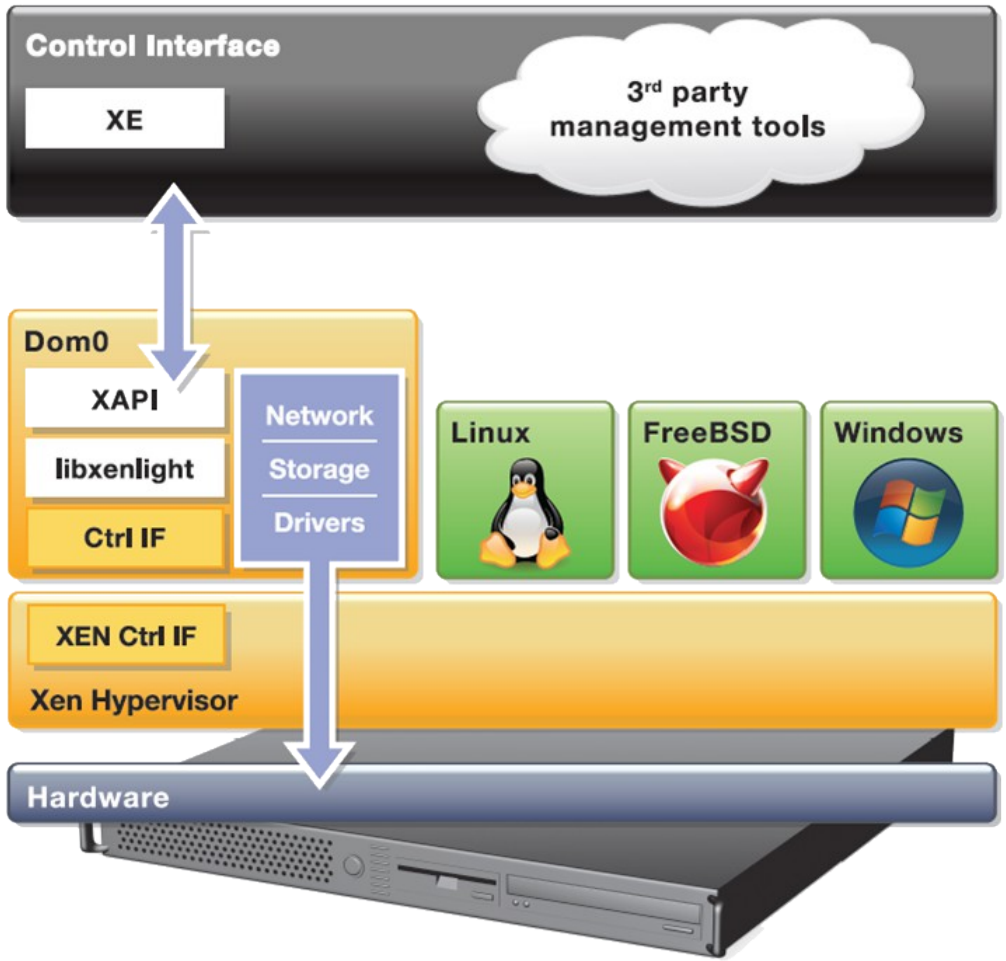

cloudstack
open source cloud computing

OpenNebula.org

openstack™

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 & XAPI
  - Ready for use with cloud orchestration stacks
  - XCP-XAPI packages in Linux distros: flexibility and choice
  - Lots of additional improvements for cloud coming in 2012
- Open Source with a large community and eco-system

- **IRC:** ##xen @ FREENODE
- **Mailing List:** xen-users & xen-api
- **Wiki:** wiki.xen.org
  - Beginners & User Categories
  - XCP Category
- **Excellent XCP Tutorials**
  - A day worth of material @ xen.org/community/xenday11
- **Ecosystem pages**

Questions …