Xen and the path to Ubiquitous Virtualization

Ian Pratt, XenSource-Citrix and

Chairman of xen.org

Outline

- Xen Project Goals
- Virtualization Benefits
- Xen’s Architectural Advantages
- From Server to Client to Mobile
- Roadmap

Xen Project Mission

- Build the industry standard open source hypervisor
  - Core "engine" that is incorporated into multiple vendors' products
- Maintain Xen's industry-leading performance
  - Be first to exploit new hardware acceleration features
  - Help OS vendors paravirtualize their OSes
- Maintain Xen's reputation for stability and quality
  - Security must now be paramount
- Support multiple CPU types; big and small systems
  - From server to client to mobile phone
- Foster innovation
- Drive interoperability

Xen Community: Strong & Productive

- Over 200 contributors to the 3.x series
- Vendors optimize Xen for their products
  - CPU and I/O vendors; OSVs; Mgmt vendors
- Research community
  - Develop new Xen features
  - Explore entire new uses of virtualization
  - Many Universities, IBM, HP, Intel, NSA
- User community
  - Amazon, Google, Oracle, MySpace, hosting providers
- Xen.org and the new Xen Advisory Board
  - Management oversight, trademark policy etc

First Virtualization Benefits

- Server consolidation
  - Consolidate scale-out success
  - Exploit multi-core CPUs
- Manageability
  - Secure remote console
  - Reboot / power control
  - Performance monitoring
- Ease of deployment
  - Rapid provisioning
- Disaster Recovery
- Ease of hardware upgrade/replacement
  - Portability: no need to upgrade OS due to new h/w

2nd Generation Virtualization Benefits

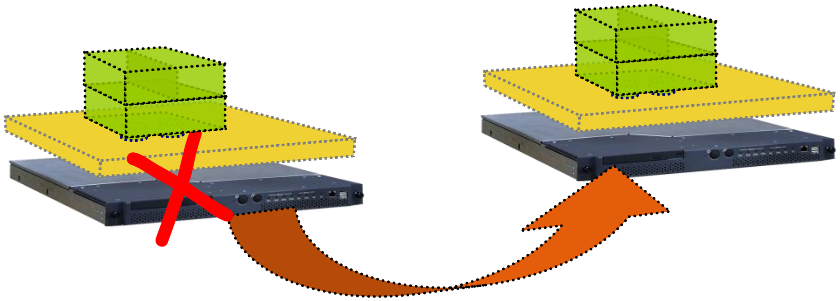

- ***Avoid downtime*** **with VM Relocation**

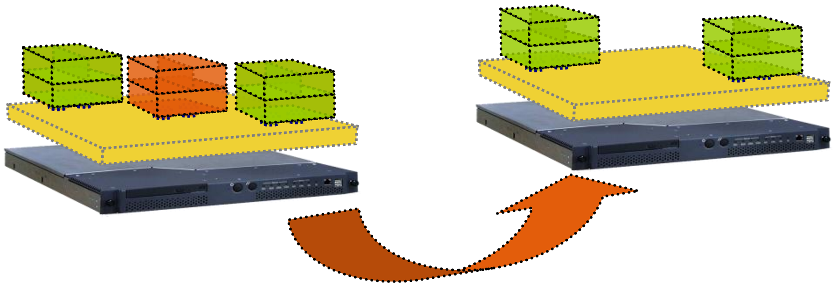

- **Dynamically *re-balance workload* to guarantee application SLAs**

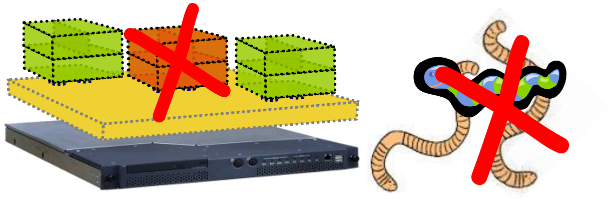

- **Enforce *security* policy**

2nd Generation Virtualization Benefits

- Resource pools
  - Zero-downtime maintenance
  - Load balancing
  - High Availability / Fault Tolerance
- Administrative policy enforcement
  - Backup, Firewalls, Malware scanning etc.
- Abstracting physical world complexity
  - E.g. multi-path storage and networking
- Simplifies application-stack certification
  - Certify app-on-OS, OS-on-hypervisor, hypervisor-on-h/w
  - Enables Virtual Appliances
- Excellent performance
  - Using hardware extensions and OS paravirtualization

Unlocking Hardware Innovation

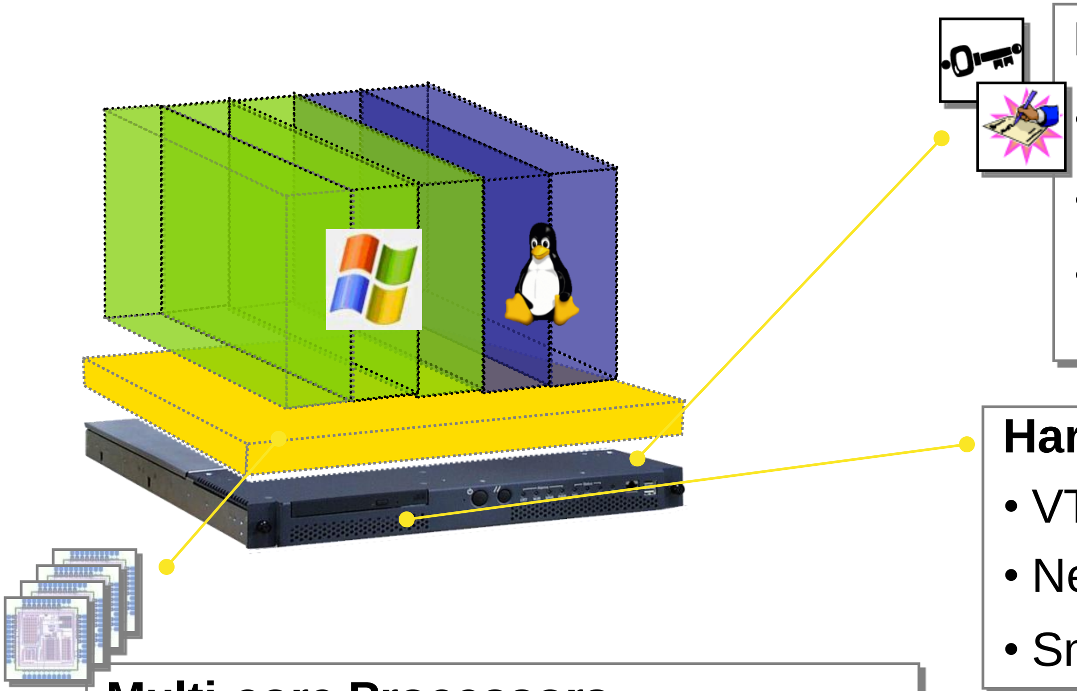

Enhanced Security

- TPM and secure boot (TXT)
- IOMMU and VT-d
- Integrated IDS & security features

Hardware Virtualization Support

- VT/AMDV
- Nested Page Tables (EPT/VMI)
- Smart NICs and HBAs

Multi-core Processors

- More efficient utilization
- Use to hide complexity from guests
- Xen supports SMP guests

Xen always first to take advantage of new hardware features

OS Paravirtualization

- Marketing term: “OS Enlightenment”
- An OS that understands it is running virtualized can be much more co-operative and will thus achieve better performance
  - – Network, disk, memory, time, SMP
- Now adopted by all major OS vendors:

Microsoft Windows redhat Novell. Sun microsystems

- Complements hardware virtualization assistance to yield excellent performance

Realizing Xen’s Architectural Advantages

- Xen’s true hypervisor architecture enables excellent security and scalability
- Lightweight service VMs
  - I/O driver domains and utility domains
  - Device emulation domains
  - Domain building / measurement domains
- Allows efficient large SMP scalability
- Minimum privilege, small TCB
  - De-privilege and disaggregate “domain 0”
- Hypervisor necessary for secure boot (Intel TXT)
- OS agnostic

Ubiquitous Virtualization

- The overhead of virtualization is getting smaller:
  - Through hardware assistance
    - CPU : VT/AMDV, NPT/EPT, ASIDs, APIC
    - Chipset : IOMMU
    - I/O : multi-queue NICs, self-virtualizing NICs and HBAs
  - Through targeted paravirtualization of OSes
    - Particularly higher-level paravirtualization

➔ Near-zero overhead

- Allows always-on virtualization
- Even for a single high-performance VM

- Xen’s goal : be the best choice for ubiquitous deployment

Xen Roadmap for Servers

- Performance and scalability optimizations
  - Larger numbers of physical and virtual CPUs
- Native Microsoft Enlightenment support
- Security hardening
  - Domain0 disaggregation
  - Automated penetration testing
  - Immutable memory
- Enable Smart IO devices
  - Key to reducing IO overhead, particularly for Network

Hardware Accelerated I/O

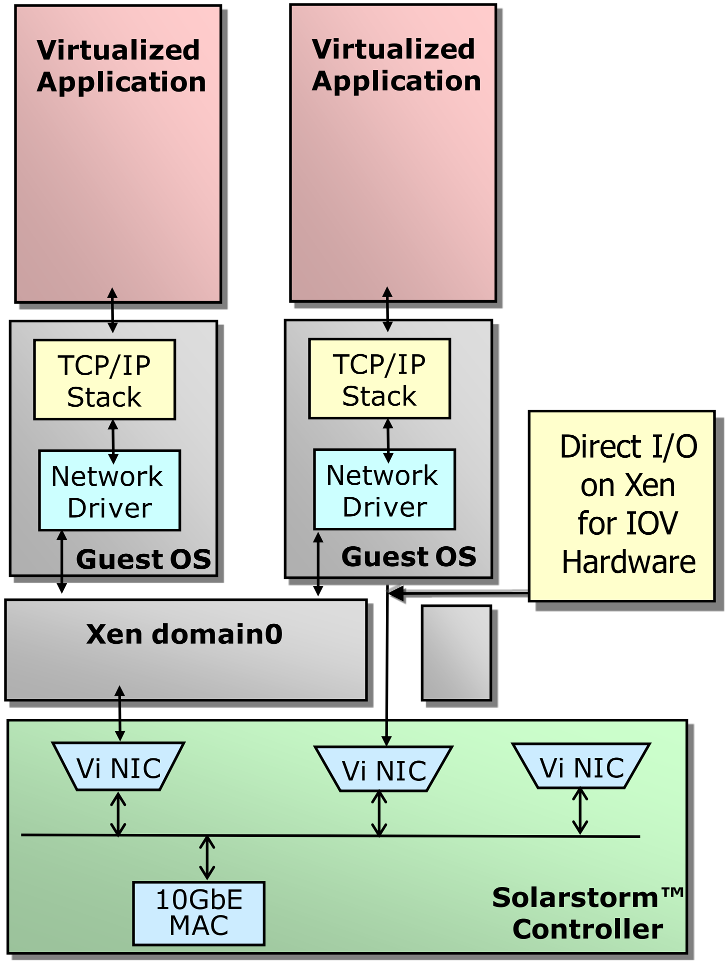

Guest-direct I/O for performance-sensitive workloads

✓ Hardware enforced protection, isolation and virtualization

✓ Hardware assist for routing IP flows to guests

✓ Driver supports both IOV and traditional host-multiplexed I/O

From Servers to Clients

- Security and manageability are key drivers for client virtualization
  - Service partitions; multi-level secure VMs; “BYOPC”
  - “Instant-on” VM’s for web browsing, email etc
- Hypervisor needs to be able to attest information about the platform to guests (using TXT/TPM)
- Preparing Xen for client
  - IOMMU device pass-through
  - Enhanced power management
  - USB device pass-through
  - 3D graphics virtualization

Xen Research Roadmap Highlights

- VM Streaming
  - Migrating full VM state between machines efficiently with content-addressable storage network synchronization and logging
  - E.g. desktop to laptop to compute cloud “GoToMyVM” and back to laptop
  - Instant provisioning, disconnected operation, online backup
- Mobile phones and tablets
  - E.g. Xen ARM port by Samsung
  - Three VMs running on one CPU:
    - one for controlling the radio, one for vendor-supplied s/w, one for user downloaded software

Xen Research Roadmap Highlights

- Storage optimized for VMs
  - Supports high-rate snapshots for continuous data protection, high space efficiency
  - Advanced caching and re-layout optimizations
- Hardware Fault Tolerance for VMs
  - Near-instantaneous on-line failover between VMs on different servers
  - Continuous check-pointing vs. deterministic replay techniques

Conclusions

- Xen is becoming a key platform component, embedded in firmware
- The path to Ubiquitous Virtualization
- Xen Roadmap brings exciting new uses for virtualization
- Get Xen from http://xen.org
- (Or try XenServer Express)

Interested in a job at xen.org or Citrix? We're looking for great devs, sysadmins, techwriters etc. Email me!