

Virtualization for the rest of us

May 20, 2009

Tim Chipman – Fortech I.T. Solutions

<http://FortechITSolutions.ca>

Disclaimer & Copyright Note:

- Please be advised, this material is provided as a courtesy to attendees of the seminar.
- While every effort has been made to ensure all content is accurate, no assurance is made nor liability implied in the event of any inaccuracies.
- Please do not reproduce the content of these slides in any way.
- If you are interested in having a presentation, demonstration, or consultation on Virtualization, please don't hesitate to contact Tim.Chipman@FortechITSolutions.ca

Topics Overview

- (Very brief) History of Virtualization
- General approaches
- Applicability: When to do it? When not!?
- Benefits (and Drawbacks?) of Virtualization
- Compare & Contrast - ~6 virtual platforms
- Live demos – ‘welcome to the jungle’ and ‘key features’
- Q&A

This presentation is NOT:

- “The complete story”
- Vendor specific (hardware or software)

Presentation Goals:

- Interesting, interactive, informative, relevant
- Unclear ? Stop me and ask questions !

Why am I doing this ?

Independent I.T. Consulting Services:
Fortech I.T. Solutions


<http://FortechITSolutions.ca>

Info@FortechITSolutions.ca

(902) 442-6633

Definition: Server Virtualization

- One or more independent instances of Operating System(s) on a single piece of server hardware
- Achieved via some form of hardware abstraction
- All instances fully isolated from each other
 - Independent root / administrator access
 - Independent in event of crash, reboot, hack



A brief history

- Not exactly 'new concept':
- IBM Mainframes, 1970's
- What has changed:
 - Reduced 'barriers to entry'
 - Virtualization on 'commodity hardware'
 - Affordable bulk storage (iSCSI SAN, etc)
 - Multi-core CPUs and cheap RAM
 - Competition among virtualization solutions

Good external Reference: <http://en.wikipedia.org/wiki/Hypervisor>

General Approaches

- Early days: (pre-2005/06 VT Assist HW)
 - XEN: paravirtualization (guest OS tweaks)
 - VMWare (and others): 'clever tricks'
 - binary translation – some performance issues
 - "Container" based: (OS *not* HW virtualized)
 - Solaris 10 / 2005
 - Virtuozzo / OpenVZ (Windows or Linux / 2005)
 - ... and various others....

General Approaches.....

- Later Days: (VT Assist hardware, etc)
- AMD and Intel each have their own ..
- Gradually more HW assist features:
 - I/O, memory performance enhancements
- Full virtual machines:
 - VMWare 3.X, Citrix Xen 5.0, KVM
 - (VirtualIron, Sun, Oracle, Microsoft ...)

Applicability

- Good Fit: (*)
 - 'Modest (*)' CPU load; Disk & Network I/O reqs.
 - No special input-output / hardware needs
 - Video capture, serial port data/logger, etc

- Bad Fit: (*)
 - 'Extreme (*)' CPU workloads (esp. those benefiting from SMP >2-4 core)
 - 'Extreme (*) I/O requirements'

() Footnote: Definition varies with day of month, year, vendor, budget, etc etc etc. ☺*

Benefits of Virtualization

- Access to new features:
 - “Server Portability and independence”:
 - live migration, power management (“on demand” migrations), image/template based deployments, DR
 - Image Management: snapshots, (backup/rollback)
 - Very structured infrastructure, mgmt model
 - Capacity planning & growth:
 - Scalable / incremental for optimal \$CapX spending
 - Resource overbooking (CPU/Disk) easily possible
 - If well managed can facilitate density, capacity planning

Drawbacks to Virtualization(!?)

- New technology; learning curve; FUD
- “Increased complexity” and “more parts to break”
- Costs:
 - Many vendors: free loss-leader products and very (!) not-free full feature set for ‘enterprise’ deployments
 - Vendor lock-in / portability / issues ?

Why 'all the rage' ?

- ~Y2K 'server per service proliferation'
 - Many(!) under-utilized servers
- Significant burden on IT / Budgets
 - Acquisition, deployment, management
 - Infrastructure: HVAC/Power/Rack Space
- 10:1 consolidation of such equipment is 'no brainer'
- "Bonus" : access to "new feature set" is 'invaluable'

Compare and Contrast

- VMWare ESX 3.5i
- Citrix Xen 5.0
- Proxmox VE 1.0 (KVM/OpenVZ)
- Others:
 - Virtual Iron
 - MS-HyperV
 - Sun,Oracle (Xen core)...

	VMW	CtXen	Prox	V.Irn	HpV	X.oth
\$	Free, then \$2900 us\$ per socket	Free, then \$2600 us\$ per server	Free/Open Source	\$800 usd per socket	'Bundled with OS'	Free
Features	+++++	++++	++	++++	+?	+?+
Limitations	No \$, no Live Migrate	'seen as second best' ?	KVM: 1 cpu Windows VMs	Underdog, less well known	Stability? Mature product ? Features..?	Lifespan ?
Best fit ?	Less price sensitivity, "can't go wrong"	SMB / cost conscious	Linux appliance template deployment	mid to large deployment & \$ conscious	Windows only shop	Oracle / other only shop

Live Demos -



- Focus on 3 products :
 - VMWare ESX 3.5i:
 - Snapshots
 - Citrix Xen 5.0
 - Live Migration
 - ProxmoxVE
 - Template based deployment

For all three - will look at:

- “Typical virtual server access” illustration
- General management interface

Before we proceed ...


- Any questions ?
- Comments ?
- THANKS go out to those who helped make this presentation possible:
 - TARA
 - John McHugh - Canada Research Chair in Privacy and Security at Dal.
 - All the Virtualization Projects!
"Competition is good for the customer" ☺



(Material past this point is 'extra discussion' slides)

Live Migration – “The Trick”

- Requires that “virtual machine system images” reside on a ‘shared storage device’ (typically NFS, iSCSI SAN or Fibre SAN). (*)
- “Migration” of a VM between virtualization nodes - 3 steps:
 1. Management tools – establish who is source, target in the transaction, once sysadmin request “pulls the trigger”
 2. “Prep work” – copy running memory image of the VM; move this data via gig-ether mgmt network from source to target
 3. “Final move” – suspend VM momentarily, final diff and transfer of suspended memory image; resume VM instance on new physical host



(* This data *does not move at all*)


Emulation vs Virtualization

- Emulation: Complete virtual PC hardware emulated in software, host platform indep. (ie, run x86-Pc on 68030-Mac, Mac-on-PC, etc)

(figure: <http://www-jpc.physics.ox.ac.uk/Emulation.html>)

Issue: “massive performance hit”

- QEMU – x86 emulator written in C (ports to multiple platforms easily); can play a role as virtualizer (x86 on x86) and is used for low-level bios and related features for KVM virtualization



Virtualization vs Emulation

- Virtualization 'encapsulates' native commands for execution on native CPU; via paravirtualization (OS tweaks); binary translation (hypervisor tweaks) or Hardware VT-assist: avoids conflicts with hypervisor (bare metal OS) or other VMs
- Minimal performance hit over native bare OS when using VT hardware assist

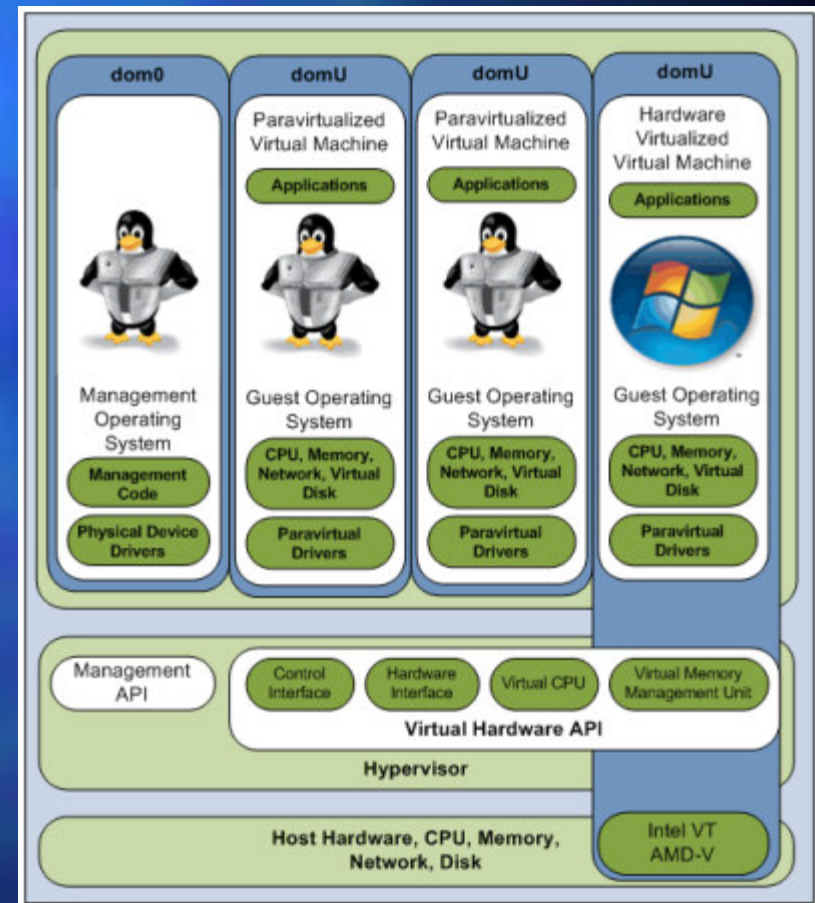


Figure Credit: Oracle

http://download.oracle.com/docs/cd/E11081_01/doc/doc.21/e10898/ovmsserver.htm

Comparing XEN and KVM


- Both are open source virtualization platforms, but with differing approaches to virtualization.
- XEN – ‘type 1 hypervisor’ (runs on bare metal): has been around longer (2003); more complex, larger code base
- KVM – ‘type 2 hypervisor’ (runs on top of host OS): newer solution (linux kernel module, 2007), simpler structure / smaller code base (each VM runs as a single user-space process)
- RHEL5.0: Bundled Xen; Looming 5.4 release – KVM (Redhat bought Qumranet – developers of KVM)
- Things get interesting: Libvirt – generalized VM management tool set – works with Xen, KVM (others...)
- “The future is diverse”

KVM, Xen – Links of interest ?

- Xen is Dead?
<http://tinyurl.com/ddus25>
- A bit more about KVM: (project homepage and wikipedia) and
<http://tinyurl.com/d93obv> and
<http://tinyurl.com/t4r4k>
- RHEL and KVM; RHEL-H release pending
<http://tinyurl.com/djhy8y>
(RHEL-H is "Redhat Hypervisor" – think ESX3i; 128mb ram footprint; boot from USB drive or via network)
- Brief history of Xen and XenSource:
<http://tinyurl.com/d7x7xx>


Example Virtualization Site 1

- Before: 10 x 1u single CPU P4-era physical servers, windows 2000/2003 (20 ghz / 20gig RAM / ~10% utilization profile)
- After: 1 x 1u physical server, dual-quadcore (20ghz, 16gig RAM / ~50% utilization overall)
- Benefits: Consolidation, reduced rack/power/heat; lower HW support contract costs
- New features: Snapshots/rollback
- Comments: No migration; no shared storage needed (NFS/iSCSI/SAN); needs decent fault-tolerant server (dual PSU, raid local storage) to be tenable




Example Virtualization Site 2

- Before: 10 x 1u single CPU P4-era physical servers, windows 2000/2003 (20 ghz / 20gig RAM / ~10% utilization profile)
- After: 2 x 1u physical server, quadcore (20ghz, 16gig RAM / 50% utilization overall)
- Benefits: Consolidation, reduced rack/power/heat; lower HW support contract costs
- New features: Snapshots/rollback; live migration
- Comments: Requires shared storage (NFS/iSCSI/SAN); needs decent fault-tolerant server (dual PSU, possibly raid local storage) to be optimal



Virtual Iron – Classic Deployment

- Summary of Components present:
 - 1 x management server (for netbooting Virtualization nodes)
 - 2 or more Virtualization nodes (no local storage / disks)
 - 1 or more iSCSI target raid disk array
 - Management, Public, and iSCSI subnet ethernet switches



Windows 2008 HyperV Notes

- Virtualization bundled free with Server 2008
- Product Maturity; Features & Stability issues ?
 - “Quick” (not “live”) migration ‘fiddley’ at best
 - Enterprise management – via SystemCentre Machine Manager (~\$900 USD)
- Windows Virtual Servers ‘price fun’:
 - Srv2008 Ent.Edn (\$4k USD): 4 free Win VM seats
 - Srv2008 DataCentre (\$3k USD/CPU) – unlimited windows VMs

Features Reference Links

- CitrixXen Feature Breakout:
<http://tinyurl.com/aal72>
- VMWare Features:
<http://tinyurl.com/5bh9km>
- ProxVE Features:
<http://pve.proxmox.com>
- Virtual Iron Feature Breakout:
<http://tinyurl.com/6zjrzu>
- MS 2008 HyperV Features:
<http://tinyurl.com/ck2sbz> and
<http://en.wikipedia.org/wiki/Hyper-V>

“Cloud Computing”

- Model: “Totally distributed” VM infrastructure
- Issues:
 - Vendor lock-in / interoperability ?
 - “Single point of trust / single point of failure”
(large consolidated data centres outside of your control)
 - Data persistence issues / Storage & access issues
 - Data security concerns (Clearly)
 - Best fit – for ‘web services’ servers which are inherently scalable in design and mesh well with Amazon cloud
 - Typically higher cost than straight co-location hosting (monthly rates); best fit is for VERY ‘peaky’ use
 - Server management / config / still an issue; just a perpetual rental cost (often CPU load / net traffic based) instead of a capX & facility operating cost.

VDI: Virtual Desktop Infrastructure

- New answer to old problem: “How to manage these pesky desktops with less sysadmin time / \$\$ / etc” ?
- Sells more virtualization licenses
- Provides dedicated abstracted remote desktop, template based deployments and lots of snazzy management features
- Why is this better than MS_Term Server? (High-demand desktop users?)