Virtualization (& paravirtualization), [Intel]

Background, Risks, Controls, Audit Steps

NEbraska CERT Conference 2006
Computer Security and Information Assurance

August 9, 2006

Michael Hoesing ***cissp, cisa, ccp*** *cia, cpa cma*

m-hoesing@cox.net

(402) 981-7747

Disclaimer, I never said THAT, if you heard THAT, it wasn't from me. None of the content of this presentation can be attributed to any of my employers, family members, acquaintances, conference sponsors, quail hunting partners past present or future.

Contents

- Drivers – why virtualize
- Practical Applications and History
- Definitions – virtualization, paravirtualiztion
- Tools – XEN, VMWare, MS & Recent News
- Architecture, Methodologies, Components
- Installation & Configuration (Xen)
- Risk, Controls & Audit (Xen)
- Demo, Resources Questions
- Hour 2 = VM ESX (howto, security, defaults)
- Hour 3 = SuSE 10.1 pro Xen “built-in”

Drivers (why are we even talking about this)

- Reduced TCO
  - 1 (or more) CPU can support many servers
  - 1 Storage Device & KVM can support many servers
  - less footprint (rent, utilities,..)
  - (generally no memory savings)
- Cheaper redundancy increasing continuity options
- Development testing
- Support
- Legacy application migration

Practical Applications

- Testing – run a version in a sandbox before deployment
- Testing – have multiple OS's and browsers and see how the website looks in different environments
- Academic – build a cheap network the students can take home on a disk
- any other cost saving opportunity

History

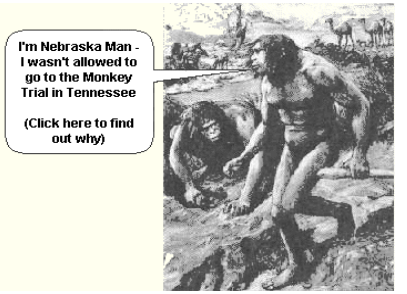

- one person, one machine life was good
- one person 2 machines (expensive)
- one person, one machine , dual boot (more choice, but only one choice at a time)
- (para)virtualization - many choices all available concurrently

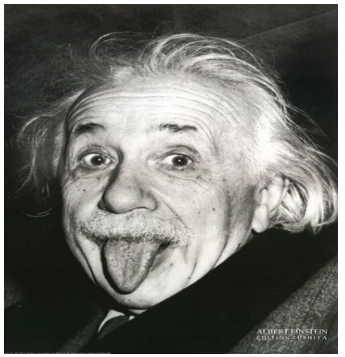

Workstation Versions

XEN $0 3.0.1 2/1/2006	**VMWare** $189 (workstation)	**Virtual PC** $0
Paravirtualization	Virtualization	Virtualization
Domain0	Host	Host
DomainUs	Guests	Guests
Kernel xen0 modified	Host kernel unmodified (sw layer)	Host kernel unmodified (sw layer)
Kernel xenU unpriviledged	Each guest unmodified	Each guest unmodified
No MS * (has been done VT)	MS and LINUX, hosts and guests	MS and OS/2 No Linux * (has been done)
Files,LVM or Partitions	Files	Files

Enterprise Versions

XEN $0 3.0.2 2/1/2006	**VMWare** $1,000 (ESX) (3.0 out soon)	**Virtual Server** $free standard R2 (+$1,000 OS)
Paravirtualization	Virtualization	Virtualization
Multi, dual core & VT in 3.0.2 , no max CPUs	**Dual processor minimum, dual core support,4 physical max**	**Single Processor Minimum,, 32 physical max**
64bit processor (since 3.0.0)	64 bit processor, dual core supported	64 bit processor
IDE minimum	SCSI minimum	IDE minimum
no max (PAE and SMP)	16 gig/guest 64 gig maximum	3.6 gig/guest 64 gig maximum
up to 32 domU's	8 guests optimal	??
Hot Move Guests (Enterprise $?,???)	Hot Move Guests & P2V $ More CPUs SMP add $2,750, HA Vmotion DR add $2,000	Hot Move Guests

Recent News

- Dec 05- VMWare Player, free, Free, FREE, you can not build a guest but you can run guests built by others
- Jan 06 – xen 3.0.1 , 64 bit processors, unlimited memory, dual core support
- Mar 06 – Intel VT (ring –1)
- May 06 – MS Virtual PC & Server free
- July 06 – VMWare Server 1.0 buil 28343 – FREE, old GSX, lots of tiered networking possibilities with 8 built-in switches, can be totally free if host is Linux distro,
- July 06 – VMWare Enterprise 3.0
- Soon – Xen Enterprise, mgmt console $?,???

Stolen from Ian Pratt of Cambridge & XenSource

Xen 2.0 Architecture

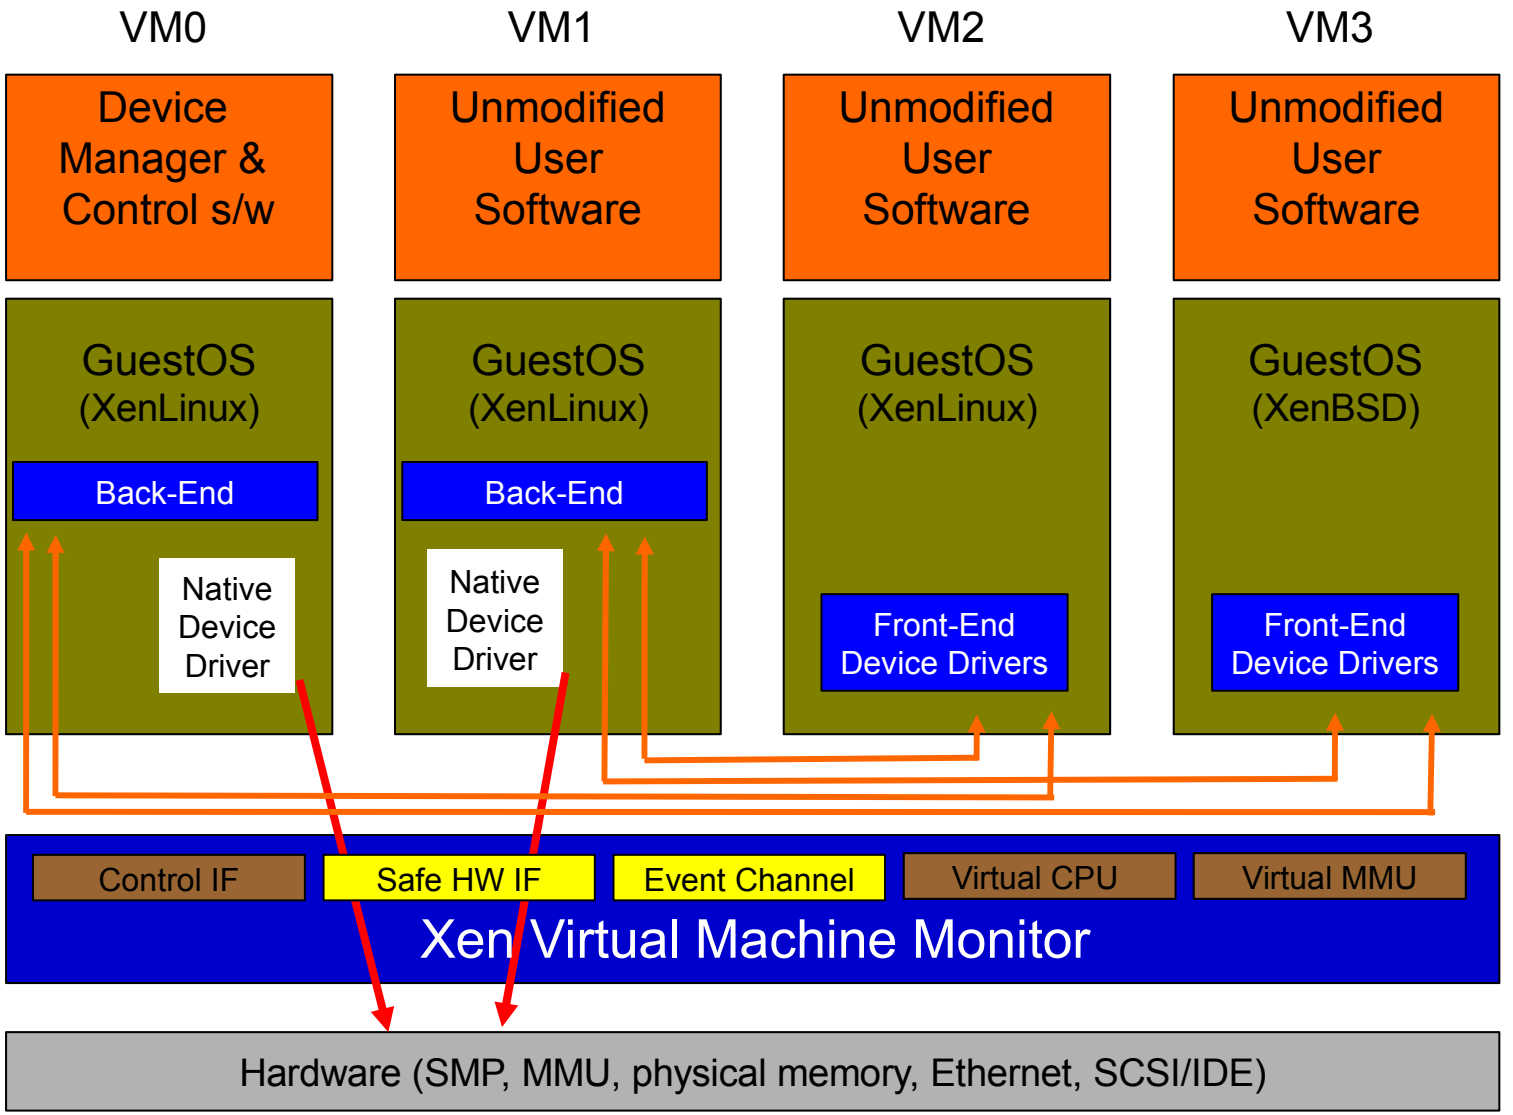

2 Methodologies

- Paravirtualization
  - Faster?
  - Altered kernel fulfilling requests rather than an app sitting on top of the kernel
  - User space applications need no modification
  - http://www.cl.cam.ac.uk/Research/SRG/netos
- Virtualization
  - Safer?
  - A software component sits between the guest OS and the host OS interpreting resource requests

Components

- VMWare + OS (MS or Linux)
- MS Virtual PC (runs on MS &OS/2 only)
- XEN (runs on Linux & netBSD only) [all can be free]
  - xen-x.x.x (paravirtualization tool)
  - twisted-x.x.x (networking framework [whatever that means])
  - linux -2.6.x.x (the kernel I virtualized)
  - bridge-utils (layer 2 protocol free bridging)
  - sysfs-utils (file system virtualization)
  - Zope-interface, iproute2, libcurl, zlib

XEN Installation

- http://lists.xensource.com/archives/html/xen-users
- (Richard Hamel-Smith 3.0 instructions) www.hpl.hp.com/techreports/2004/HPL-2004-207R1.pdf
- (Andreou and Walji sponsored by HP)
- http://lists.xensource.com/archives/html/xen-devel/2005-01/msg00434.html
- http://www.fedoraproject.org/wiki/FedoraXenQuickstart
- (Jeremy Katz)
- Plan and partition before hand
- Can use LVM or NFS also
- Can also live migrate

XEN Configuration

- Grub – sets xen0 memory, can also boot to unaltered kernel
- /etc/xen/xend-config.sxp xen config script
- /etc/xen/xmdomainname domain config script, memory, VIFs
- /etc/xen/xm commands, create, console, destroy
- /var/log/xend.log guess what
- /etc/xen/scripts network and vif-bridge scripts

XEN Risk, Control & Audit

- **RISK** - virtualization creates a single point of failure (dom0, host) for the guests
- restrict access to config files /etc/xen/
- restrict access to xend.log files
- check routes carefully, twisted and bridge-utils are powerful, can send packets anywhere
- Continuity – copy domains, have an extra machine (probably one of the ones retired)
- Fancy Audit Tools = ls –l, route, ps aux, cat

Risk Control & Audit - cont

- St_R0nG3r root password
- Use SUDO
- /etc/xen/xend-config-sxp
  - xend-address '' - any host can connect
- Check /etc/xen/auto for authorized domains at start-up
- Patch, harden (I have to say this every speech)

Risk Control & Audit – xmdomU config file

- memory = xxx (too small crashes, too big and other domains crash)
- vif = define virtual MAC numbers and assign them to bridges, duplicates cause problems
- disk = where to look for this domain's OS and apps, wrong pointer and things go bad
- extra = x this is the runlevel, why they call it extra beats me, avoid “0”

Risk, Control & Audit - /etc/xen/scripts

- network - builds bridges and VIFs at xend start
- network-route – sets /proc/sys/net/ipv4/ip_forward to “1”
- vif-route – sets interface routes up or down
- vif-bridge – associates vifs to bridges

Resources

- XEN modules and manuals
  - http://www.cl.cam.ac.uk/Research/SRG/netos/xen/
- XEN user group archives
  - http://lists.xensource.com/archives/html/xen-users/

Demo

- Show=
  - /boot/grub/menu.lst
  - /etc/xen/xend-config.sxp
  - /var/log/xend.log
  - /etc/xen/xm1firewall
  - /etc/sysconfig/network/ifcfg-eth0
  - /var/log/xend-debug.log
- Dom0 Ifconfig uname -a
- brctl show
- brctl showmacs xen-br0
- xm list xm create xm console xm list (again)
- vncserver (in domU) vncviewer (in dom0)
- DomU Ifconfig ping

OTHER

- Questions ??

VM ESX

- VMWare Security White Paper
  http://www.vmware.com/pdf/esx2_security.pdf
  - No public interfaces
  - Minimal host installation (apache in default install)
  - Guest isolation (using files)
  - AV & Firewall recommended (but not supplied)
  - Su to root
  - Default non-promiscuous NIC
  - Code was audited (scope & methodology not stated)
  - Use VLANs and place management console on separate vlan
  - Recommends disabling logging of VM messages in guest (?!)
  - Host OS is 100% VM, only drivers are open source
  - Management Console is from Red Hat 7.2
  - Users & Groups within VM mgmt console, home directory throttle

VM ESX (cont)

- VMWare ESX Other
  - Logical Access Control Provided at the OS level in addition to MUI users
  - Can overprovision memory , but throttle with wieghts called “shares”
  - (min host mem 192mg for 8 guests)
  - Watch routing, eth0 DHCP default install
  - /etc/vmware the goodies like **hwconfig** and **vm-list**
  - VMotion requires a SAN
  - Provide for swap or core dump on a separate partition
  - “bonded NICs” teamed interfaces, management access on the guest subnet through vmxnet_console
  - IBM blade:
    - USB CDROM won’t work on RDM installed guests
    - Bonded NIC failure of both, fix with Net.Zerospeedlinkdown 1

VM ESX (cont 2)

- VMWARE ESX More
  - **Console OS – host operating system**
  - **Service Console – administers host & guests, do not run X**
    - **VMWare Management Interface – http browser based controls the host and guests, 509 certificated, SSL, 90 second refresh window possible multi-user conflict, DOS possible with:**
      - **/usr/lib/vmware-mui/apache/conf/access.conf vmware_SESSION_LENGTH 0**
    - **API – HP Insight, Veritas,**
    - **SNMP – feed other tools**
    - **Remote Console – control the guest, MIME,**
    - **Check /proc/vmware for allowed methods**
  - **.vmx the guest configuration file /root/vmware/ , text editor can alter**
  - **.vmdk the guest image file VM MUI has a file manager**
  - **Admin manual suggests “flagship” user that is never on vacation**
  - **Install manual requires at least one non-root user**

VM ESX (cont 3)

- VMWARE ESX Still More
  - PXE Install – from a stored image, test then lock the image
  - Cannot downgrade from dual processor to single processor
  - LSI Logic SCSI adapter – see 30 pages of howto
  - VMware-console-2.x.x-xxxx.exe check authorized use
  - Reinstall VMware Tools overwrites the power level scripts
  - Move a vm, check the backup software
  - Dual CPU requires VMWare Virtual SMP
  - Backup from Service Console requires guest shutdown

VM ESX (cont 4)

- More more
  - No USB on Guest (2 factor impact?)
  - NT can only run on a single processor machine
  - Guest event log , user is not indentified
  - /etc/pam.d/vmware-authd
  - /etc/vmware-mui/ssl/mui.crt and mui.key
  - Security Config:
    - Medium – mgmt and remote encrypted, telnet & FTP are not encrypted
    - Low – no connections to host are encrypted
    - Custom -

VM ESX (cont 5)

- More again
  - VMFS 2.11 file system, public shared
  - Physical extent aka partition
  - SPAN joins across partitions creating a volume, first “span” formats thus wiping out existing data
  - Logs /var/log/vmkernel and vmkwarning
  - /etc/snmp/snmpd.conf trapcommunity public (rename this)
  - vmkload_mod –l to list loaded modules
  - /etc/vmware/hwconfig and vmkmodule.conf

VM ESX (cont 6)

- More stuff
  - LUN masking, only allow guests to see what they need
  - vmkmultipath -q where the data goes

VM ESX 2.5 Default Installation

- LILO without a password
- MOTD empty, no login banner
- gopher, news, mail, finger, ftp, samba 2.2.7, telnet 0.17
- login as root , su not required
- 2.4.6 kernel 3/17/05 last update
- cracklib present, but no pword strength enforcement
- /proc/sys/net/ipv4/conf/all/accept_redirects 1
- ports 902 8222 8333

SuSE 10.1 Xen “Built-in”

- Partition the drive first, guests will be installed in in extended partitions hda5, hda6, hda7, in YAST make the mount points data1,2,3 they will be built into fstab in dom0
- Disable the autostart of SuSE firewall
- Xen is on the distribution media, but not part of the standard installation, use YAST2 check the box for Xen
  - Xen-kernel, xen-kernel-nongpl, xen, 2 UML files
  - 3 doc howto files that ……
- Re-uses the xen kernel for both dom0 and

SuSE 10.1 Xen “Built-in” (2)

- Reboot, the normal build will mount mount data1
- Yast2, Software, “install into directory for xen”
- Yast2 VMM, don’t config domU’s with this
- Select /data1 as the guest target directory, do not install “image”
- Use distribution DVD media
- Select the 6 xen packages to install in the guest target directory also (do not select tomcat5)
- Select other SW, accept, wait, exit YAST

SuSE 10.1 Xen “Built-in” (3)

- While /data1 is still mounted
  - Edit dom0 /etc/fstab, comment out the data1,2,3 drives, then copy to /data1/etc/fstab
  - edit /data1/etc/fstab so the boot drive is /dev/hda1 (not /dev/hda5, because this will be logically re-mapped in the xm<yourname> start file)
  - Copy the 6 security files, both normal and YAST2 versions (password, shadow, groups) to /data1/etc/ (the xen install forgets to ask for a root password)
  - /etc/sysconfig/network change FORCE PRESISTENT HAMES to “no”
  - mv /data1/lib/tls /data1/lib/tls.disabled and mv /lib/tls /lib/tls.disabled
  - Change /data1/etc/HOSTNAME, motd, bashrc.local, copy wallpaper

SuSE 10.1 Xen “Built-in” (4)

- Copy my appendix A vm1 start and edit
  - Change the guest name, nics (if more than 1), memory
  - Bootloader /usr/lib/xen/boot/domUloader.py
  - Bootentry /dev/had: /boot/vmlinuz-xen,/boot/initrd-xen
    - Single kernel update (great improvement)
  - Edit vif(s) to assign a static mac to a virtual bridge
  - Map real partition to /hda1 disk = [‘phy:hda5,hda1,w’]

SuSE 10.1 Xen “Built-in” (5)

- Reboot into SuSE Xen
- xm create /etc/xen/xm<yourname> xm list
- xm console YourName
- Root, password, vncserver (note the TTY number)
- On another machine: vncviewer ip:tty , kde , dostuff

Fedora Core 5 Xen “Built-in”

- LVM method
- Yum install kernel-xen0 (dom0 auto build)
- Local copy of FC5 iso mounted as loopback /pub/ftp
- xm console YourName
- domU install script (way too easy)
  - **/usr/sbin/xenguest-install.py – n DomainU1 **
    **-f /dev/VolGroup00/LogVolDomU1 **
    -r 256 –l ftp://xxx.xxx.xxx.xxx/pub/