



# **Virtualization: From the Desktop to the Enterprise**

**Chris Wolf**

**Senior Analyst**

**Burton Group**



# Hazard Ahead!

## Successfully Navigating the Waters of Enterprise Virtualization

*Chris Wolf*




# Agenda

Here's what you need to worry about

- Today's performance issues
- Integration with storage and data protection
- Pending technological advancements that impact virtualization planning today
- Aligning the right technology with the right problem

There are plenty of potential bottlenecks today

- CPU
- Memory
- Network
- Storage
- Host Applications



- Naturally, the more the merrier
- Be careful with Virtual SMP
  - Only useful with multithreaded applications
- Architectural
  - By default, operating systems want low level (Ring 0) access to CPU
  - Not possible with full virtualization
  - Solutions:
    - Hardware-assisted virtualization (Intel, AMD)
    - Paravirtualization (Linux)
- CPU Quality of Service (QoS) helps
- Beware CPU spikes (backup, antivirus scans, etc.)



## Memory Sizing

- Sizing Guidelines

- VMware ESX
  - Physical Host: 24MB + either 192MB (< 8 VMs) or 272MB (> 8 VMs)
  - VM: 79MB overhead + RAM allocated to guest OS
  - Example: 4 VMs with 512MB RAM requires 2,580MB
- MS Virtual Server
  - Physical host: 512MB
  - VM: 32MB overhead + RAM allocated to guest OS
  - Example: 4 VMs with 512MB RAM requires 2,688 MB
- SWsoft Virtuozzo: Overhead < 30MB per virtual environment
- XenSource: 1GB + VM guest OS requirements
- Virtual Iron
  - Physical Host: 300MB
  - Page Table Overhead: 5% of installed physical RAM
  - VM Overhead: 10% of guest OS memory




## Performance Considerations

- Memory balloon driver
  - Unused physical memory paged to physical host hard disk
  - Allows you to allocate enough memory to each VM to handle performance spikes
- Adding memory to running VMs
  - Server virtualization: VM must be taken offline
  - OS virtualization: No downtime required
- Memory virtualization impacted by client load

## The cost of virtualization

- Today – Shadow Page Tables (SPT)
  - VMM maintains separate SPTs for each VM
  - VM memory access is captured and translated
  - Performance hit: 25-75%
  - Memory access time measured in nanoseconds (1 billionth of a second)
  - Noticeable with server virtualization in high performance database, web, file serving applications (not an issue with OS virtualization)
- Tomorrow (by 2008)
  - AMD: Nested Page Tables (NPT)
  - Intel: Extended Page Tables (EPT)
  - VMs will modify own physical page tables and handle own page faults
  - Architectural memory bottleneck removed!

- Physical (raw) disk
  - VM mounts directly to FC, iSCSI, or SCSI disk
  - I/O may require translation via a paravirtualized driver
  - FC disk may appear to VM's guest OS as a SCSI disk
    - NPIV changing this
  - Excellent integration with storage virtualization and third party data management tools
- Virtual disk
  - Disk resides in a file or collection of files
  - Fully portable
  - Requires controller emulation or paravirtualization



- VMware ESX Server

- Varies (1-20%)

- IEEE 2003

[http://www.vmware.com/pdf/wwc\\_performance.pdf](http://www.vmware.com/pdf/wwc_performance.pdf)

- Microsoft Virtual Server 2005 R2

- 28% average latency

- CapitalHead.com

<http://capitalhead.com/articles/benchmarking-microsoft-virtual-server-2005.aspx>

- Xen:

- < 5% (average)

- <http://www.clarkson.edu/class/cs644/xen/files/repeatedxen-usenix04.pdf>

- Antivirus
  - Filter virtual disk files from antivirus scans
- Use fixed size virtual hard disks (prevents fragmentation on host)
- Run virtual hard disks on separate physical disks (than host OS)
- Consider storage virtualization integration
- Redundancy is critical
  - RAID
  - Multipath I/O



- Types:

- Virtual Host Only Network

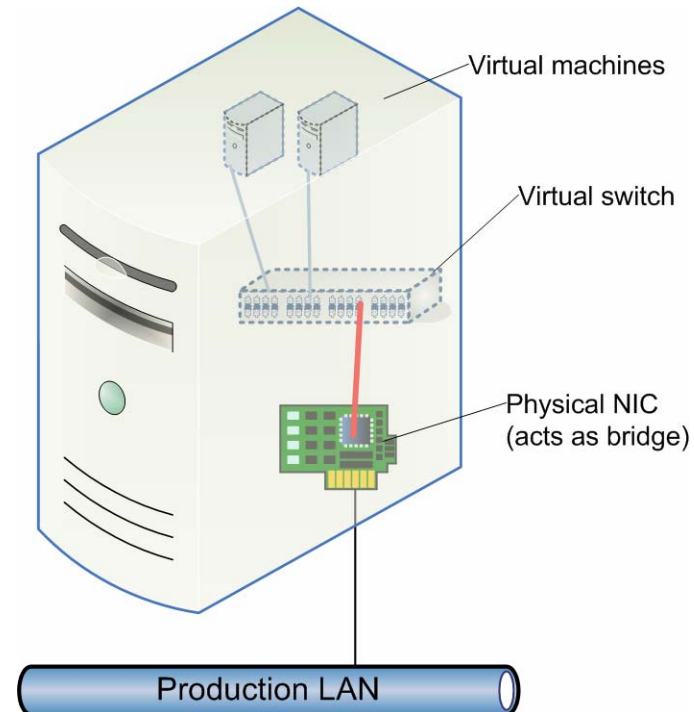
- VMs interconnect using a virtual network on the host
    - Host would need to route or NAT VMs to production network

- Bridged Network

- VMs directly connect to host's physical NIC

- Bridged Network Optimization

- Install multiple NICs on physical host
  - Bridge VMs to physical NICs based on average network utilization





# Network Performance Comparison


13

- VMware ESX 3: **2%** latency  
[http://www.vmware.com/pdf/hypervisor\\_performance.pdf](http://www.vmware.com/pdf/hypervisor_performance.pdf)
- MS Virtual Server 2005: **8-20%** (CW est.)  
No published benchmarks
- Xen: **2%** latency  
[http://blogs.xensource.com/simon/wp-content/uploads/2007/03/hypervisor\\_performance\\_comparison\\_1\\_0\\_3\\_no\\_esx-data.pdf](http://blogs.xensource.com/simon/wp-content/uploads/2007/03/hypervisor_performance_comparison_1_0_3_no_esx-data.pdf)
- SWsoft Virtuozzo: **Native**

- Disable unused virtual ports & devices
  - USB
  - Floppy
  - COM
- Disable unnecessary programs and services
- Run a single application per VM
- Install VM Tools/Additions
- Video
  - Do not run X Windows
  - Disable screen savers
- Antivirus/Backup should only run during off business hours



- Hot Backups vs. Cold Backups
- Virtual Disks vs. Physical Disks
- Local SAN Integration



## **Backs up a “live” VM**

- Most common approach - Install backup agent(s) on VM
- Alternate Methods
  - Serverless Backup
    - Options:
      - VMware Consolidated Backup
        - File level backup support only available on Windows guest OS
      - SAN-based serverless backup
        - Requires third party “advanced” or “proxy” agents
        - VMs must map directly to raw storage
    - Continuous Data Protection
      - Double-take
      - Neverfail
  - N\_Port ID Virtualization (NPIV) Improving Backup Options
    - VM gets own WWPN via a virtual fibre channel HBA



# Cold Backups

## VM placed in offline state prior to backup

- Backs up VM disk images
  - Granularity achieved by mounting disk images on alternate system and running file level backup
- Scriptable in both Microsoft and VMware
- Common approach:
  - Use a script to suspend VM(s)
  - Create snapshot or back up VM data
  - Use a script to resume VM(s)

- Back up the physical host system
  - Host system contains all VM configuration files
  - Required for DR
  - Files backed up by hot or cold VM backups should be filtered from host backups
    - Example: Create filters for .vmdk or .vhd files
- Verify hardware compatibility
  - Hypervisor-based solutions (VMware ESX, XenSource, Virtual Iron)
    - Manual installation of 3<sup>rd</sup> party drivers not supported
  - OS-based server virtualization (Microsoft, Novell)
    - Installation of 3<sup>rd</sup> party drivers supported
  - OS virtualization (Virtuozzo, Solaris Containers)
    - Installation of 3<sup>rd</sup> party drivers supported
- Purchase fibre channel HBAs that support NPIV

- Dynamic failover is essential in mission critical environments
  - Live migration is not dynamic failover
- Virtualization technology selection
  - Server virtualization
    - Legacy consolidation, dedicated application servers, simultaneous connection range (tens – hundreds)
  - OS virtualization
    - Performance-intensive apps (web, file, database), simultaneous connection range (hundreds – thousands)
- Explore dynamic hardware provisioning solutions
  - Power down/repurpose media servers when they are not needed
- Secure backup copies of VM images at time of deployment



# Additional Recommendations

20

- Use live migration to accommodate performance spikes
- Plan backups from a restore perspective
- Verify infrastructure compatibility
  - Physical devices (NICs, storage controllers, etc.)
  - Storage devices
  - Enterprise management and data protection software
- Requisition hardware platforms that support hardware-assisted virtualization
- Redundant storage and network I/O paths is a requirement
- Stay involved in any migration project
  - You know your user and server tendencies better than anyone!