Performance & Scalability Testing in Virtual Environment

Hemant Gaidhani, Senior Technical Marketing Manager, VMware

vmware

About the Speaker

Hemant Gaidhani
- Senior Technical Marketing Manager, VMware
- Capacity Management, Mastermind
- Author, “Virtualizing Microsoft Tier-1 Apps with VMware vSphere 4

vmware

Performance & Scalability in Virtual Environment

Understanding Virtualization Performance

Performance Benchmarking in Virtual Environment

Performance Monitoring in Virtual Environment

Virtual better than Physical?

Real Life Examples

Physical to Virtual Comparison

- Physical machine:
  - 2 socket, dual core server
  - 2 GB RAM
- Virtual machine: (default settings)
  - 1 virtual CPU
  - 256 MB RAM

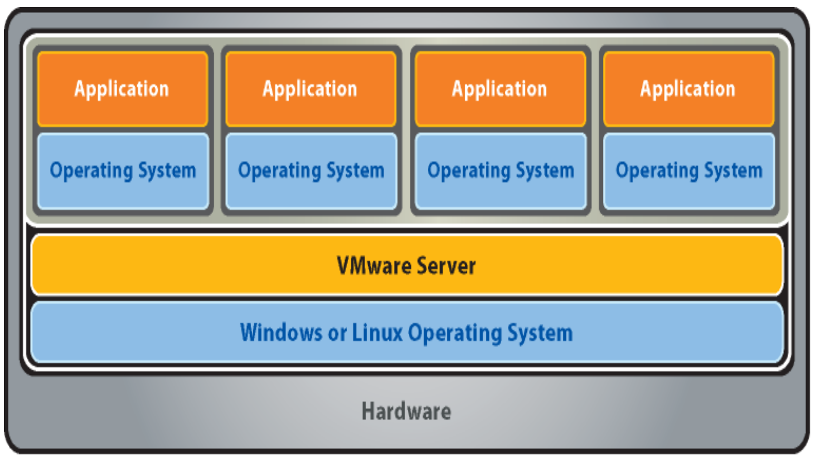

Comparing different virtualization products

- VMware Server
  - Hosted architecture
- VMware ESX Server
  - Bare-metal architecture (hypervisor)

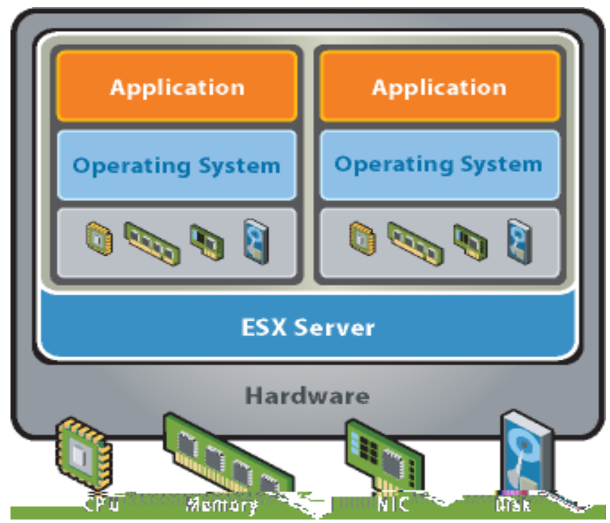

Real Life Examples (Continued)

Server Consolidation or VDI benchmarks

- Multiple under-utilized physical machines migrated to virtual machines
- Over-commitment of resources as server utilization goes up
  - CPU
  - Memory
- In over-committed environments,
  - virtual machines can start to starve each other
  - performance at the cost of another virtual machine

Performance in Virtualized Environment

Two dimensions

- Vertical performance – monitor efficiency (VMM)
  - Single VM performance - monitor overhead
- Horizontal performance – scheduling (VMKernel)
  - Multi VM performance - Scheduling overhead - often overlooked

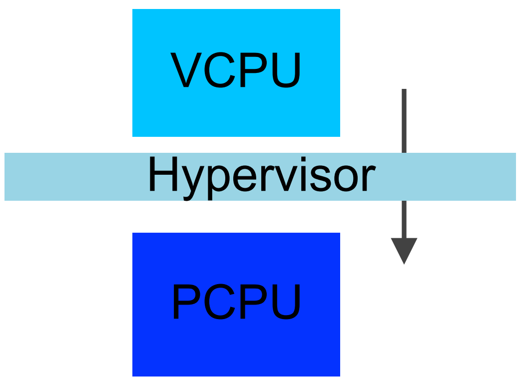

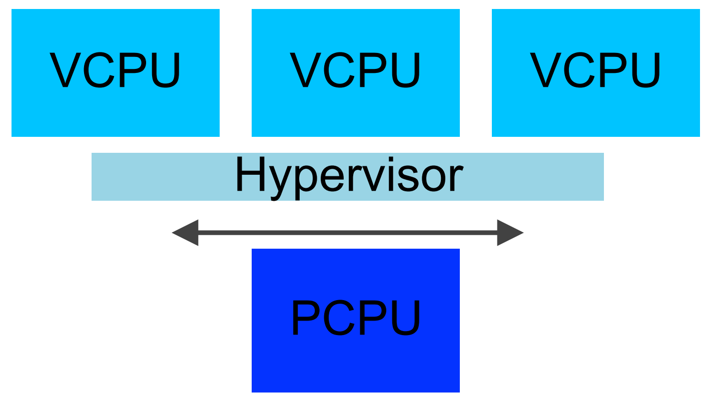

Horizontal Performance

Some examples

- Ready time
  - VMs don’t get to run all the time even if they want to
- Migration/Switching overhead
  - VMs will have to be context switched out for performing critical tasks or for scheduling another VM
- CPU cache contention
  - Multiple VMs share the same physical CPU
- NUMA effects
  - VM could be temporarily migrated to remote Node

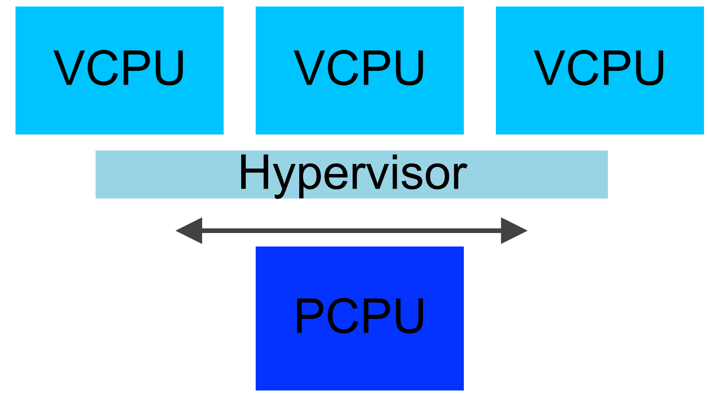

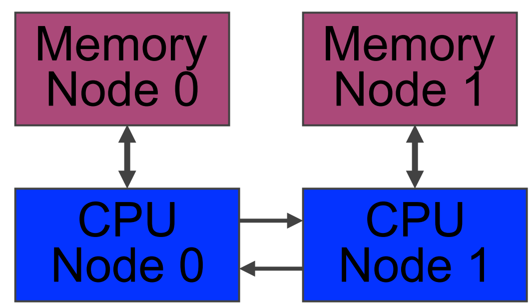

SQL Server Scale-Out Overcommittment Fairness

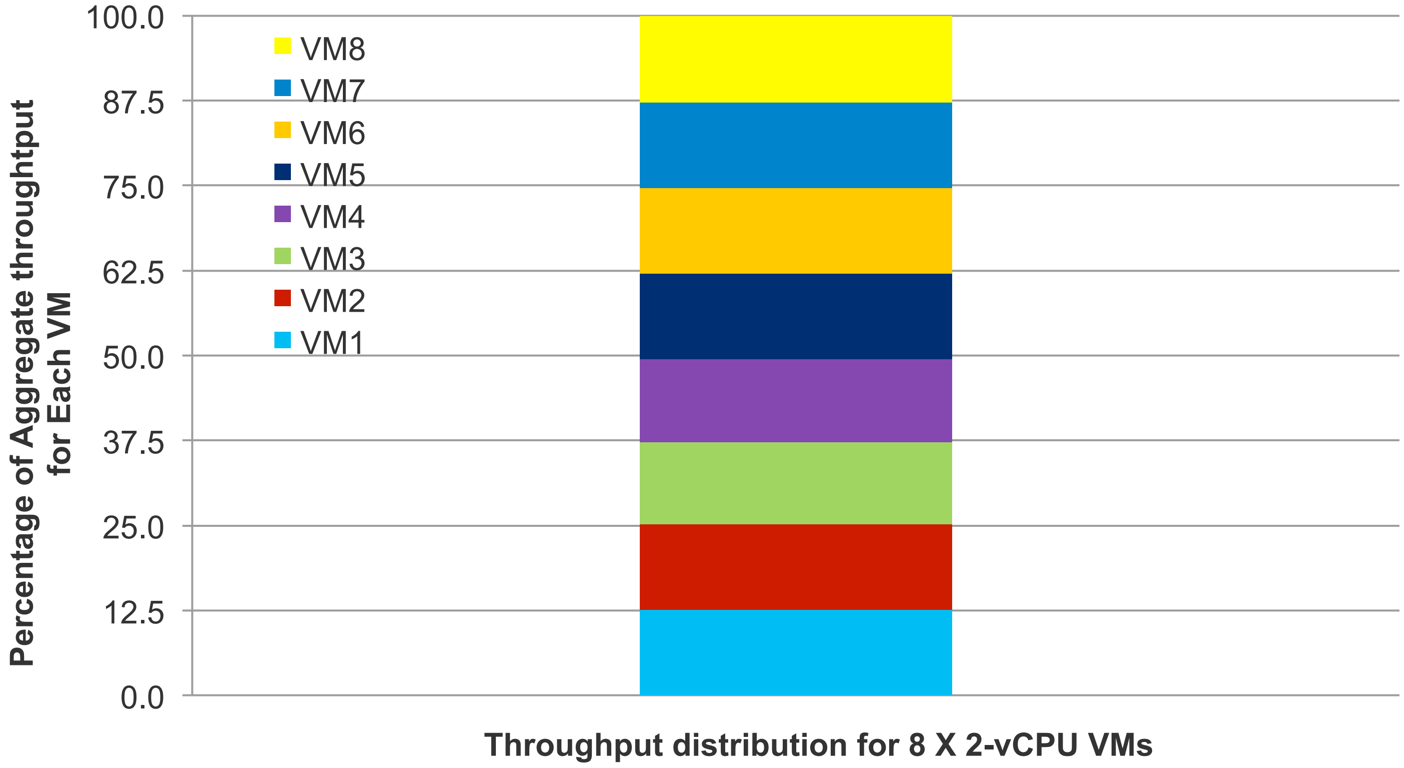

- Fair distribution of resources to all eight VMs

Third Dimension

Distributed Resource Scheduling (DRS)

- Reduces the impact of horizontal performance issues
- However aggressive migration impacts performance

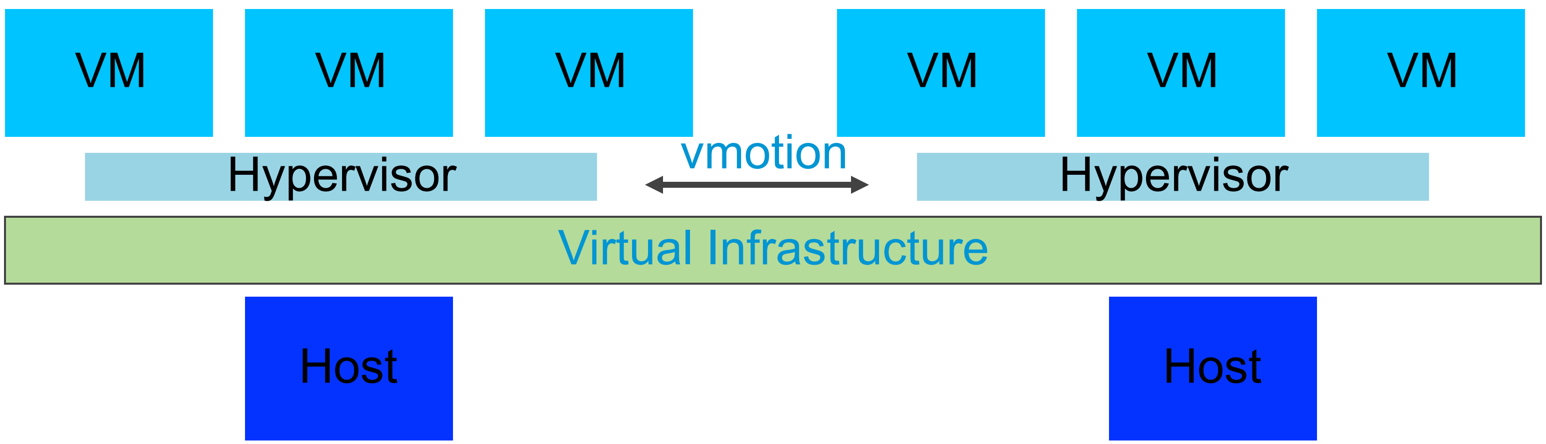

vSphere Scaling Trends

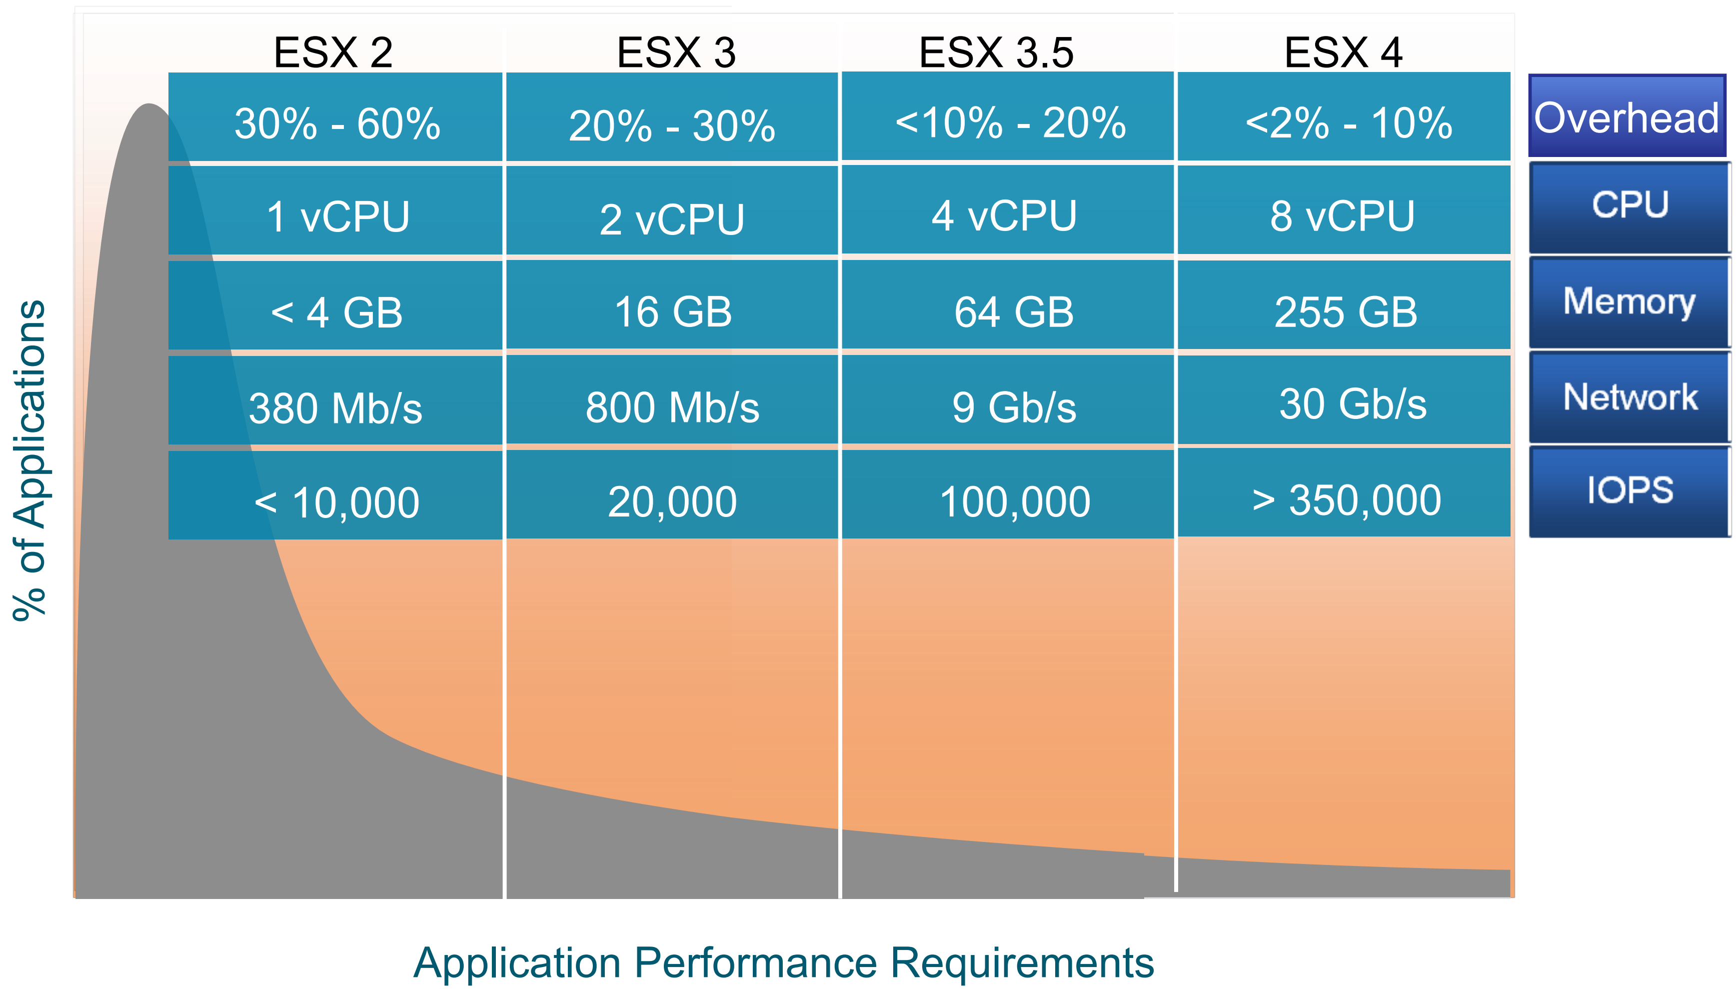

Virtualization Architectures

Hosted architecture

- VMware Server/ Workstation/ Fusion

Bare-metal architecture (hypervisor)

- VMware ESX Server
- VMware vSphere Hypervisor – free version

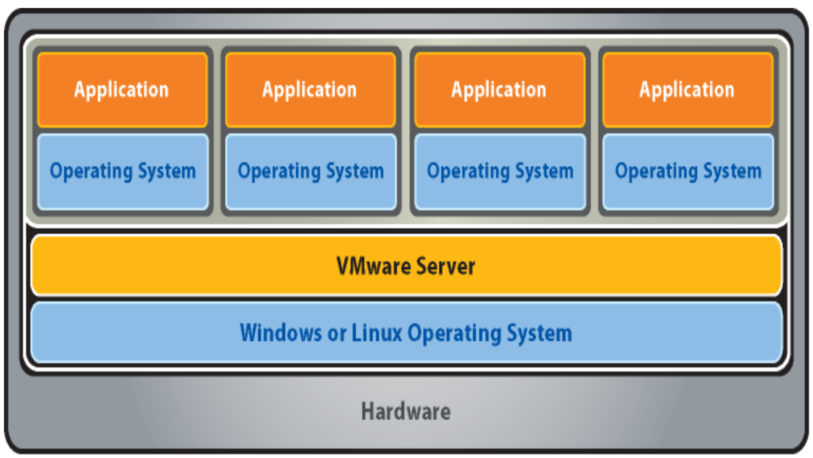

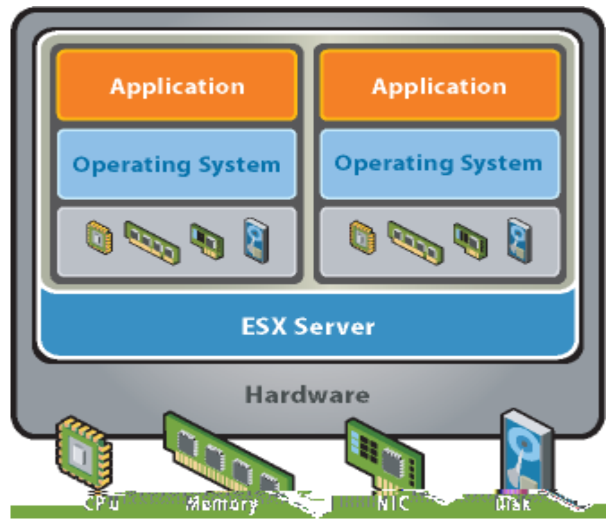

VMware ESX Architecture

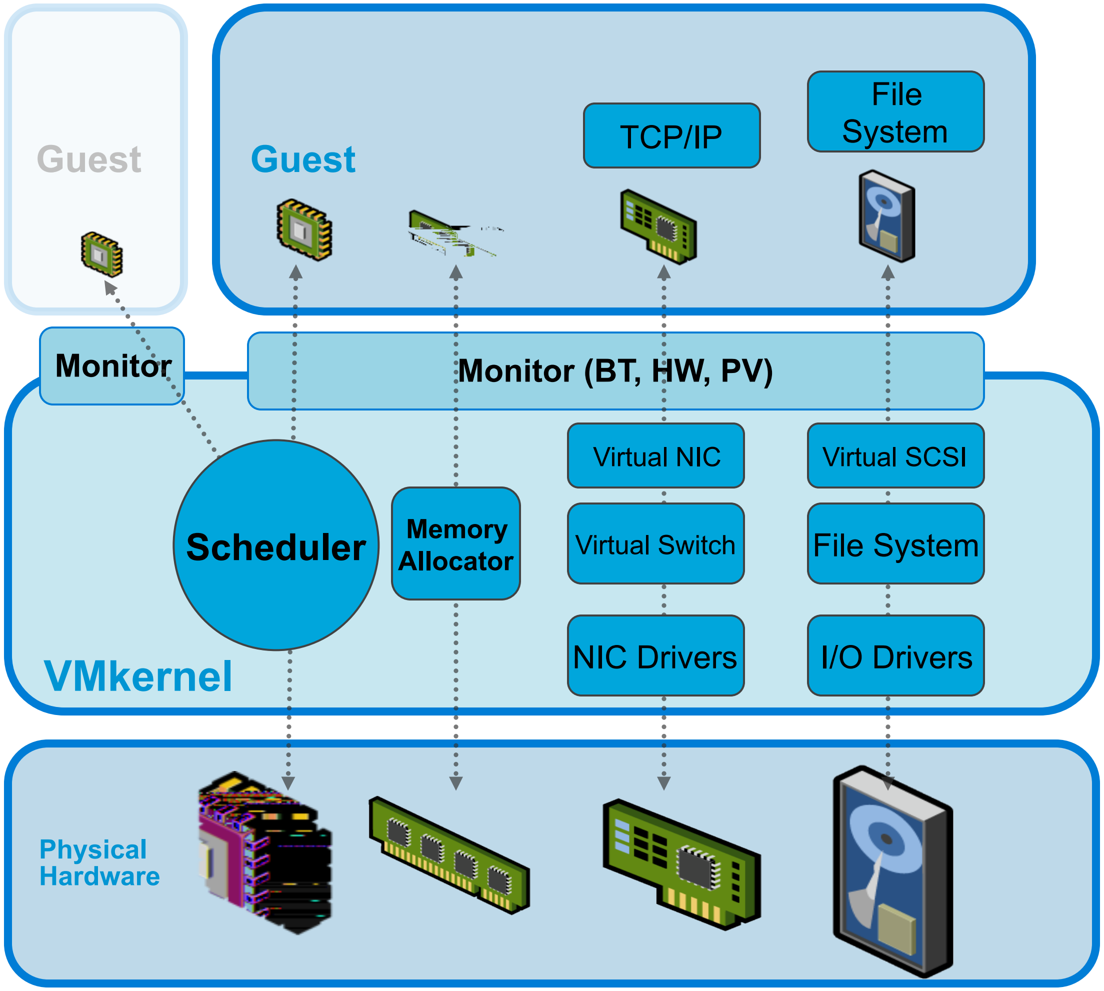

CPU is controlled by scheduler and virtualized by monitor

Monitor supports:
- BT (Binary Translation)
- HW (Hardware assist)
- PV (Paravirtualization)

Memory is allocated by the VMkernel and virtualized by the monitor

Network and I/O devices are emulated and proxied though native device drivers

HW Assist CPU Architecture

- **Intel**
  - Intel VT-x
  - Available since: 2006

- **AMD**
  - AMD-V
  - Available since: 2006

- **Prior to these hardware technologies – "binary translation" was used.**
- **First implementations of hardware assist were slower than binary translation in software.**

Hardware Assist (HWmmu) Memory Architectures

- **Intel**
  - Extended Page Tables (EPT)
  - Available since: 2009
  - Supported in ESX4.0 +
  - Nehalem or better

- **AMD**
  - Rapid Virtualization Indexing (RVI)
  - Available since: 2008
  - Supported in ESX3.5 +
  - Shanghai or better

- **Prior to these hardware technologies – "shadow paging" (or SWmmu) was used.**

vSphere Unlocks Processing Cores for Applications

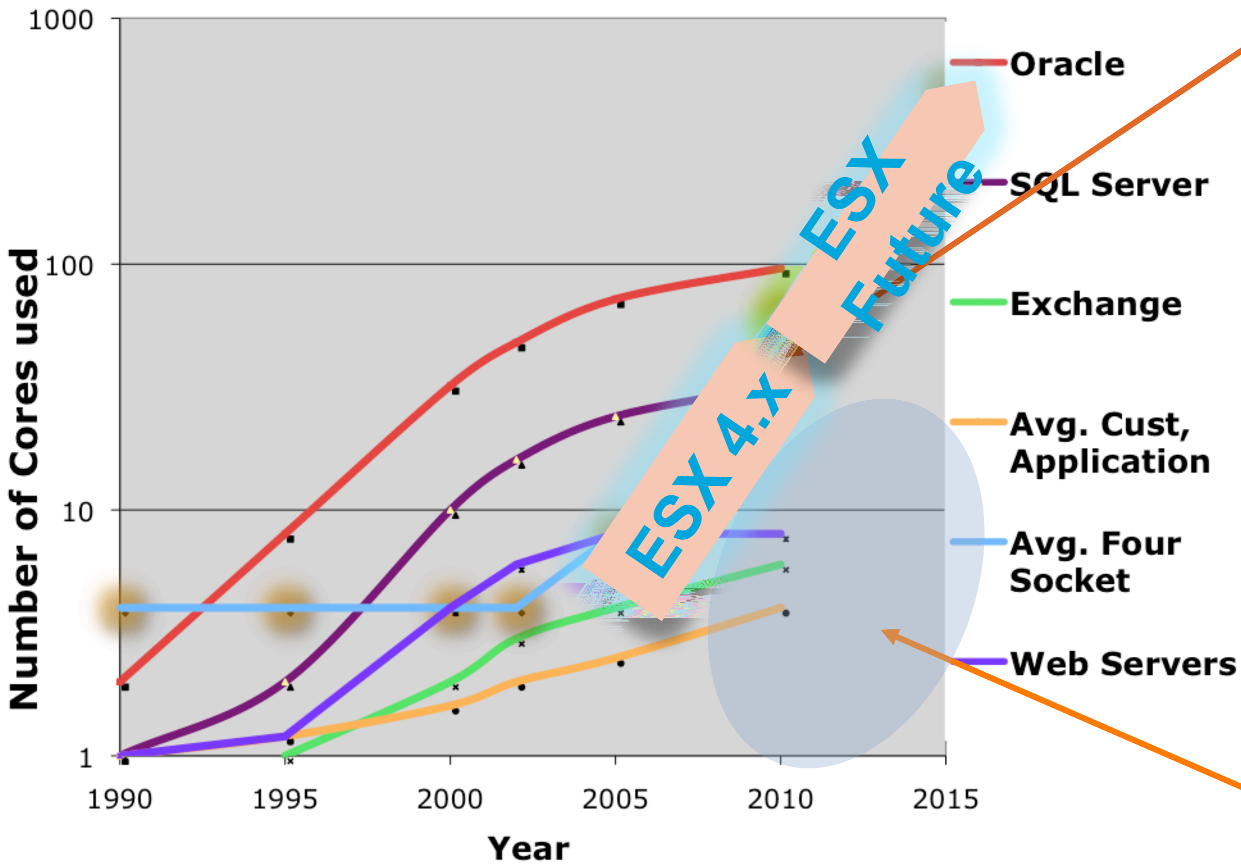

VMware ESX Scaling: Keeping up with core counts

Virtualization provides a means to exploit the hardware's increasing parallelism

Most applications don't scale beyond 4/8 way

Agenda

Understanding Virtualization Performance

Performance Benchmarking in Virtual Environment

Performance Monitoring in Virtual Environment

Virtual Better than Physical?

Performance Benchmarking in Virtual Environment

The standard performance benchmarking guidelines still hold true

Out-of-the-box performance optimization for vSphere

- No tuning required

Virtualization overhead depends upon the workload

Virtualization

- does not create new resources
- enables more efficient utilization of existing resources

Cannot break the laws of physics

Virtual Environment Implications

Guest OS metrics

- Performance metrics in the guest could be skewed when the rate of progress of time is skewed
- Guest OS resource availability can give incorrect picture

Resource availability

- Resources are shared, hypervisors control the allocation
- Virtual machines may not get all the hardware resources

Performance Profiling

- Hardware performance counters are not virtualized
- Applications cannot use hardware performance counters for performance profiling in the guest

Virtualization moves performance measurement and management to the hypervisor layer

Benchmarking Options

Benchmarking Methodology

- Fair and consistent
- Avoid “Apples to Oranges” comparison

Careful selection of benchmarks

- Several public and custom benchmarks available

Best benchmark is the one that simulates your workload

Some benchmarks may provide inconsistent, unpredictable results in virtual environments

Typical physical benchmarks are single workload benchmark

- Does not directly apply to virtual environment
- Consider VMmark: A Scalable Benchmark for Virtualized Systems

Focus on Storage

Storage Virtualization Concepts

- **Storage array** – consists of physical disks that are presented as logical disks (storage array volumes or LUNs) to the ESX Server.

> **Storage array LUNs** – formatted as VMFS volumes.

> **Virtual disks** – presented to the guest OS; can be partitioned and used in guest file systems.

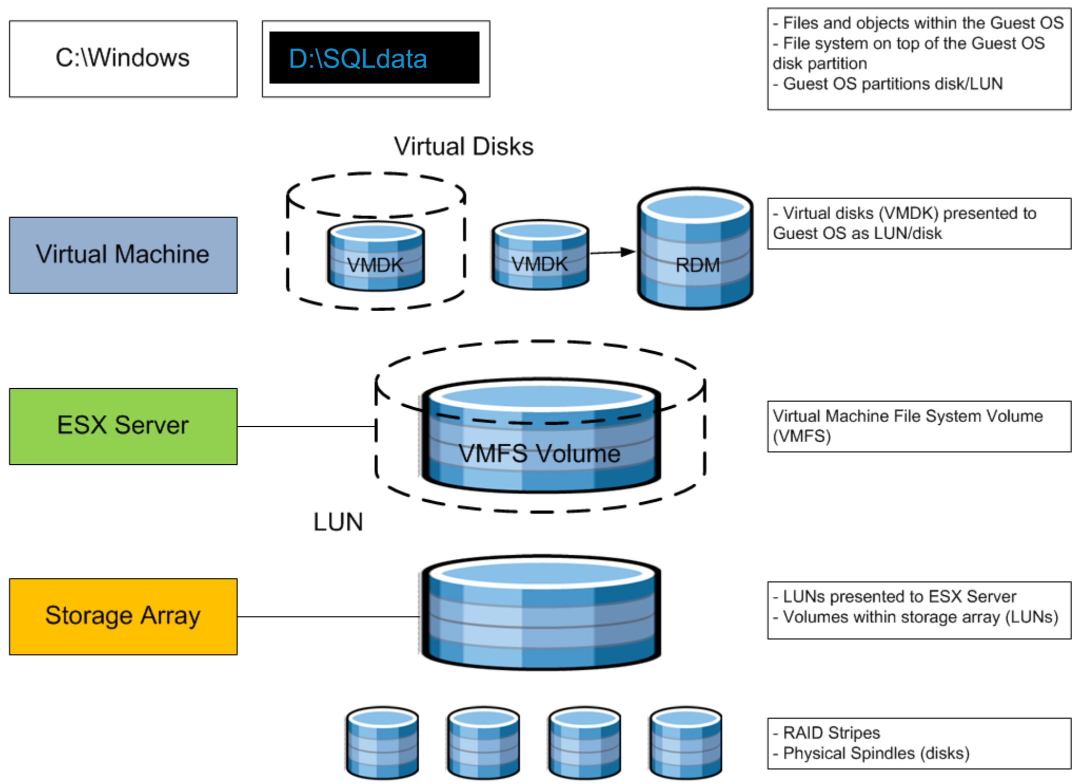

Server Consolidation: Storage Planning

Physical setup: each instance provided 5-spindle LUN

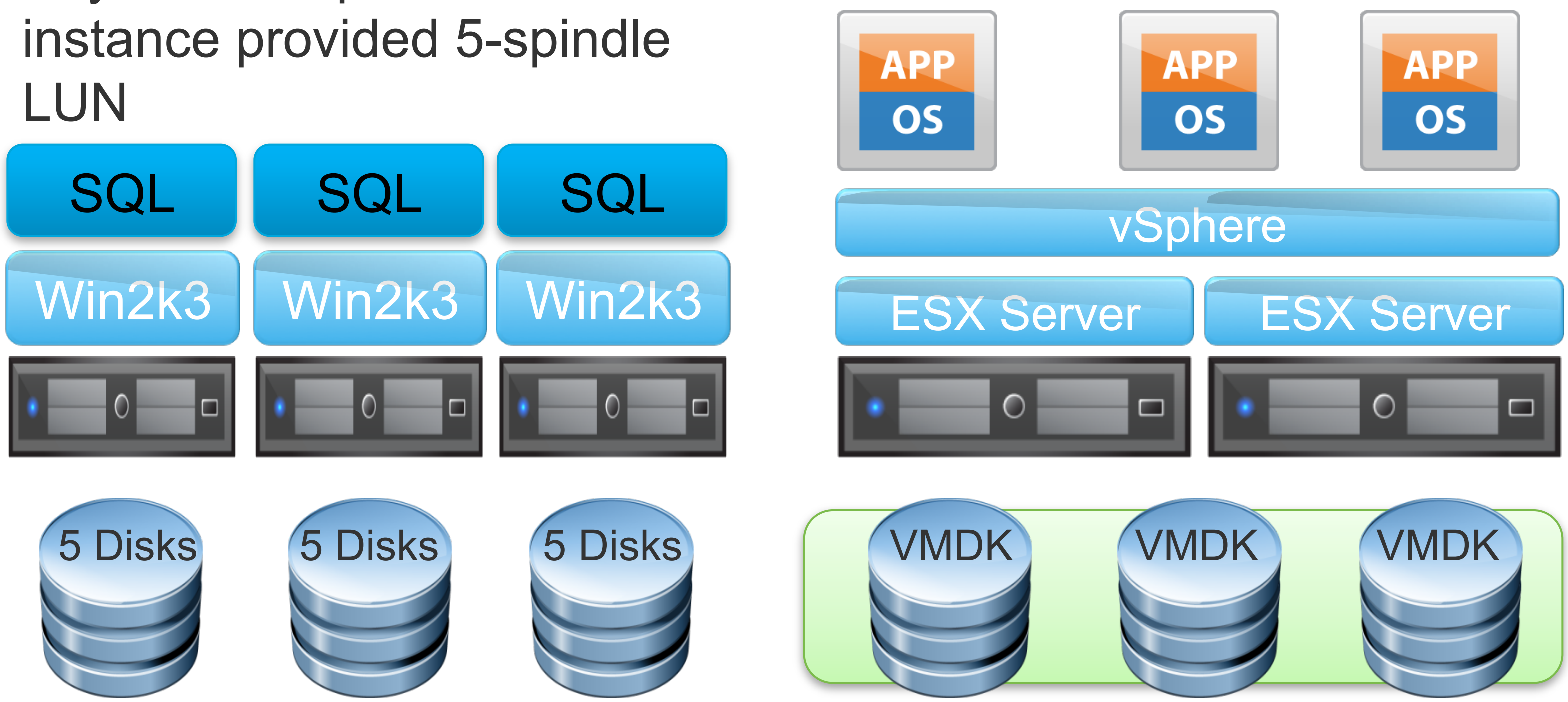

Virtual architecture: Each VM provided its own VMDK

•But now do they map to disks?

Sequential Workloads Generate Random Access

As observed in VMFS scalability tests

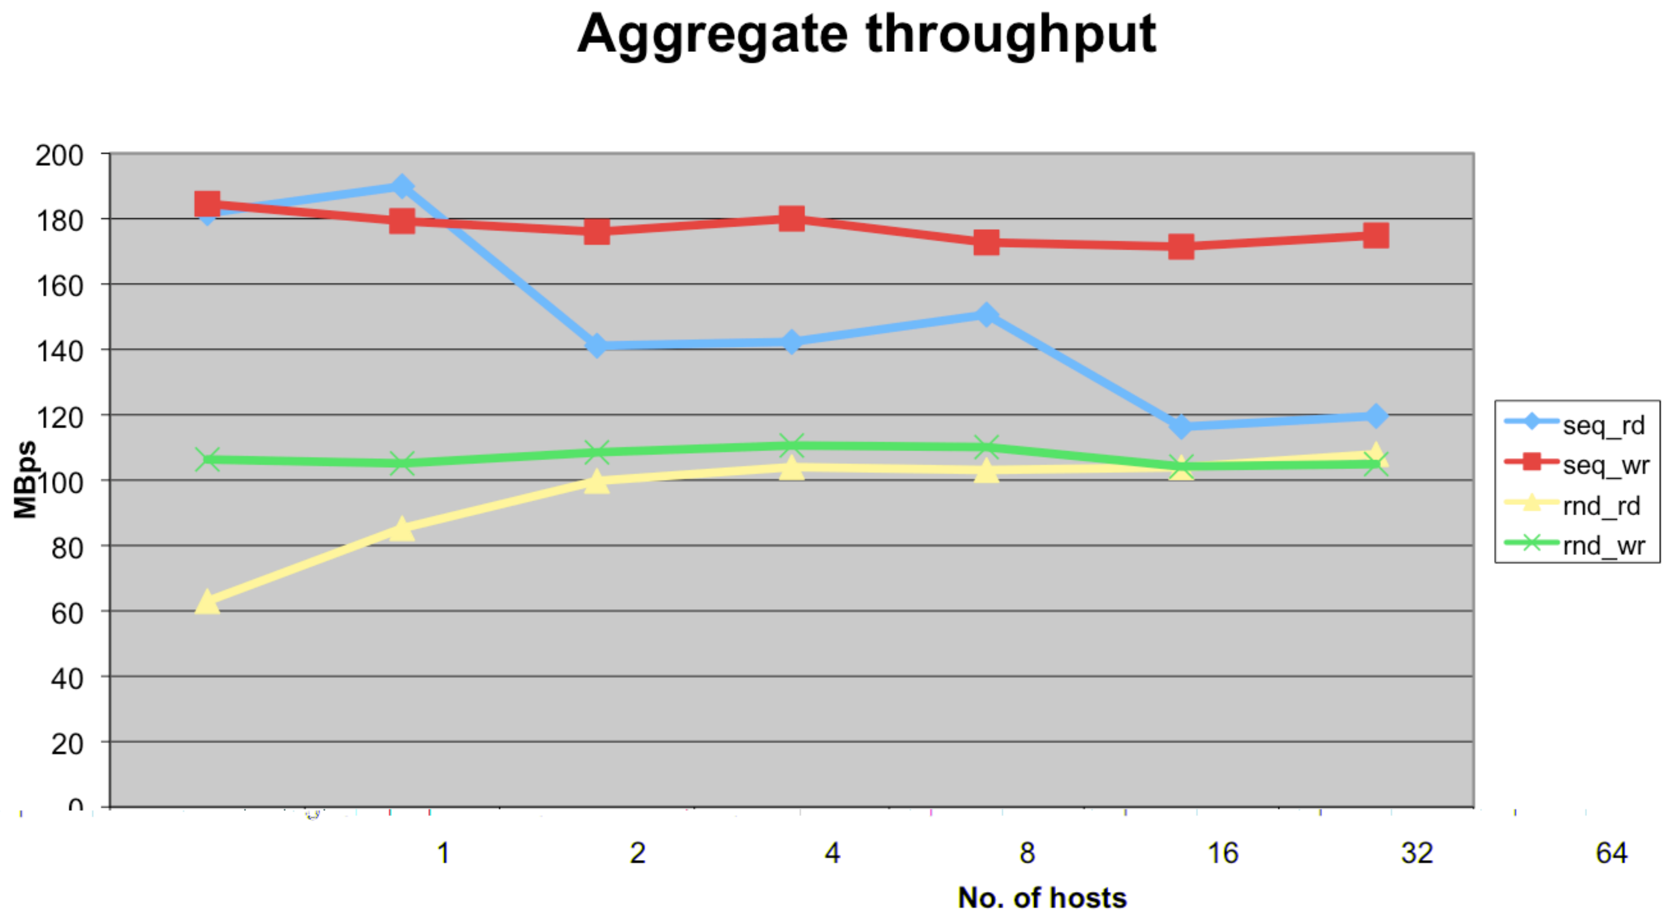

Creating Virtual Machines

Two options

- P2V
- Create a fresh VM - **Recommended**

Select correct guest OS when creating virtual machine

Install latest version of VMware tools

Disable unused devices

- Serial and parallel ports
- USB devices
- Floppy drive, CD-ROM/ DVD-ROM

Configuration Guidelines

Smaller VMs recommended

- SMP VMs have co-scheduling overheads

Avoid over-commitment

- Linear scalability until CPU or memory are over-committed

Do NOT disable vSphere optimization features

- Several optimizations for CPU, memory, Storage and Networking to improve performance in shared environment

Leverage para-virtualized drivers

- Storage (pv-scsi), Network (vmxnet3)

Run-time Guidelines

Idle Virtual Machines

- Turn off virtual machines that are not required
- Do consume resources

Benchmark Clients

- Can use virtual machines
- Pay attention to time keeping issue if VMs are used
- Measuring end user response time
  - Will include network latency
- Remote performance monitoring

Virtualizing Application Stack

Multi-tier applications

Obviously NOT virtualizing to run single VM per server

So how do you mix-n-match

Overheads multiply

Orchestra analogy

Performance benchmarking

- Virtualize one at a time
- Virtualize all at once and then optimize

Mix-n-Match Application Stack Components

Understand the load profile of the different components in the application stack

Place the VMs accordingly

- Do not combine all CPU or I/O or memory intensive VMs on the same server

Horizontal scalability or scaling out

- Do not place similar multiple VMs on the same ESX server
  - Similar load profile can result in bottlenecks
  - Lead to single point of failure in application stack

Do not have to virtualize all components

VMmark Benchmark

Consider VMmark: A Scalable Benchmark for Virtualized Systems

- Provides meaningful measurement of virtualization performance.
- Generates an easily understandable and precise metric that scales with underlying system capacity.
- Used to effectively compare the suitability and performance of different hardware platforms for virtual environments.

Employs realistic, diverse workloads running on multiple operating systems

VMware worked with SPEC to establish industry-wide credibility and standardization

VMmark: Mixed Workloads

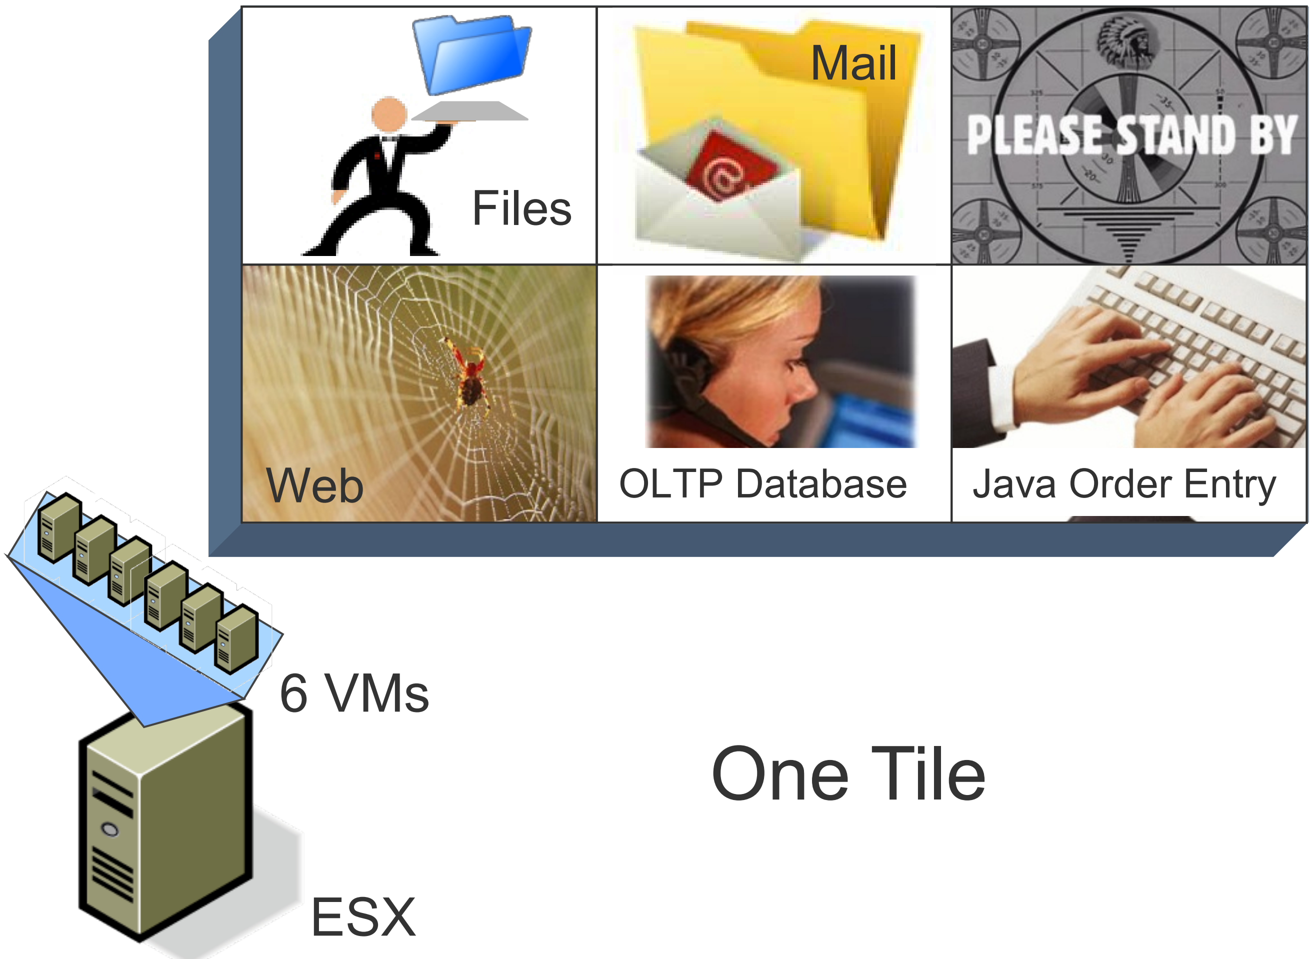

One Tile

Benchmark Topology

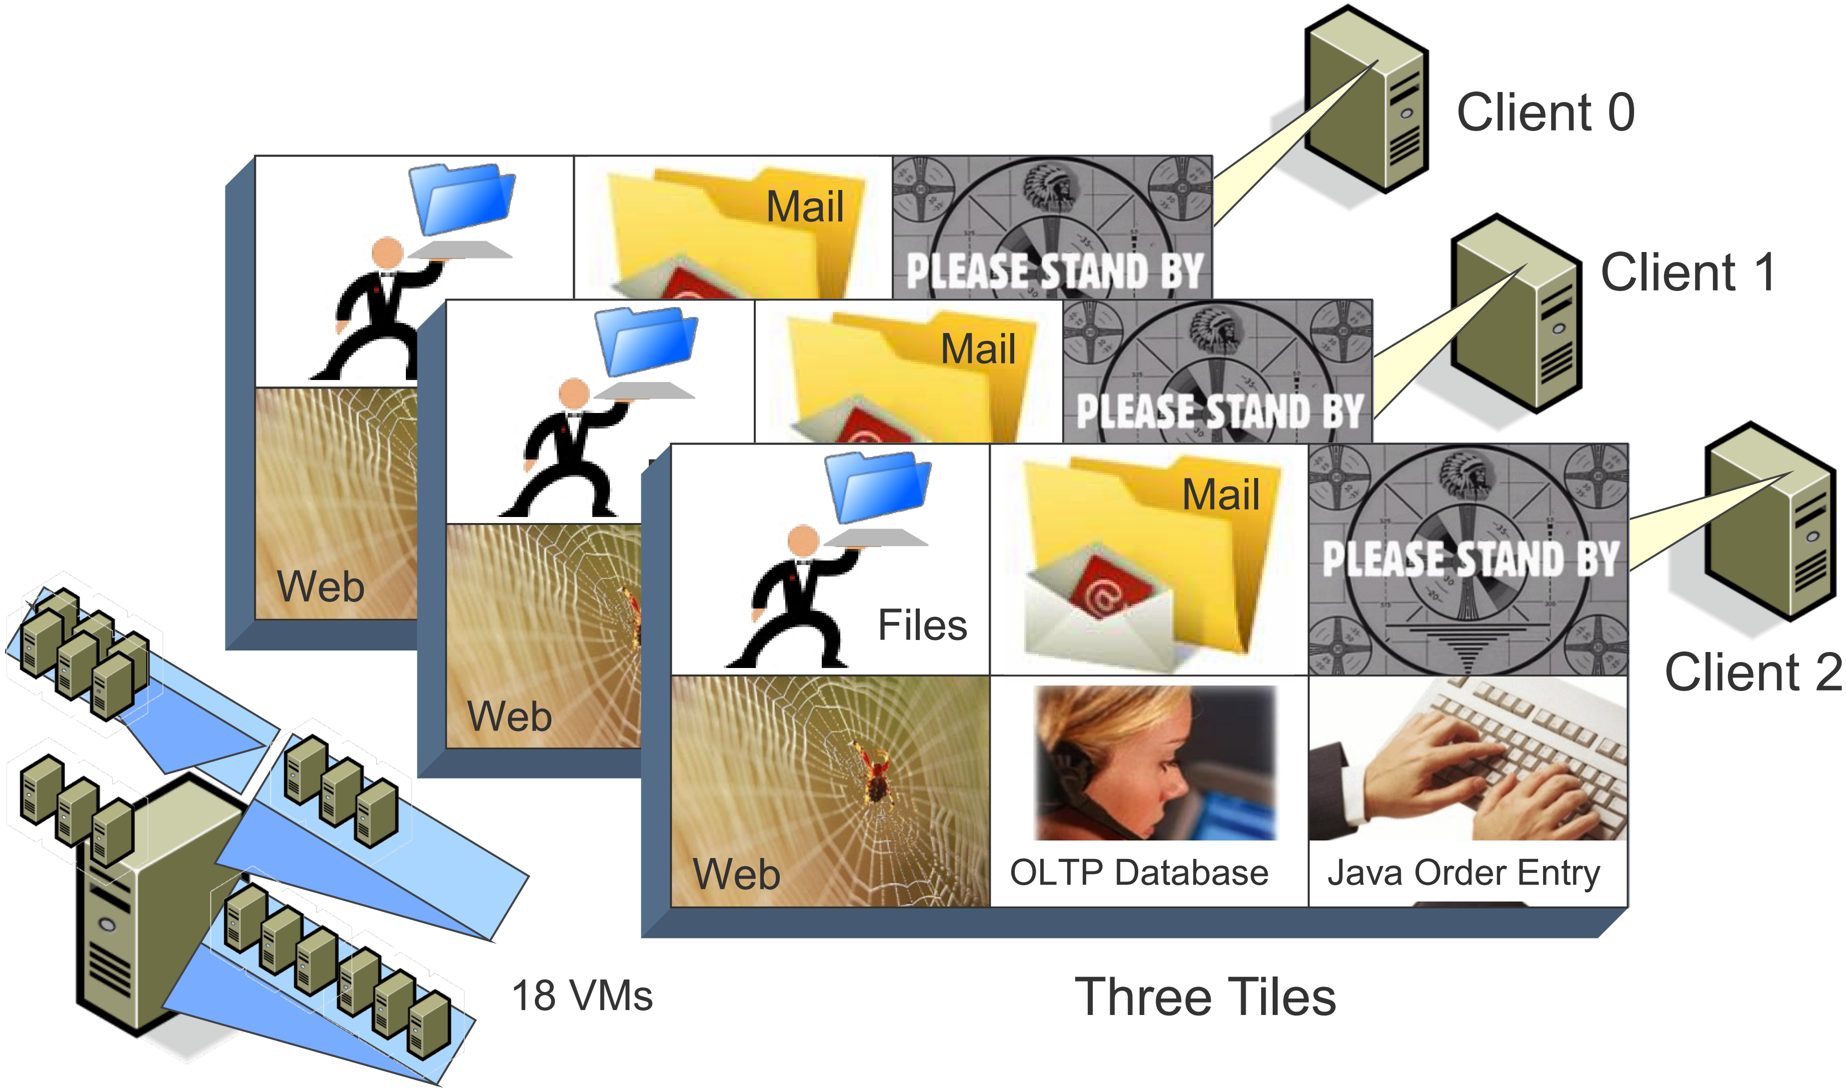

Agenda

Understanding Virtualization Performance

Performance Benchmarking in Virtual Environment

Performance Monitoring in Virtual Environment

Virtual Better than Physical?

Troubleshooting Tools

- **vCenter (high level)**
  - Historical performance data (*check statistics levels)
  - Consolidated metrics for all hosts / datastores in environment

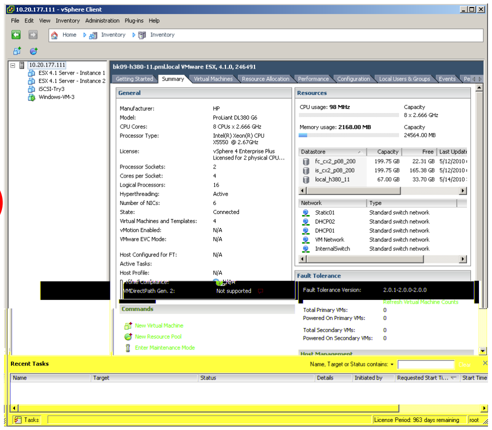

- **vscsiStats (storage guts)**
  - Detailed virtual SCSI device latency metrics
  - Seek distance, IO size, Latency
  - Displays Histograms

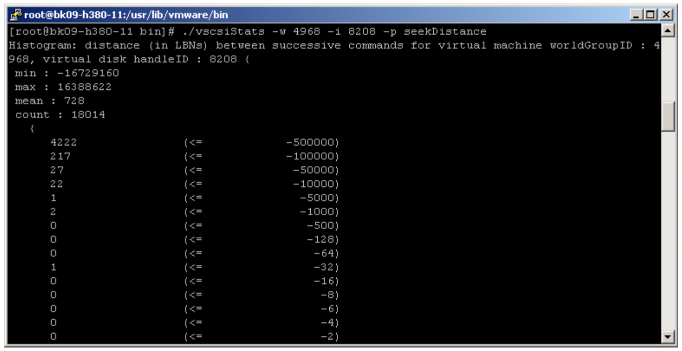

- **esxtop / resxtop (tactical)**
  - Single ESX/ESXi Host
  - Detailed performance data in real time

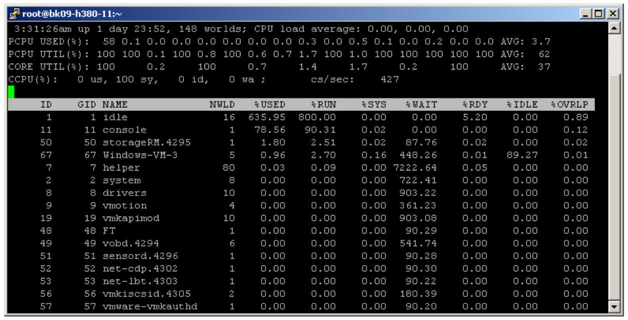

Performance Monitoring Guidelines

Agenda

Understanding Virtualization Performance

Performance Benchmarking in Virtual Environment

Performance Monitoring in Virtual Environment

Virtual Better than Physical?

Virtual better than Physical?

Workaround application stack scalability limits

Components of application stack do not scale in physical environment

Virtualize and achieve higher scalability

- Higher ROI despite virtualization overheads

Examples

- Citrix Presentation Server 32-bit
  - 32-bit operating system limitations prohibit scaling
- Java Virtual Machines
  - Heap size limits on 32-bit Windows operating system

Multi-VM Approach to Scalability: Building Blocks

Some applications have limited SMP scalability

- Web Server, Mail Server, Java Virtual Machines

Virtualization onto the ESX platform enables horizontal scaling

- Concurrently run several operating environments
- Advanced Resource Management features allow efficient sharing of compute, network and storage resources

Let's look at examples : Exchange, Apache & Citrix XenApp

Virtualization-aware Architecture: Building Blocks

Many applications lack scalability beyond certain CPUs

- Apache web server, WebSphere, Exchange
- Configure vCPUs to application scalability limits
- For additional capacity instantiate more of such VMs

SPECweb2005 Native and Virtual Scaling

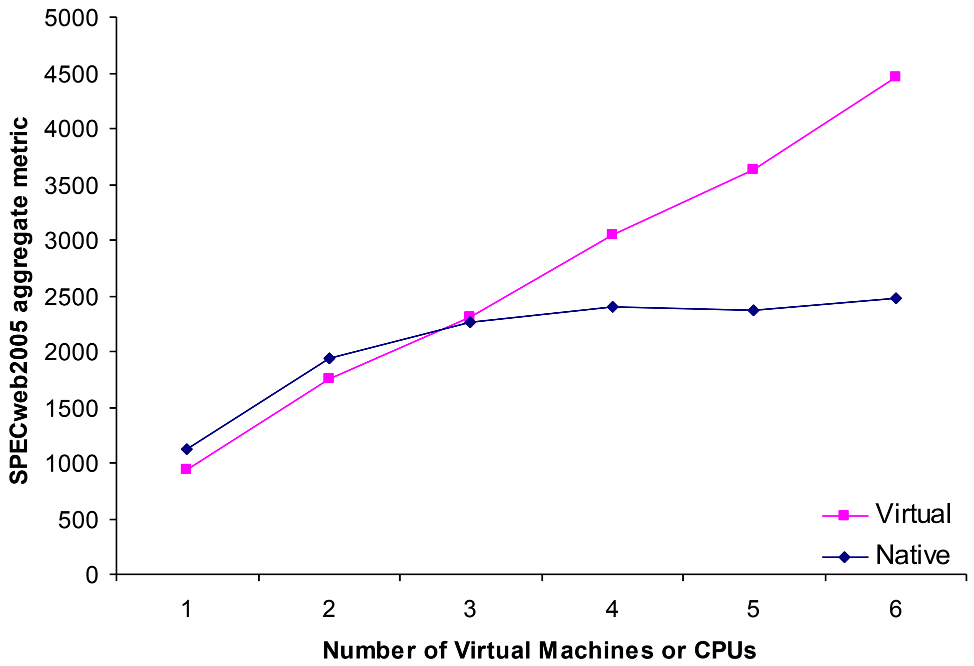

Citrix XenApp\ Microsoft Terminal Services

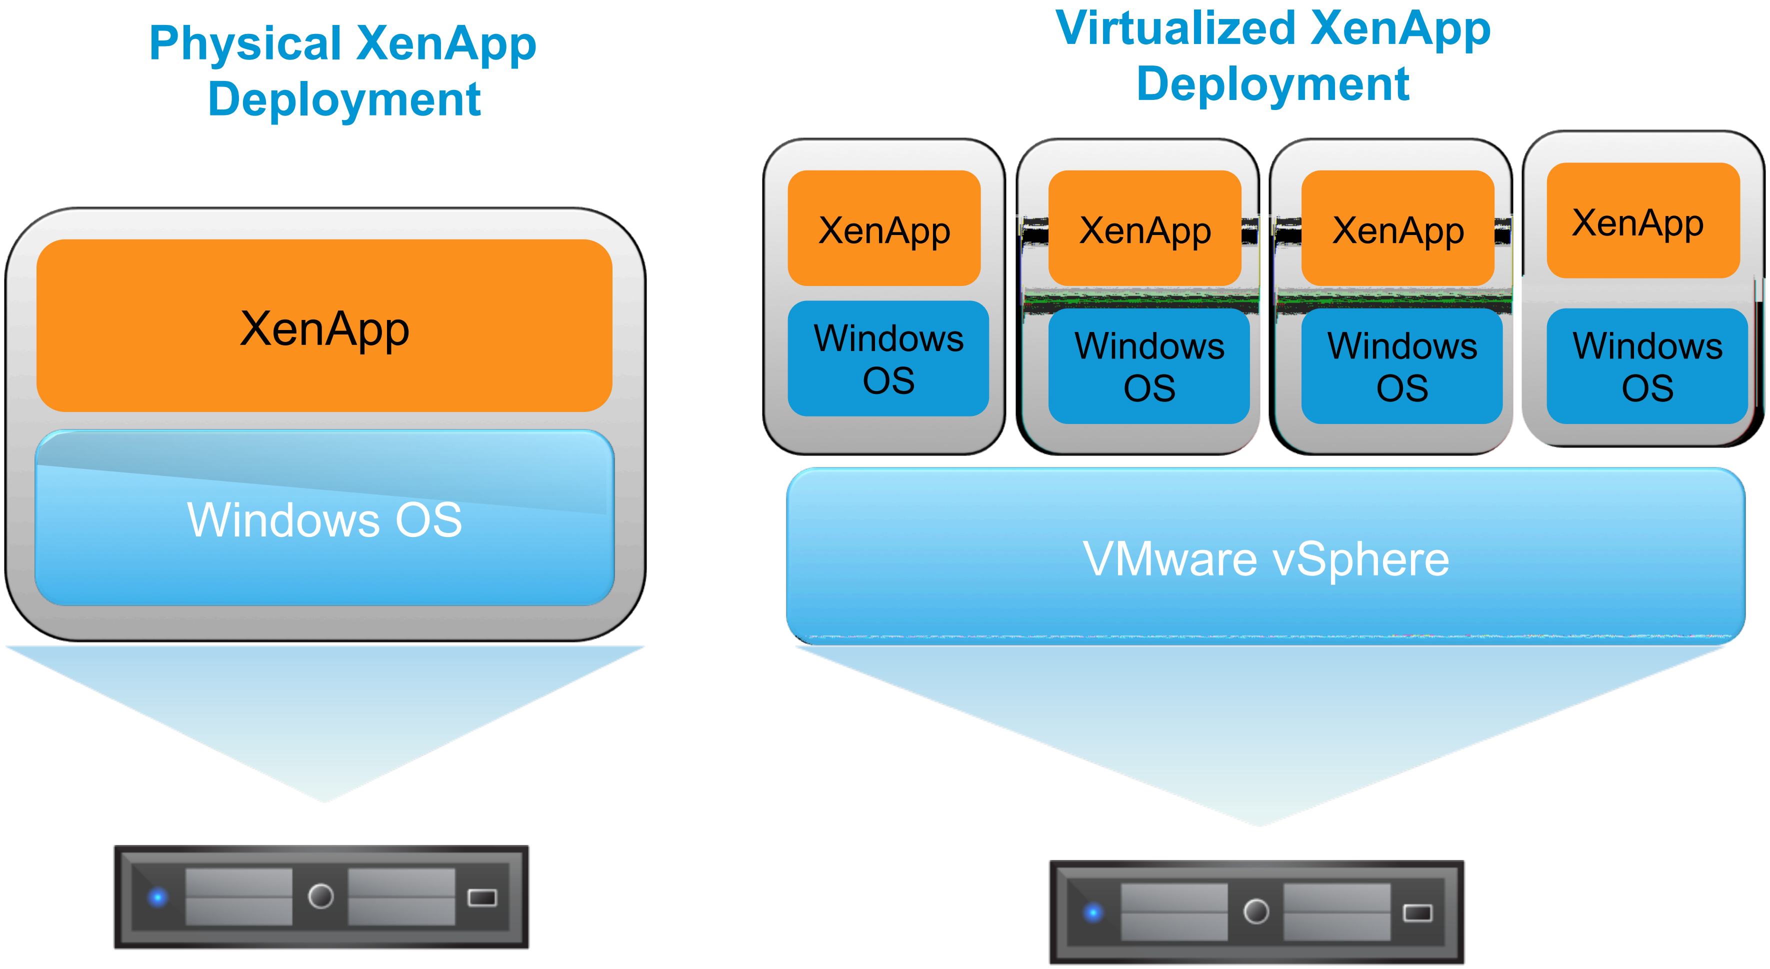

Example Migration Scenario 4_4_0_0 with DRS

Balanced Cluster

DRS Scalability – Transactions Per Minute

(Higher the better)

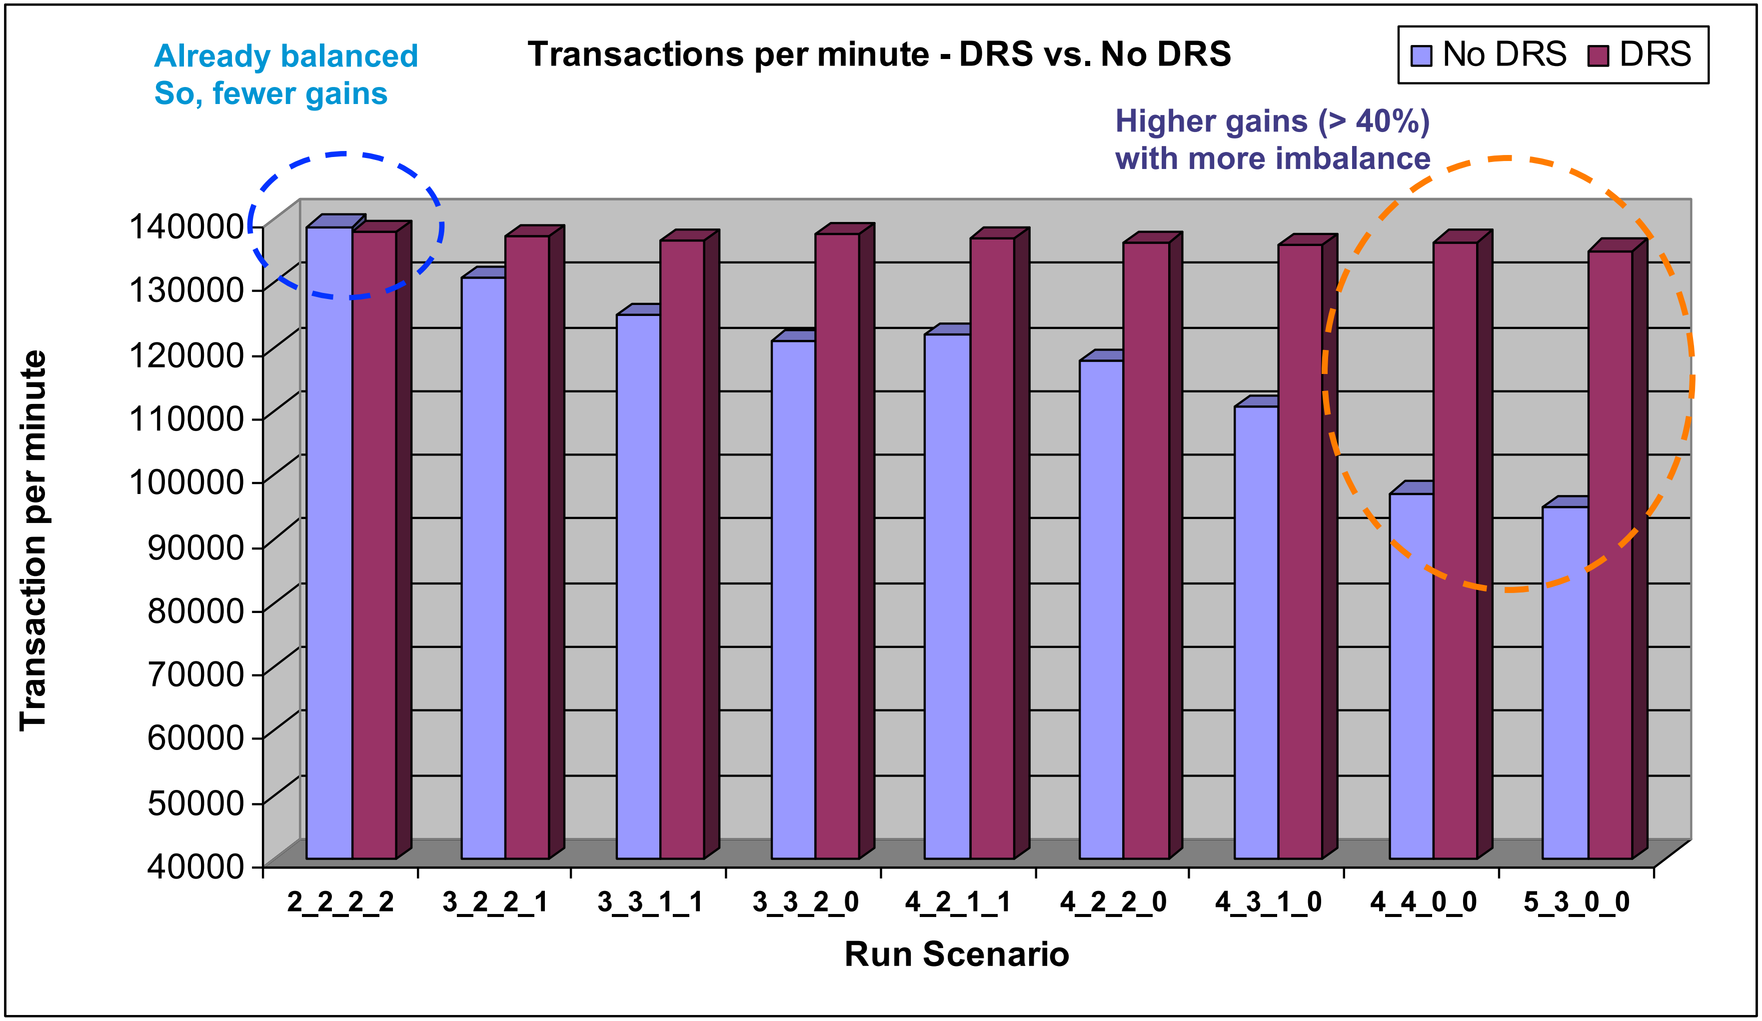

DRS Scalability – Application Response Time

(Lower the better)

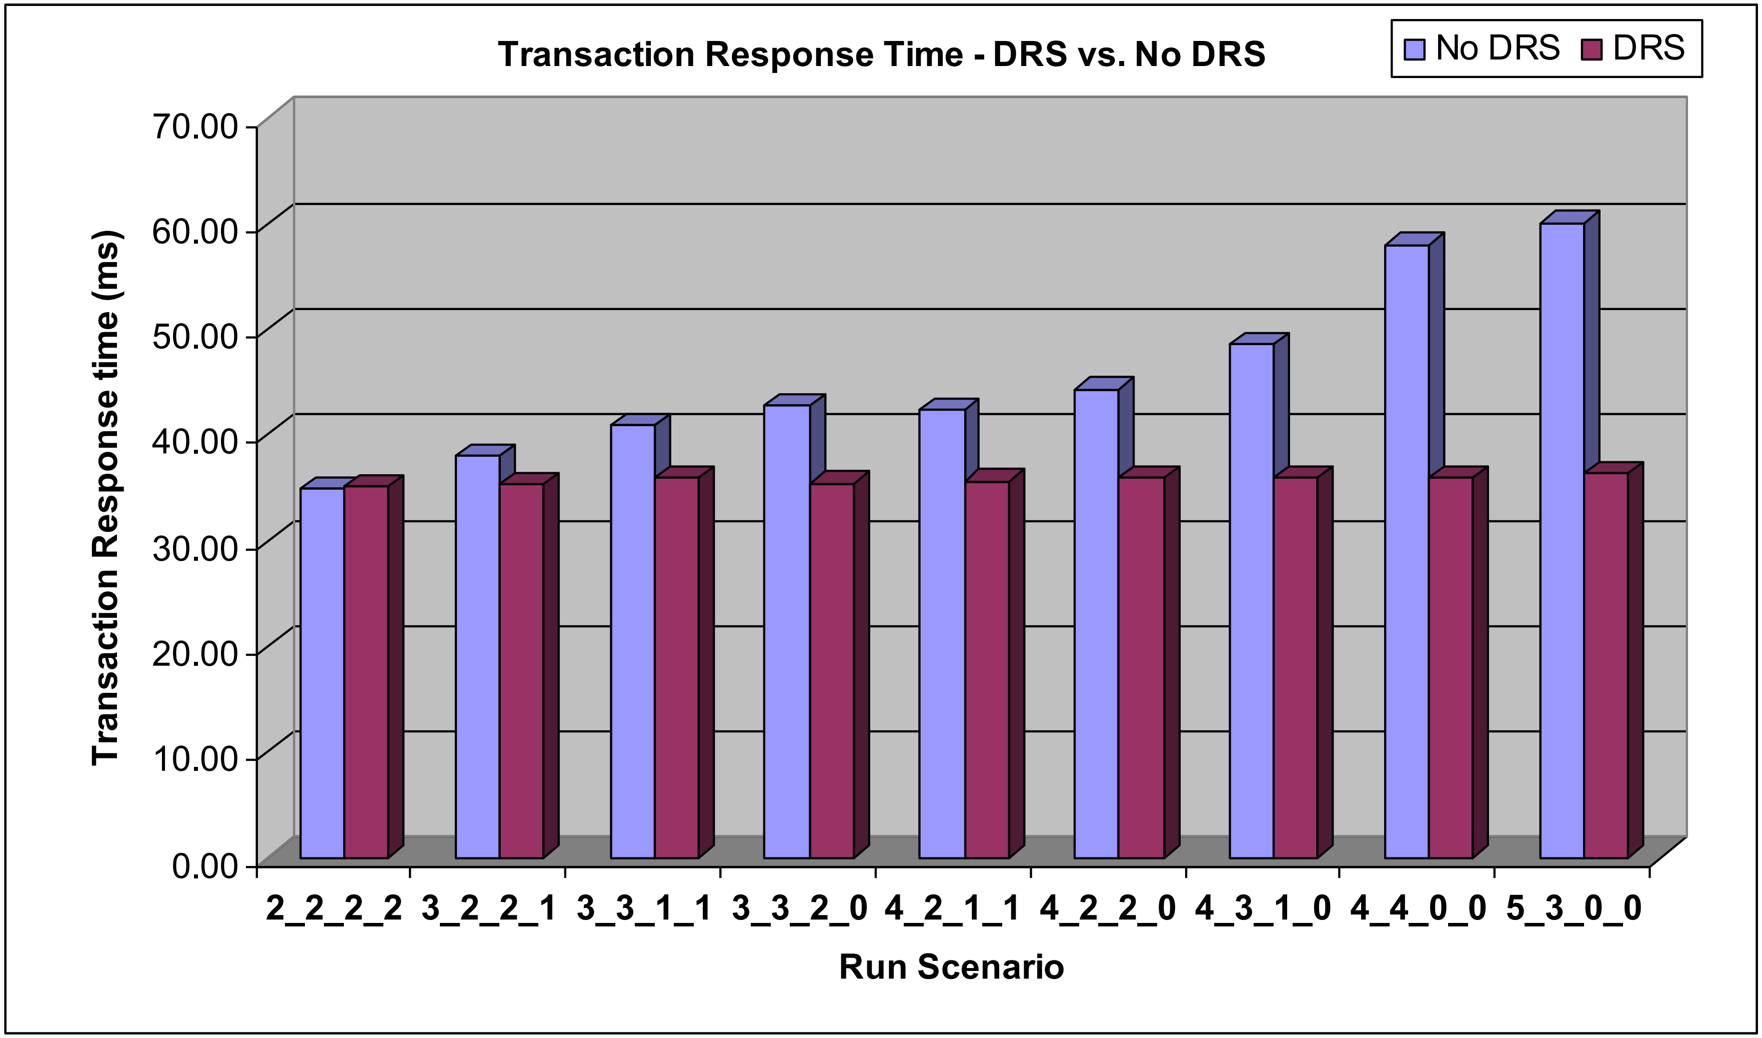

Q & A