Scaling in a Hypervisor Environment

Richard McDougall

Chief Performance Architect

VMware

VMware ESX Hypervisor Architecture

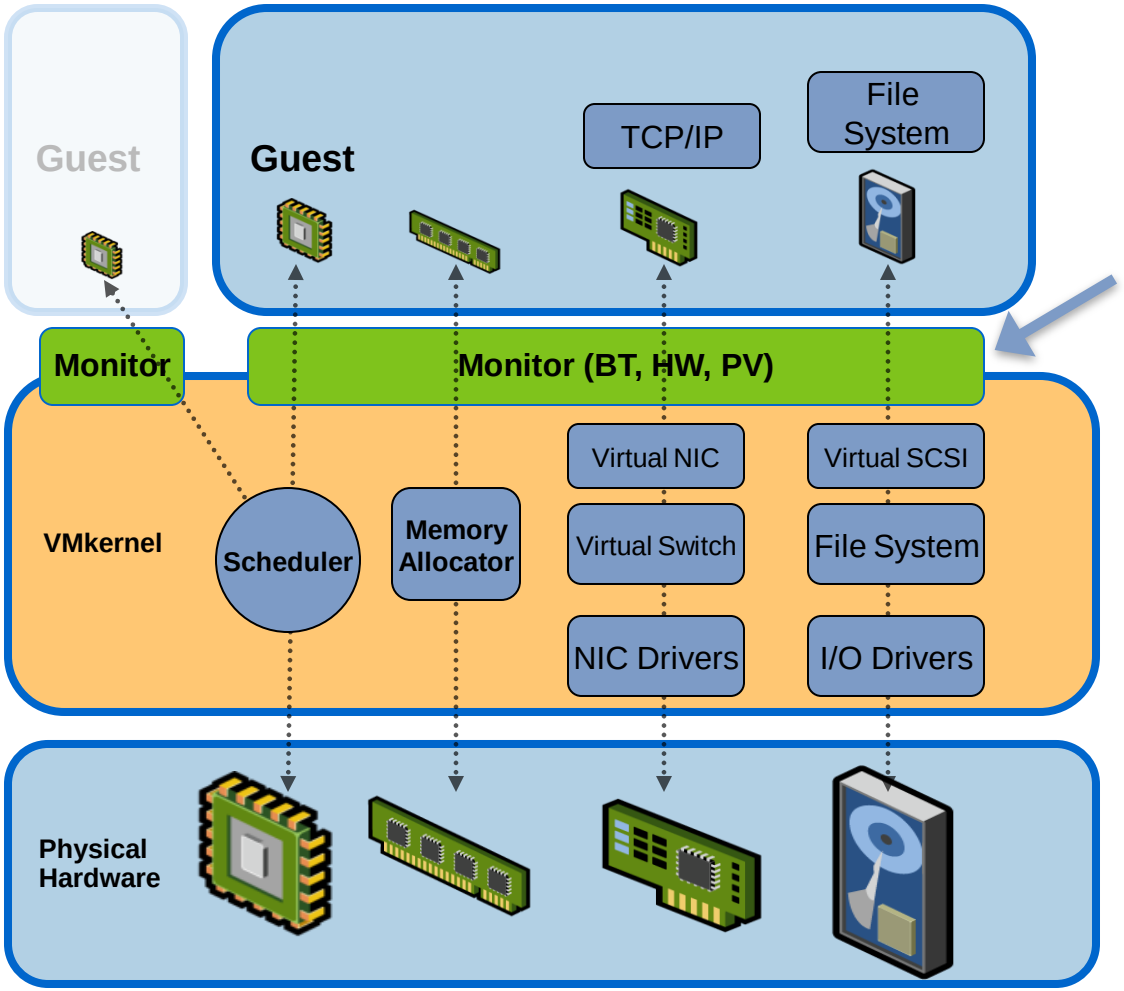

CPU is controlled by scheduler and virtualized by monitor

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

Evolution of Performance for Large Apps on ESX

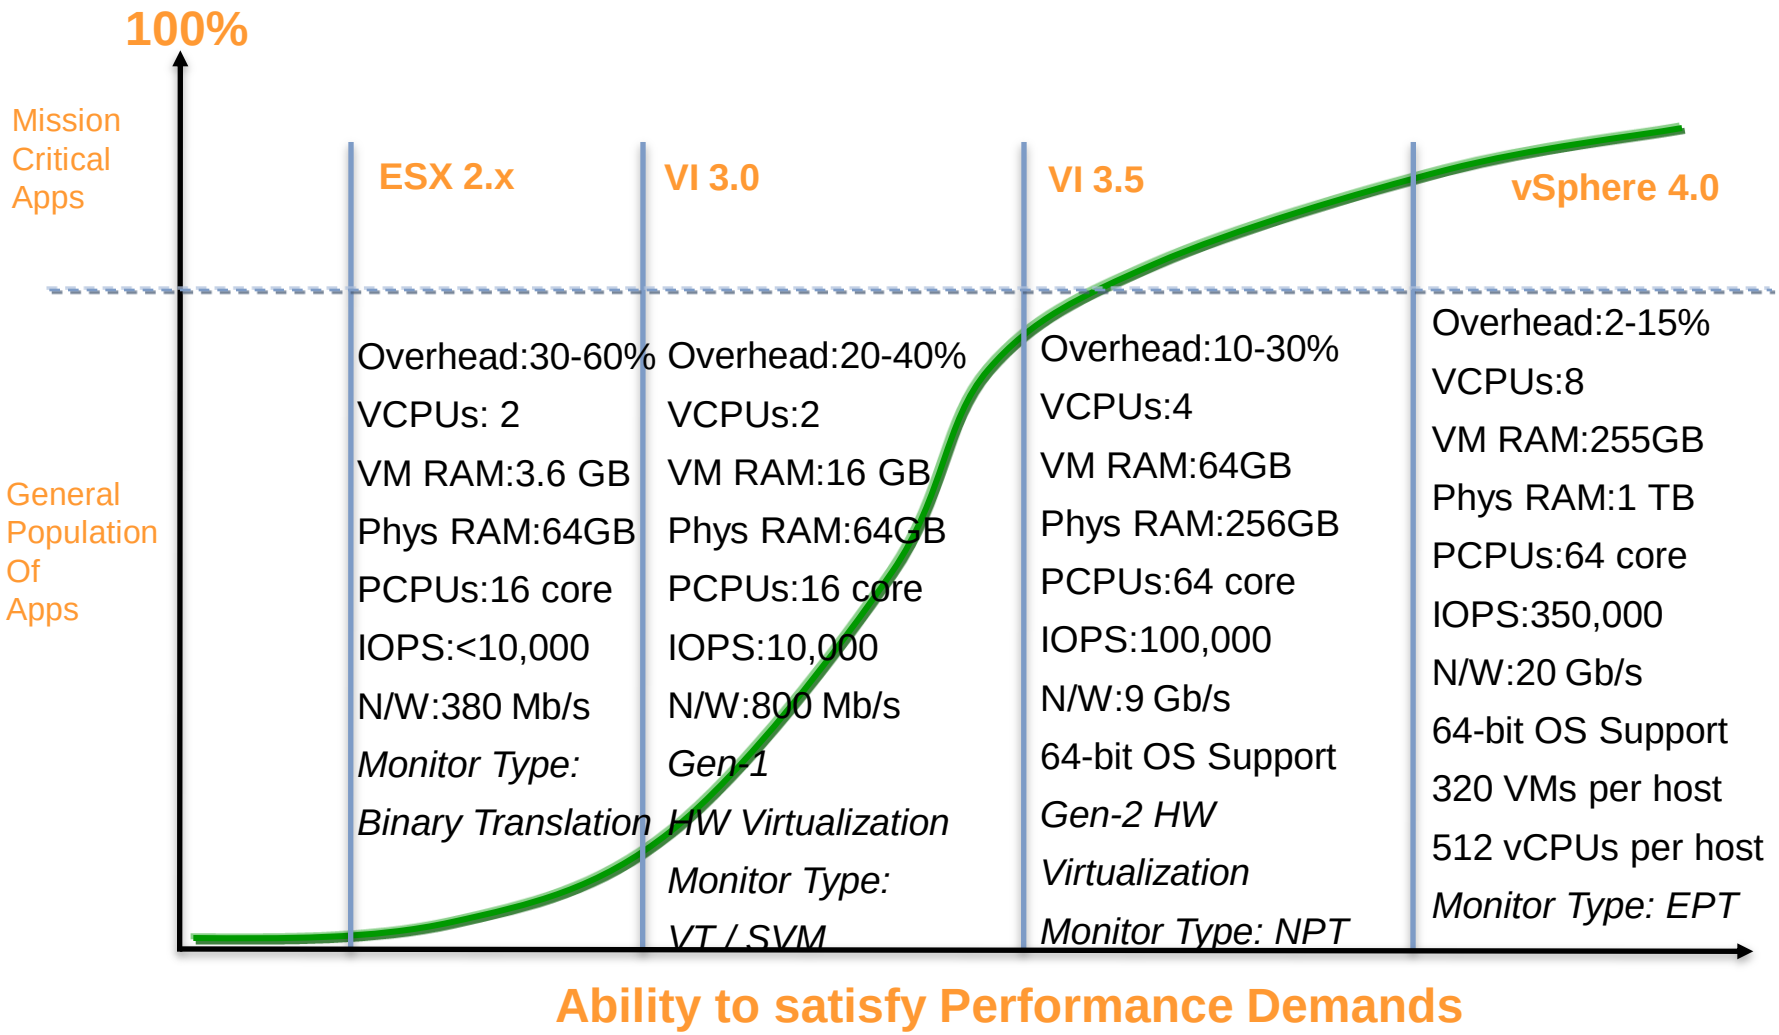

Databases: Why Use VMs Rather than DB Virtualization?

Virtualization at hypervisor level provides the best abstraction

- Each DBA has their own hardened, isolated, managed sandbox

Strong Isolation

- Security
- Performance/Resources
- Configuration
- Fault Isolation

Scalable Performance

- Low-overhead virtual Database performance
- Efficiently Stack Databases per-host

Large Oracle Transaction Workload: SwingBench

Database:
- # of Users = 5,011,872
- # of Products = 5,011,872
- Db size = 5.11GB

Test
- Test duration = 10 mins
- # of Users per run = 30
- Think time = 0
- New customer - 11%
- Browse products - 28%
- Order products - 28%
- Process orders - 5%
- Browse orders - 28%

- Dell PowerEdge 2950
- Mem: 8GB
- Two dual core) Intel(R) Xeon(R) CPU 5160 @ 3.00GHz processor 4GB cache
- Storage: CX 3-40 (30 disks)

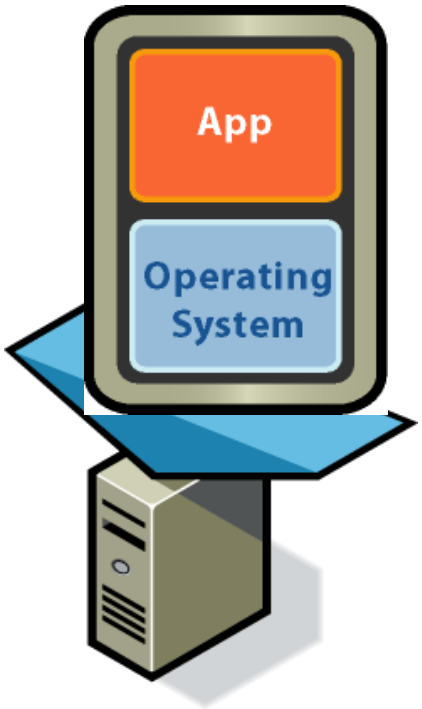

Powered By ORACLE

- Oracle 11g (11.1.0)
- RHEL4 x86_64
- SGA: 3GB
- PGA: 1GB

- RHEL5 U1
- 64 bit
- 2.6.18-53.1.13.el5

- ESX version: 3.5 build# 60217
- Number of VMs: 1
- vCPU: 4
- Mem: 6GB
- vDisk: 16GB
- vNIC: 1

vmware

Measuring the Performance of DB Virtualization

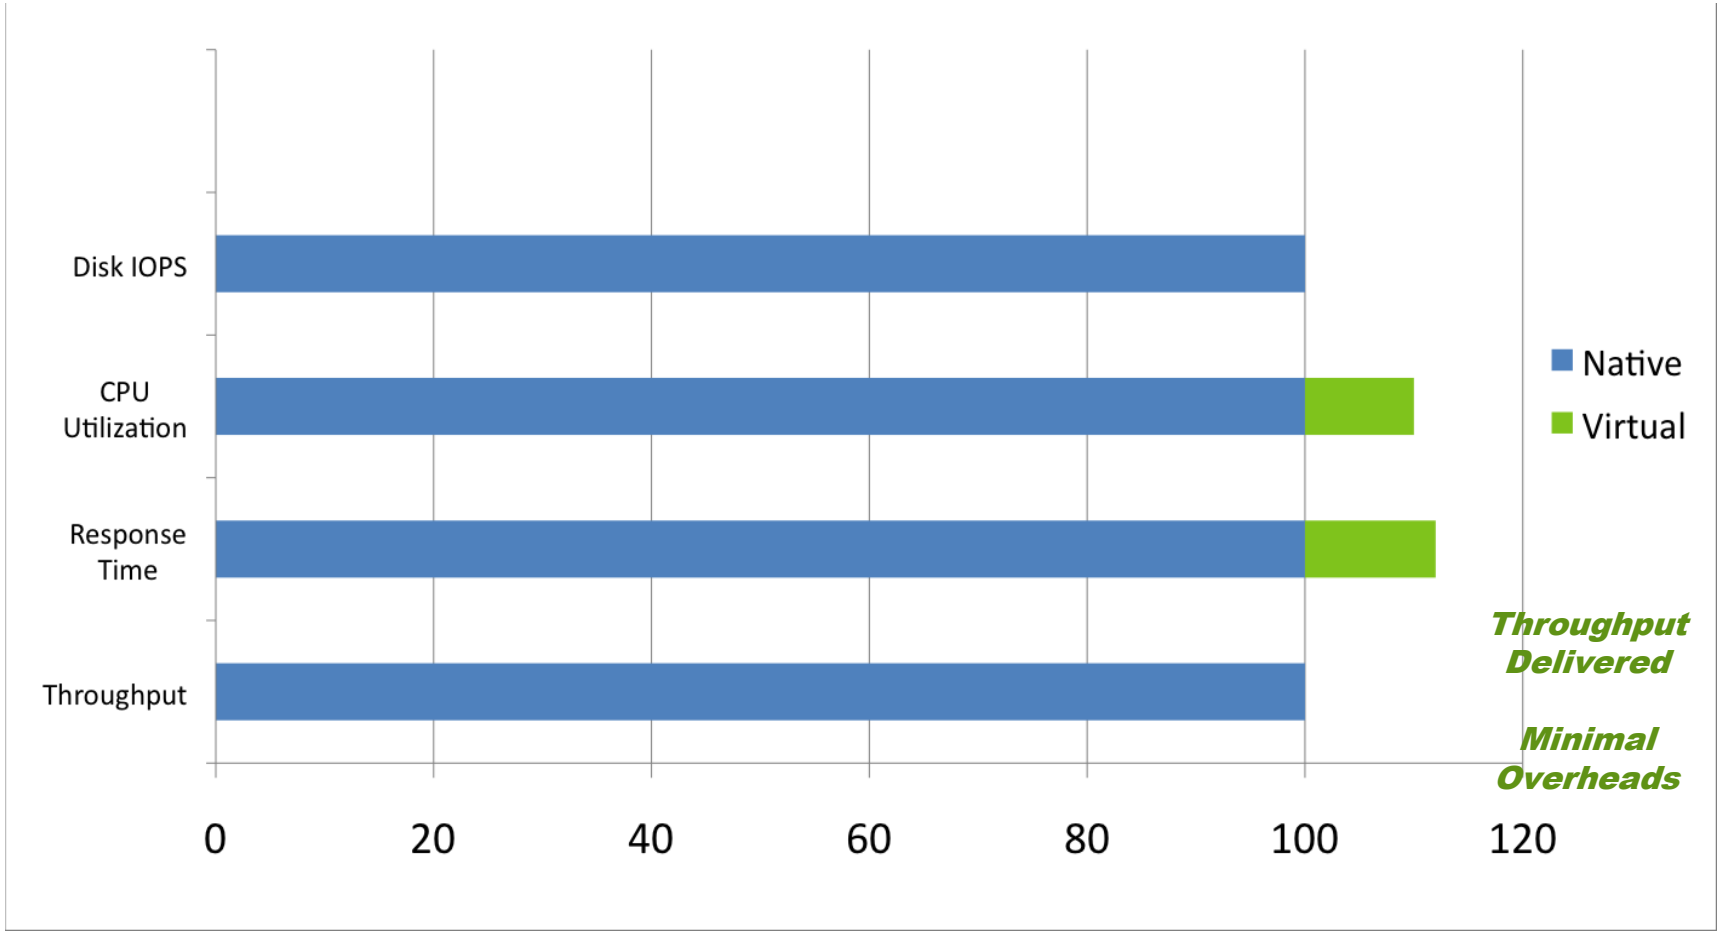

Single VM Performance: Well-Known Database OLTP Workload†

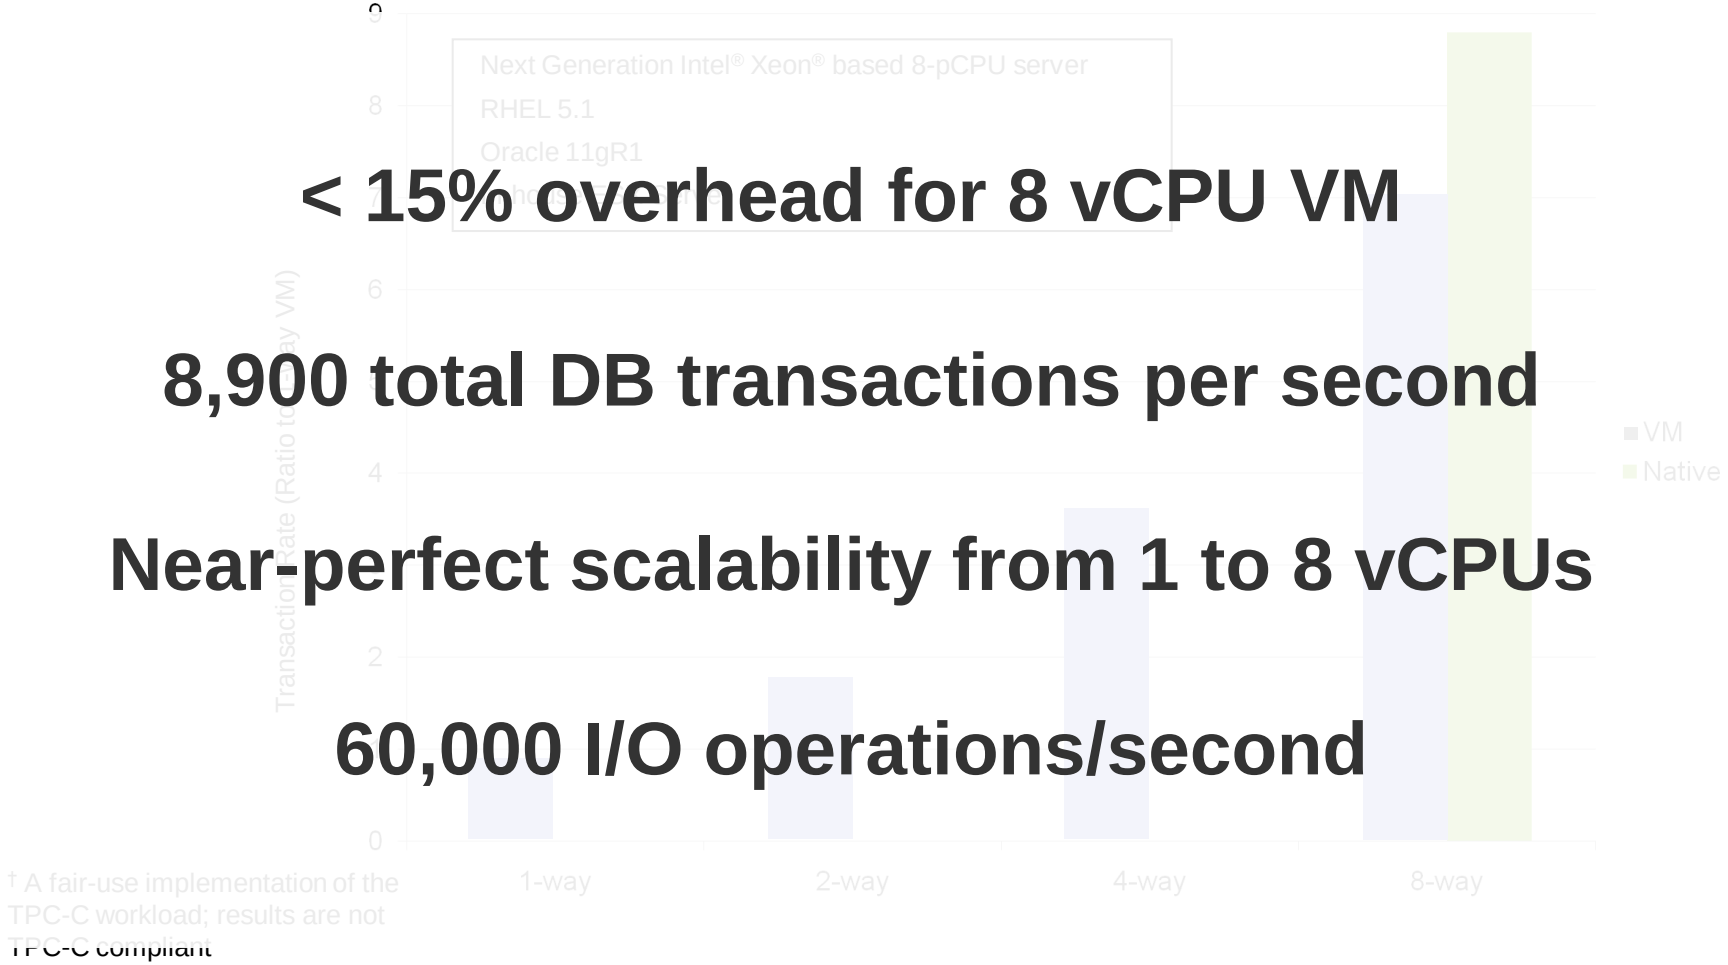

< 15% overhead for 8 vCPU VM

8,900 total DB transactions per second

Near-perfect scalability from 1 to 8 vCPUs

60,000 I/O operations/second

† A fair-use implementation of the TPC-C workload; results are not TPC-C compliant

Single VM Handles Most Demanding Applications

=

Sun Fire 15k (ca. 2002)

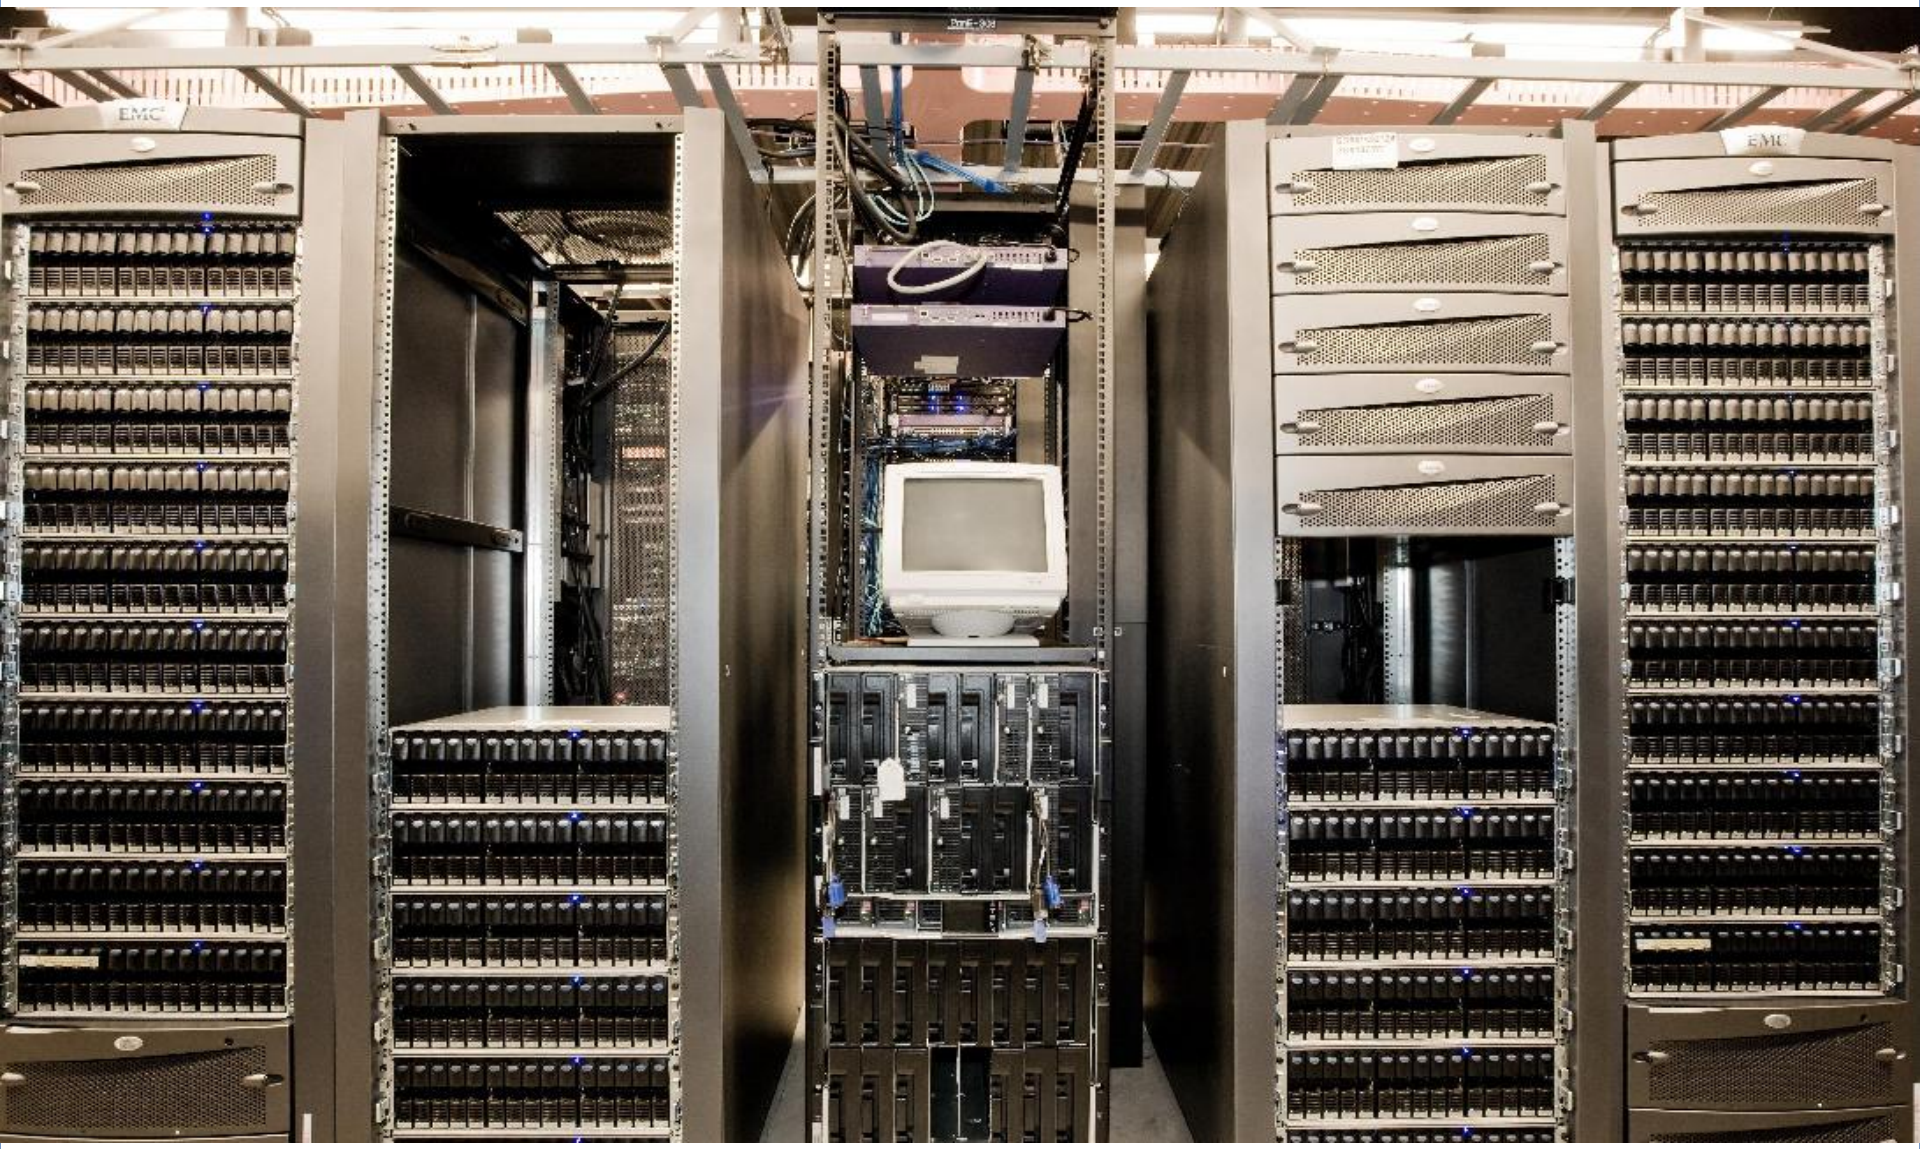
EMC
EMC

The world changed @ 3Ghz: Multi-Core Focus

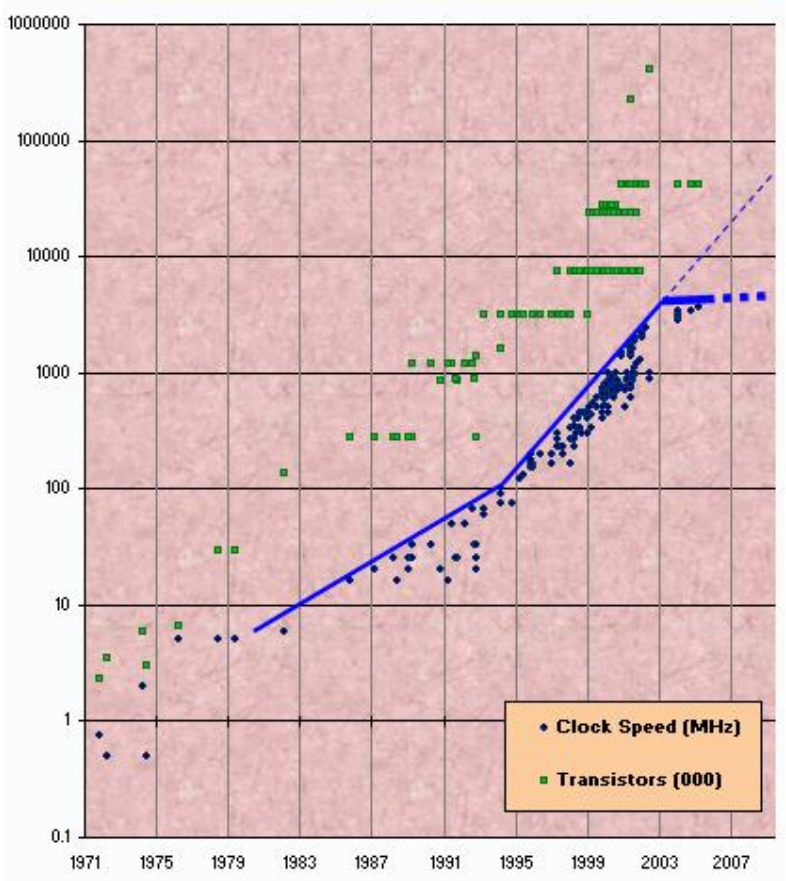

- Scaling by clock frequency no-longer possible
- Memory Latency Dominates
- Performance now achieved by implementing multiple CPU cores

Moore's Law: Direct Correlation to Number of Cores

1 Billion transistors and growing

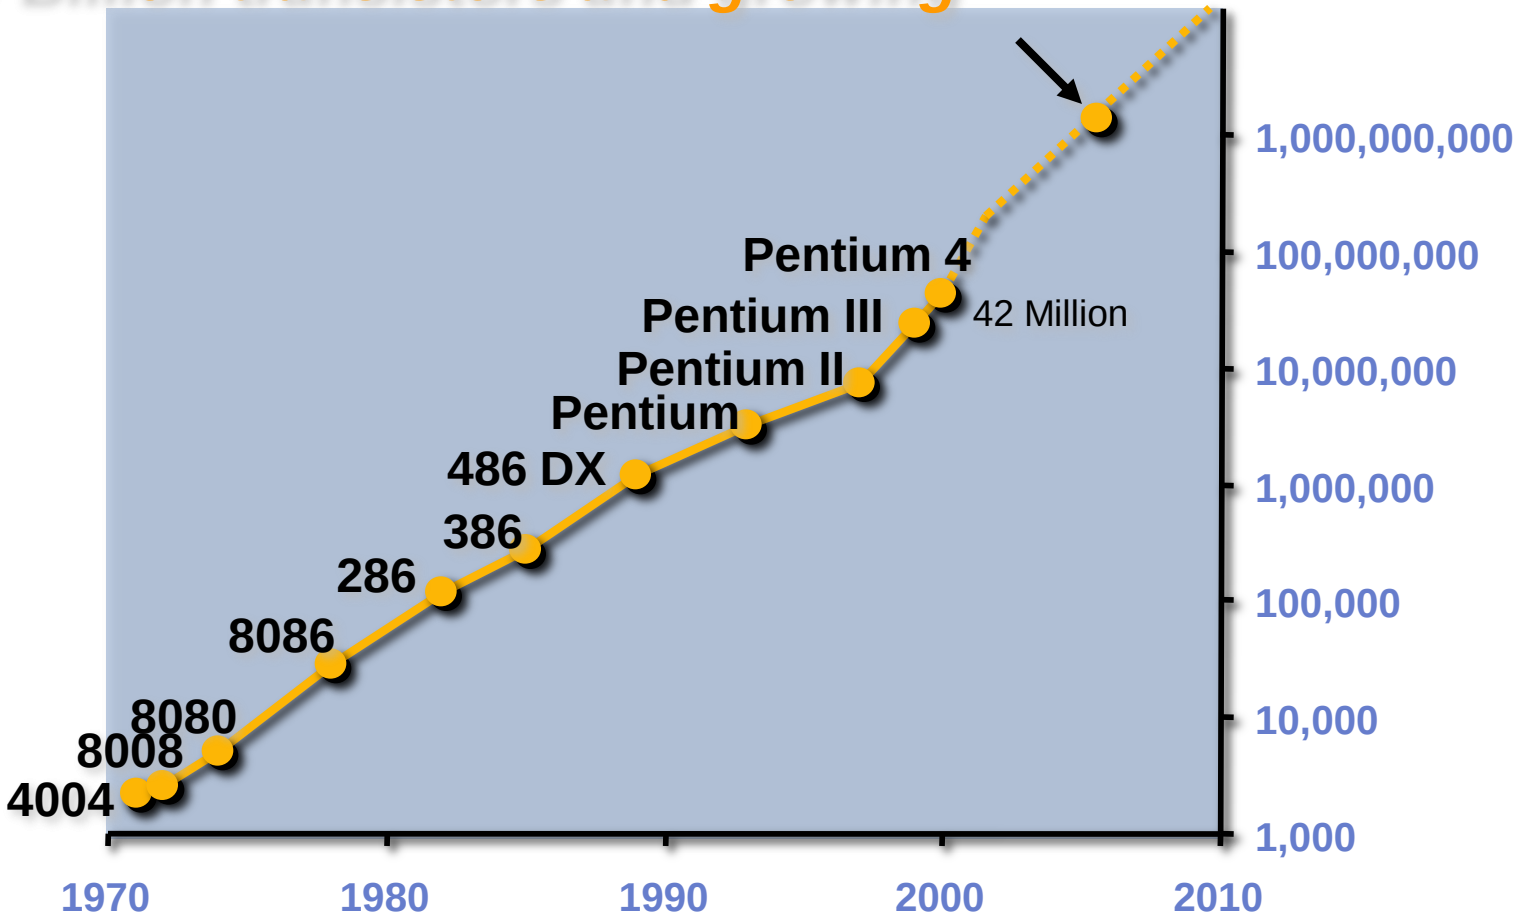

Estimated Logical CPU Counts: 2-Socket Server
(Assuming hardware threading is a given)

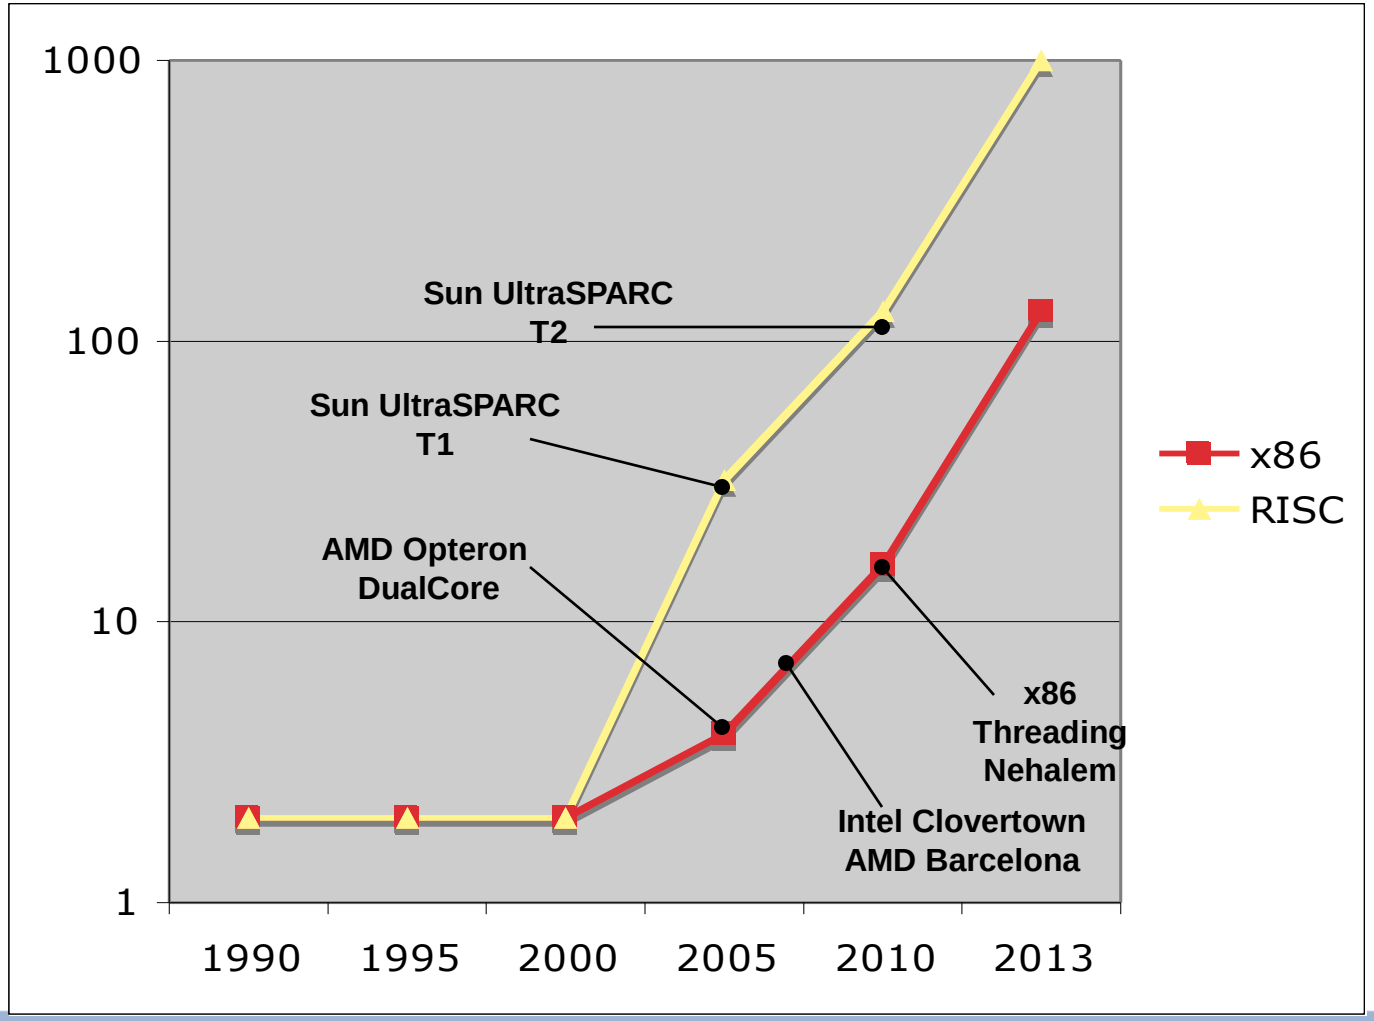

VMware vSphere enables you to use all those cores…

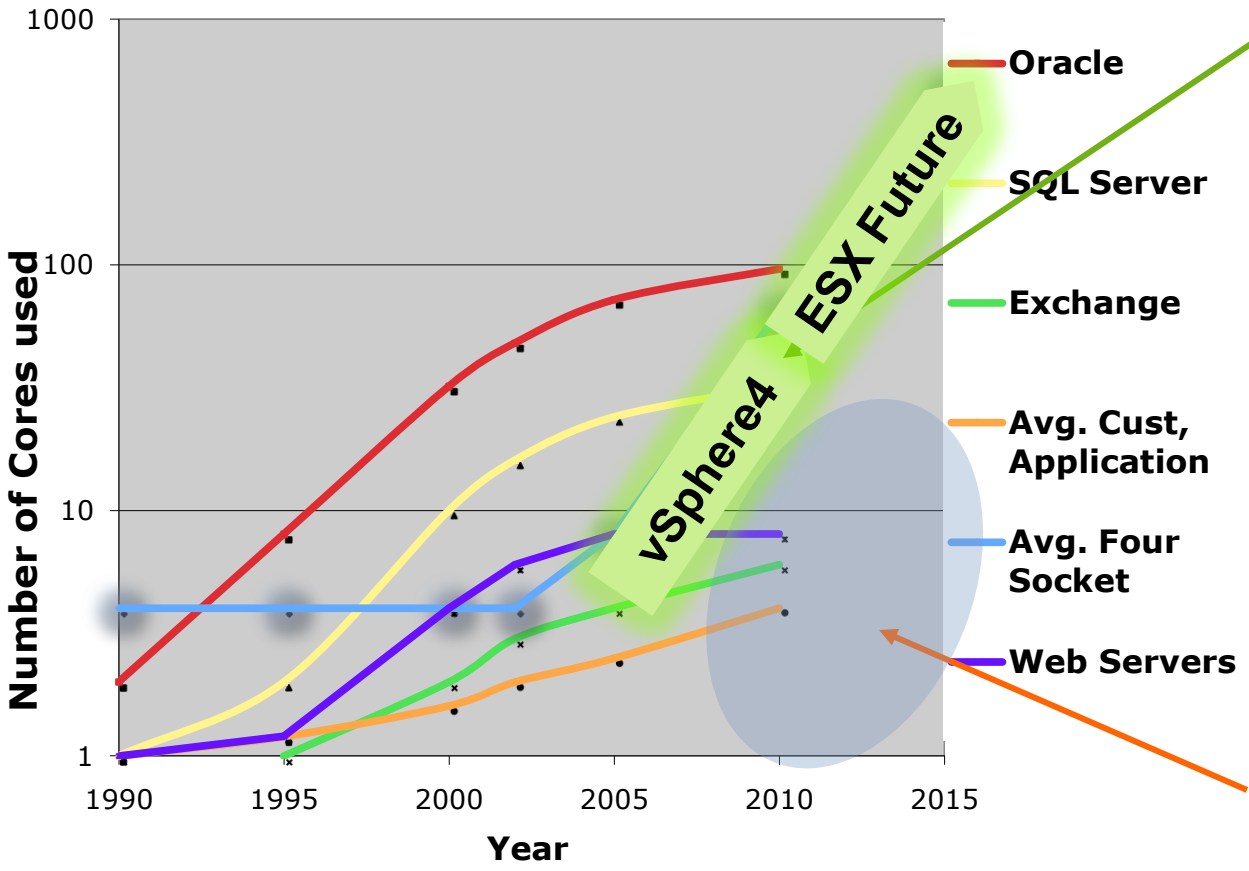

VMWare ESX Scaling: Keeping up with core counts

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

Most applications don't scale beyond 4/8 way

“Large” Database Consolidation Study

Scaling to 16 Cores, 256GB RAM!

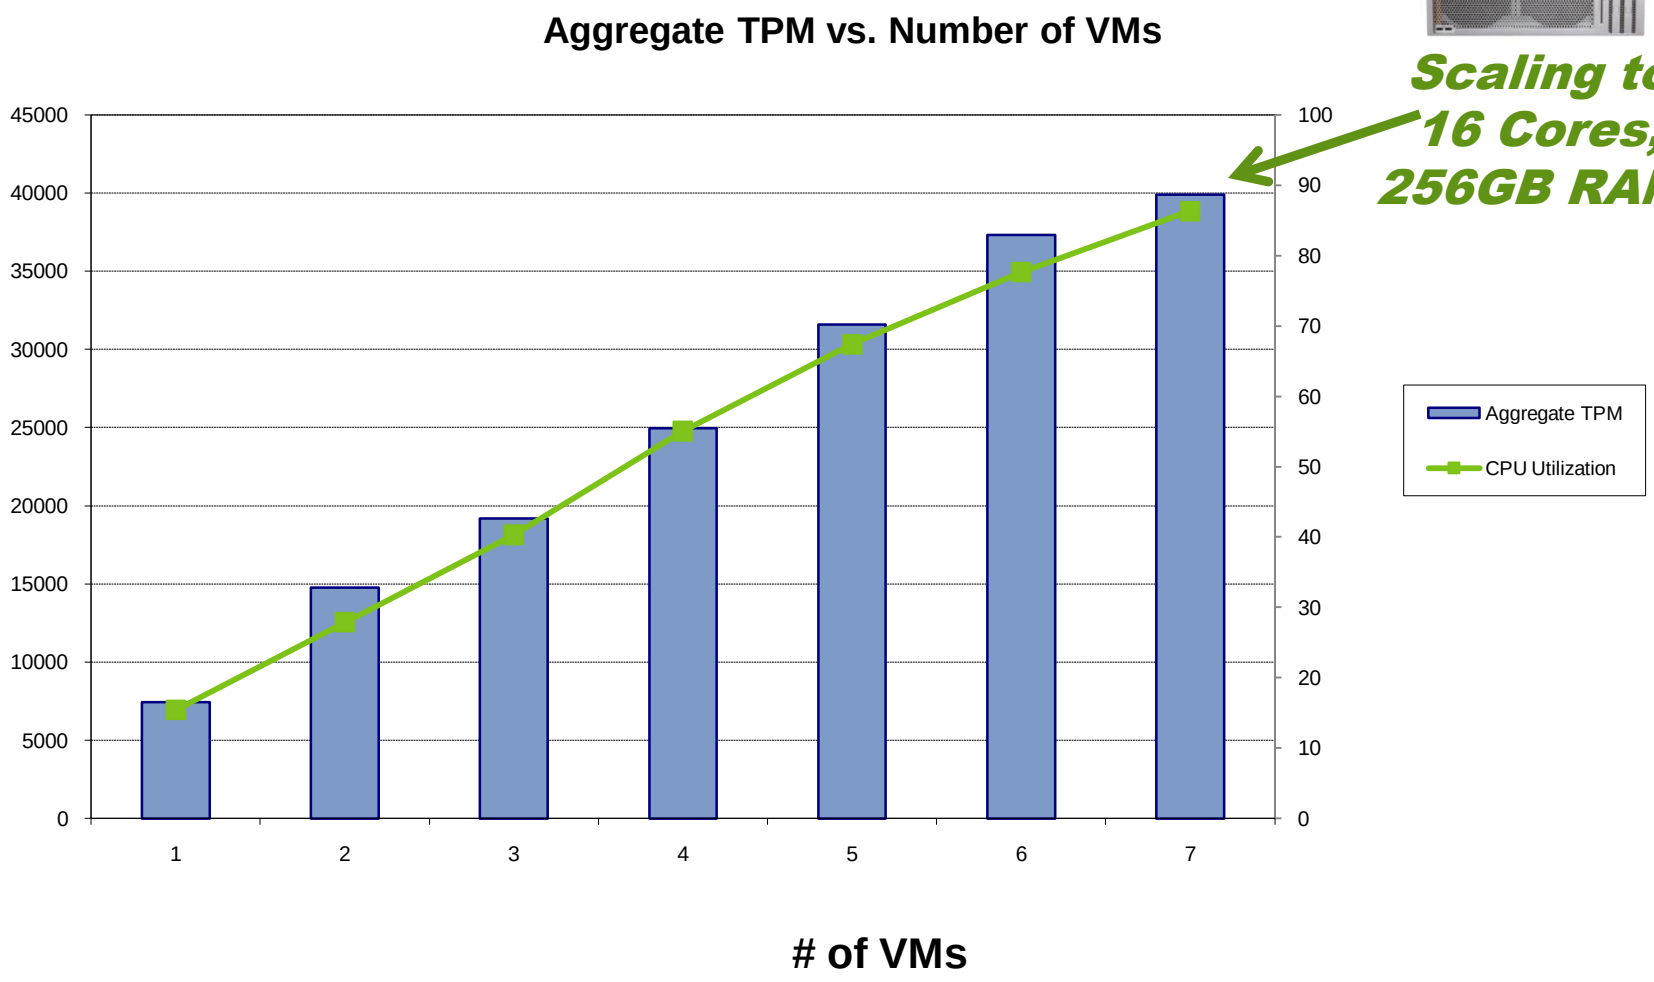

vmware

Oracle Performance (Response time)

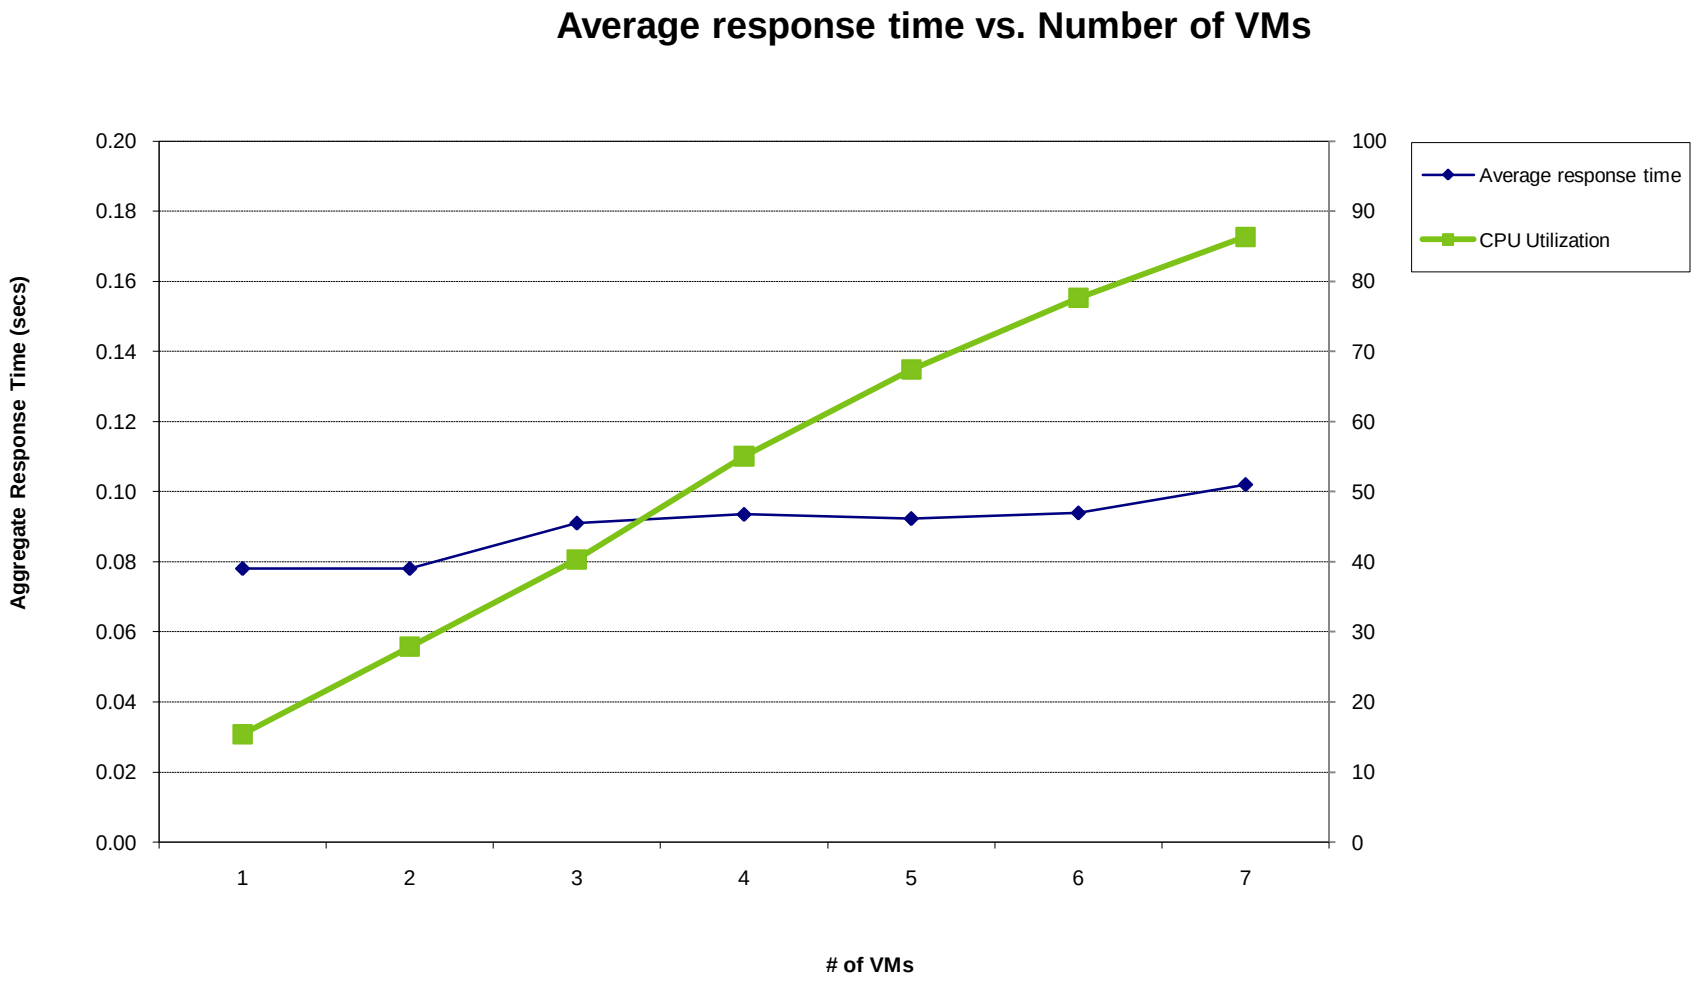

Scaling of Virtual Web vs Physical

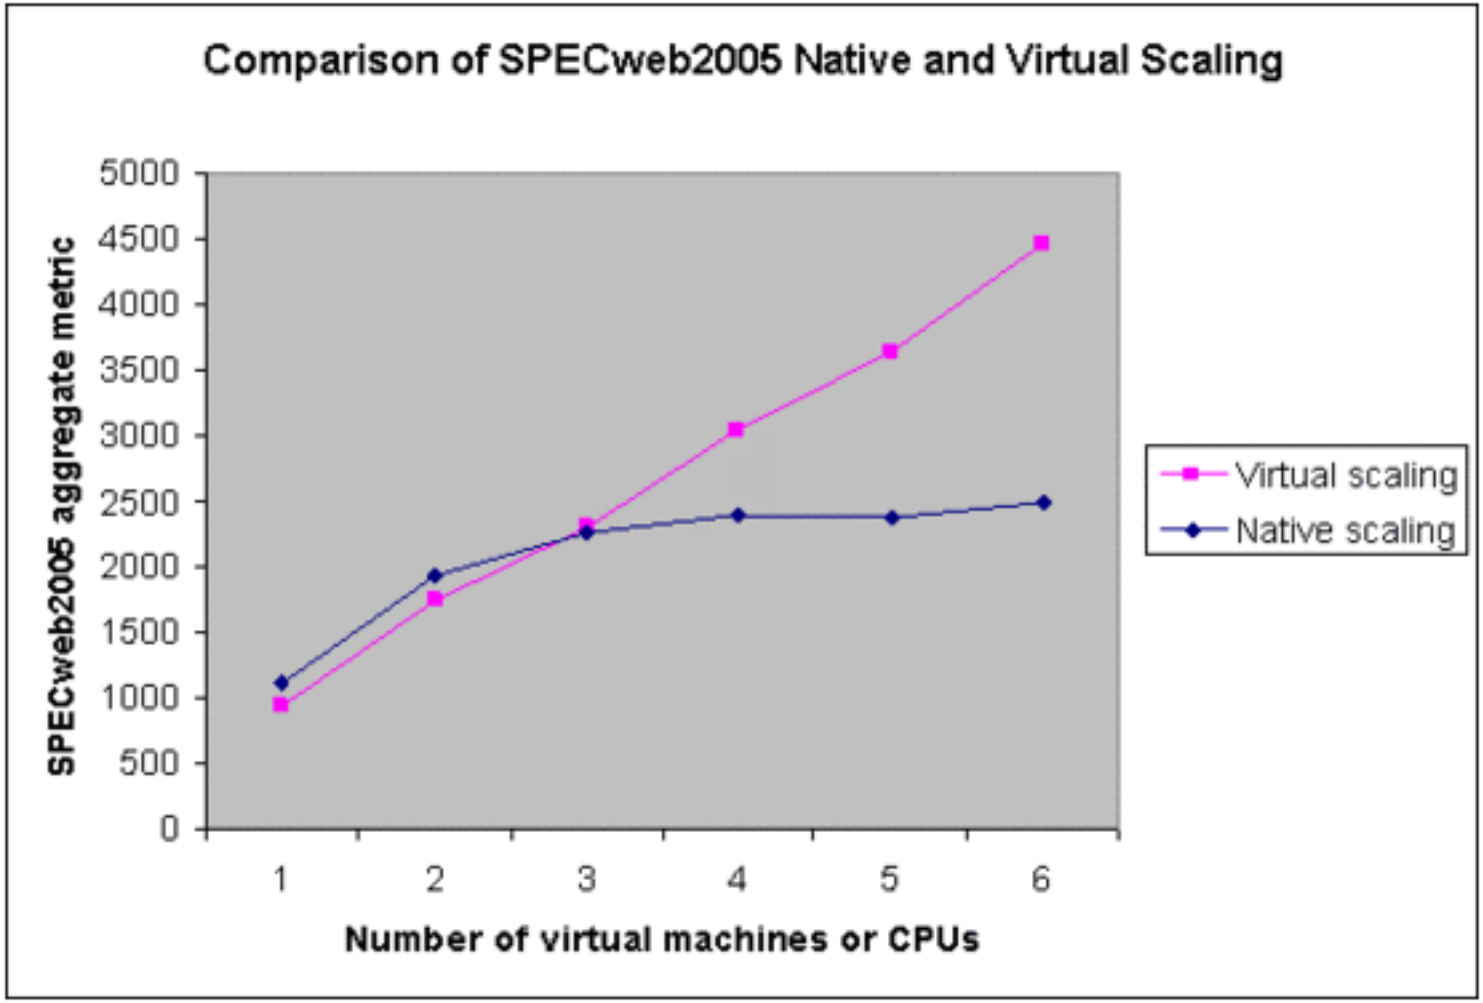

MS Exchange VM Scaling

Single Host Physical Throughput ~8000 Mailboxes (50% of Virtual)

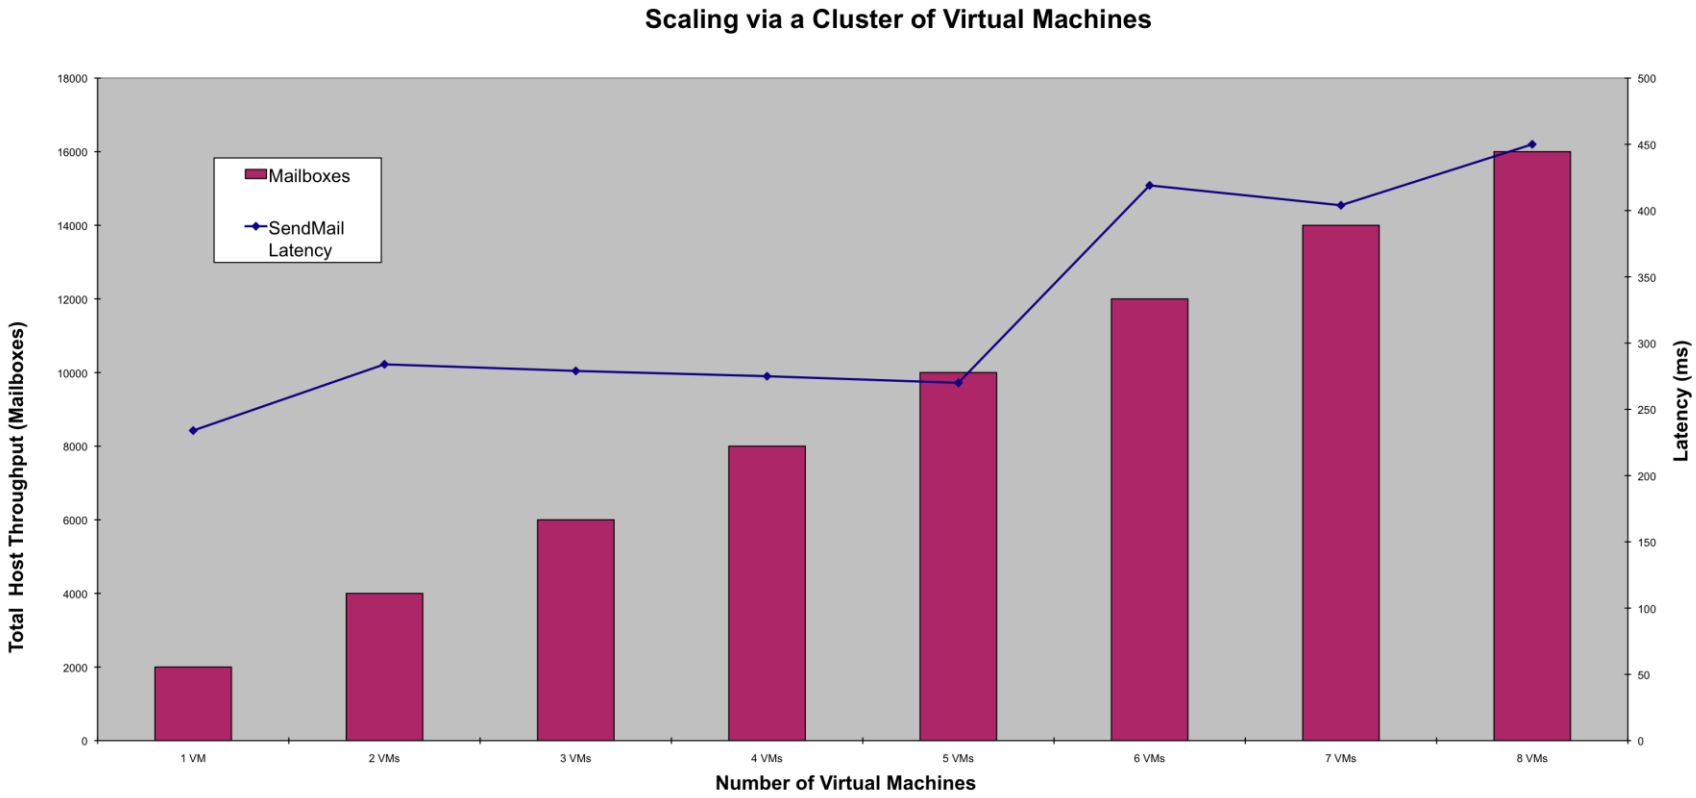

Example migration scenario 4_4_0_0 with DRS

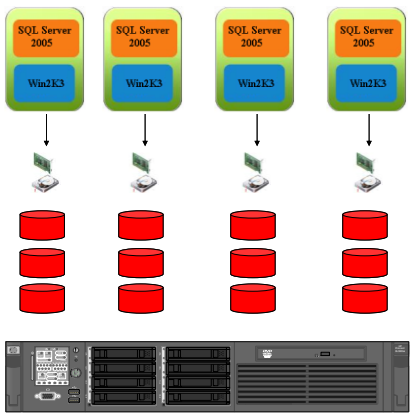

vCenter

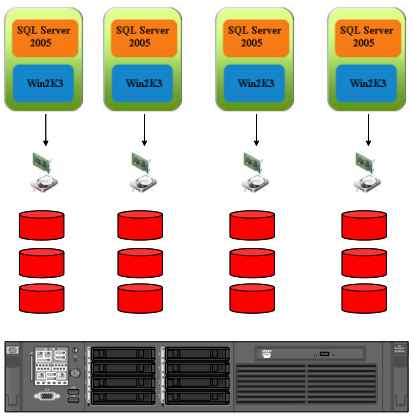

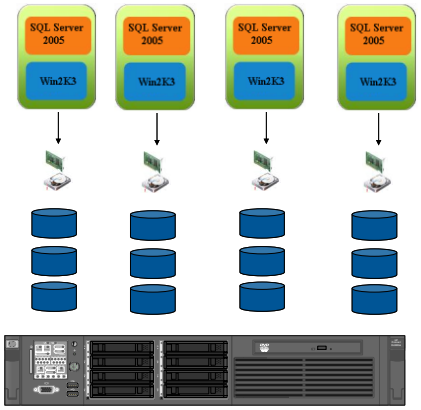

Imbalanced Cluster

Balanced Cluster

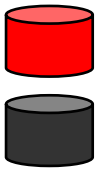

Heavy Load

Lighter Load