TA30

Scaling Your Virtual Infrastructure for Larger Workloads

Mike DiPetrillo

Principal Systems Engineer

VMware

A Level Set

- This session is more of a general overview.

- For workload specific detailed information see the
- “Workloads and Virtual Appliances” track

Scaling Apps Workflow

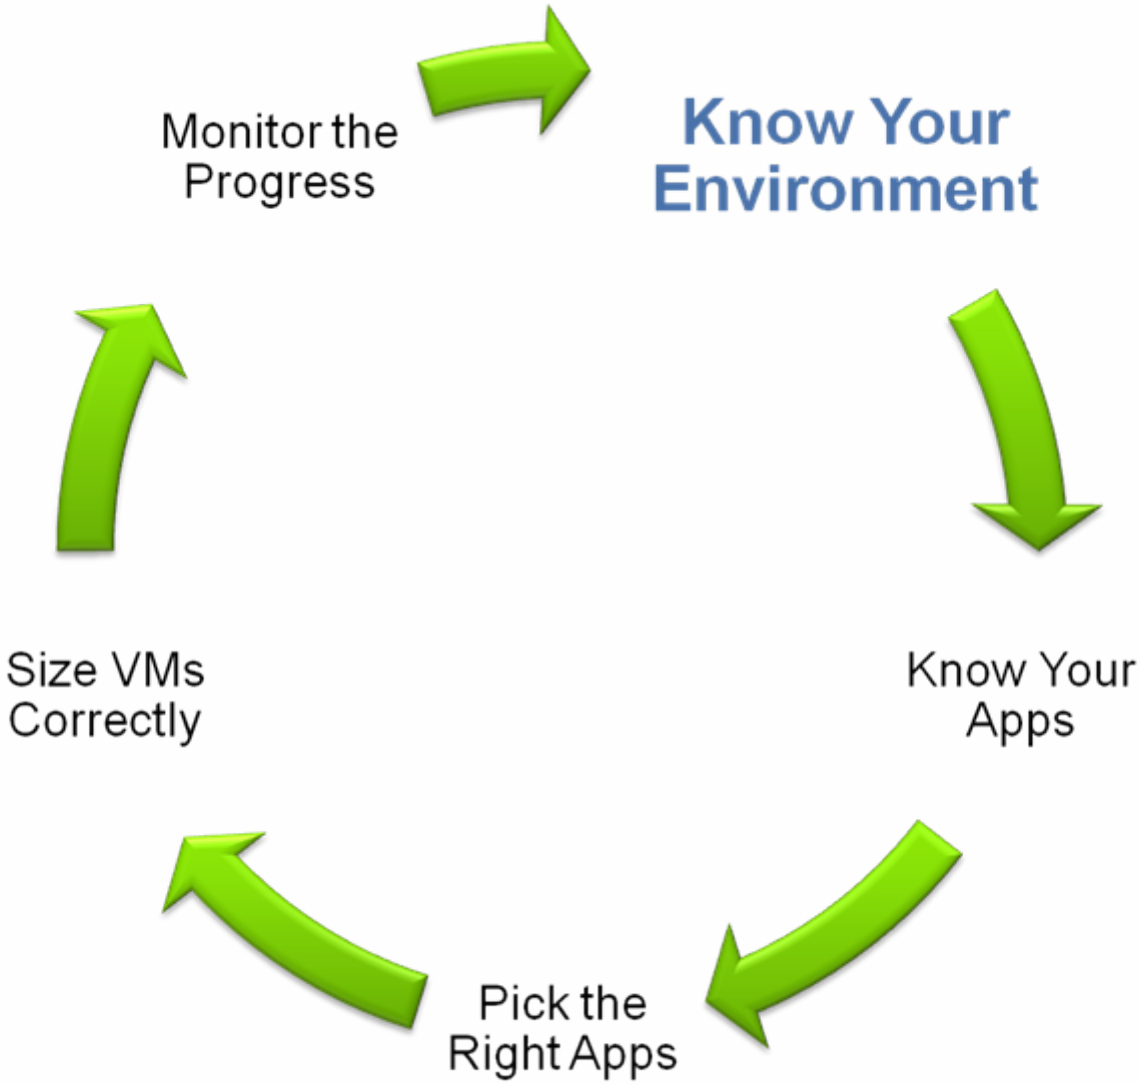
Know Your Environment
Know Your Apps
Pick the Right Apps
Size VMs Correctly
Monitor the Progress

Virtual Architecture

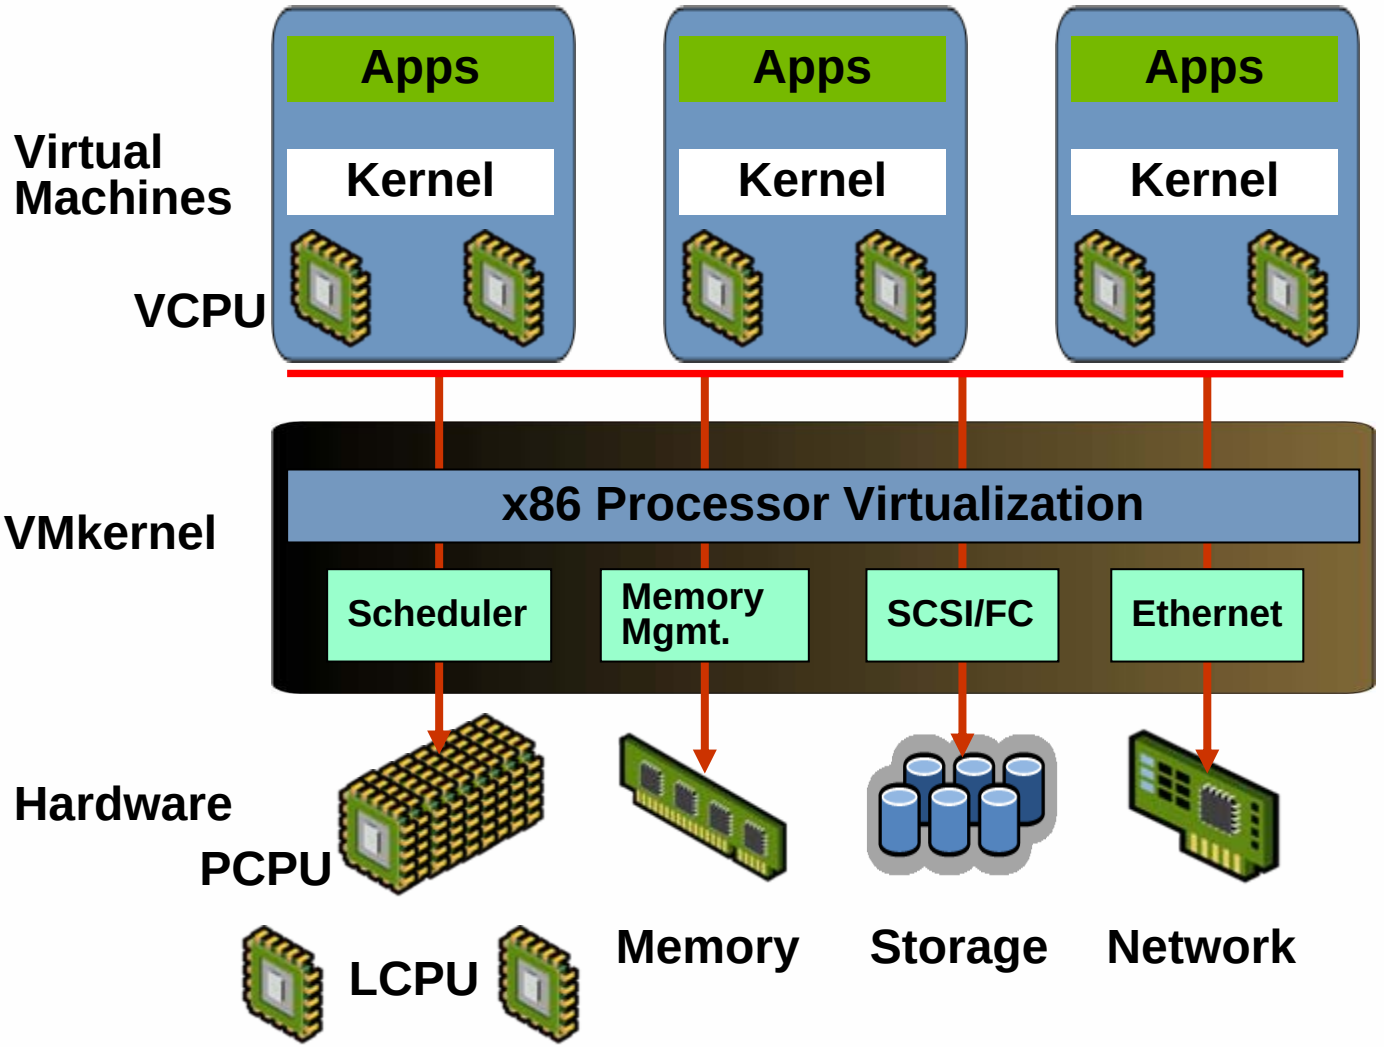

Potential Bottlenecks in Virtualization

- **There are several places within the virtual infrastructure where bottlenecks can occur.**
- **Typically bottlenecks only occur when VMs are oversubscribed.**
  - Too many VCPUs to PCPUs
  - Too much memory allocated versus physical memory
  - Too much network throughput with not enough NICs

Look beyond the virtualization layer for bottlenecks

Often times a network bottleneck in a VM is really related to the downstream receiving server's slow disks.

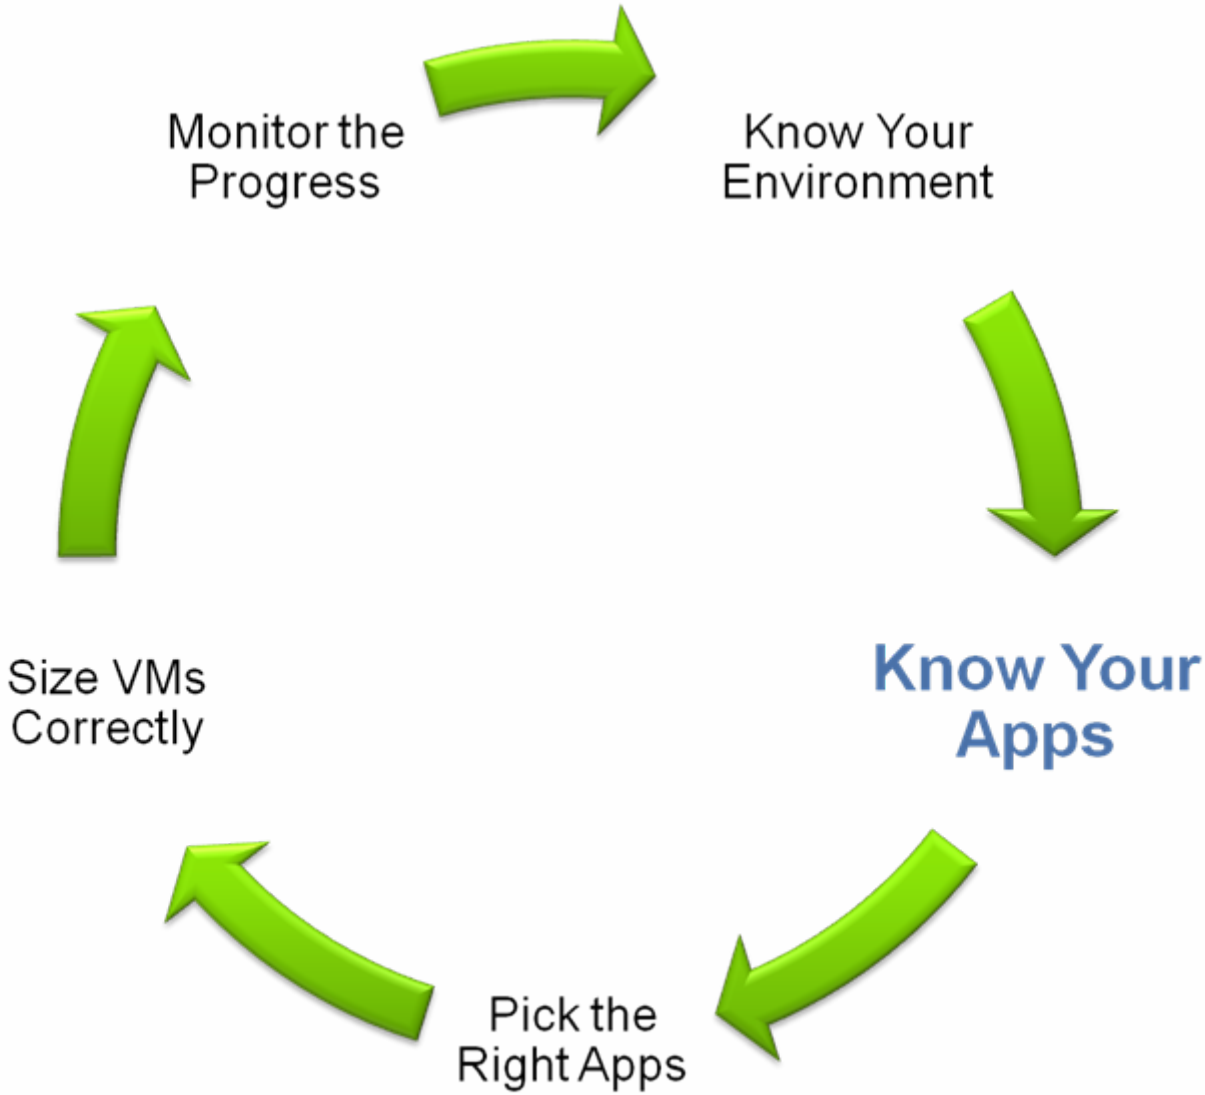
Know Your Environment
Know Your Apps
Pick the Right Apps
Size VMs Correctly
Monitor the Progress

Look at Performance Before you Start

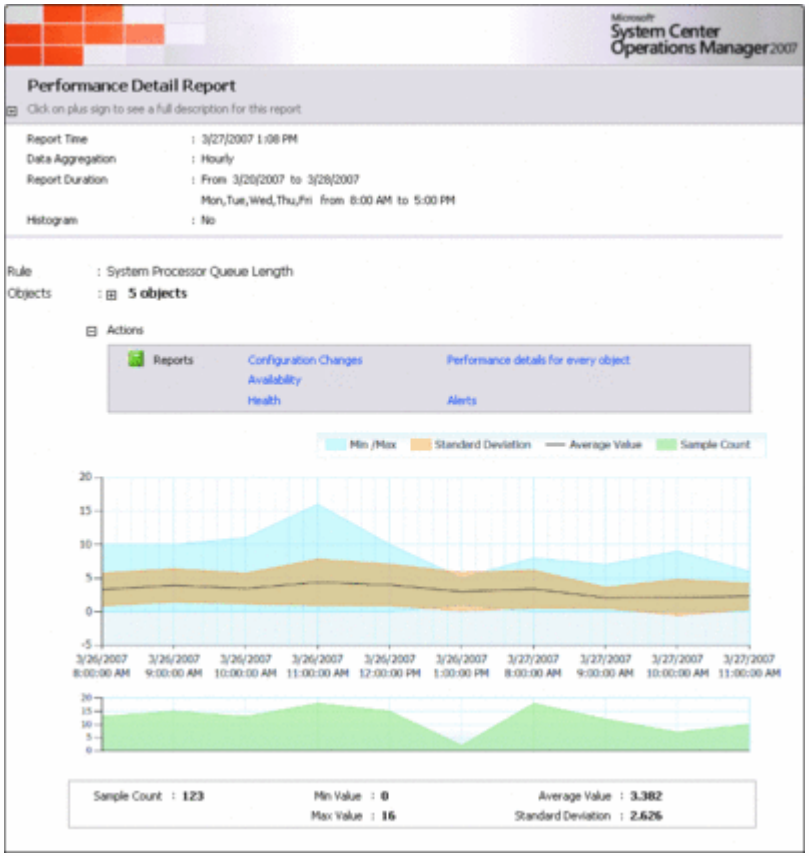

Use VMware Capacity Planner or other 3rd party solutions

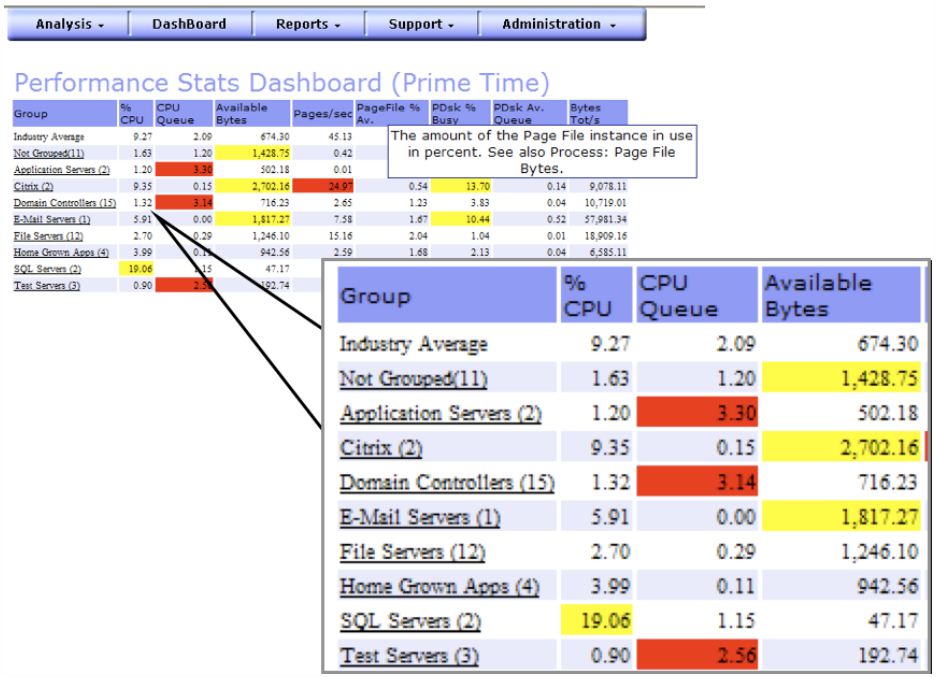


Monitor the Progress
Know Your Environment
Know Your Apps
Pick the Right Apps
Size VMs Correctly

Citrix – Good and Bad

- **Bad** Citrix Apps for VMs**
  - Well behaved Citrix Apps
    - Get 80 – 100 users on your normal physical system (2-way)
  - Traditional apps
    - Microsoft Office
- **Good Citrix Apps for VMs**
  - Kernel bound or CPU bound applications
    - Medical, finance, home grown
    - Get 10 or 20 users on your normal physical system (2-way)

** Bad does not mean they won't run – it simply means you don't get any more benefits for performance.

Citrix - Kernel Bound and CPU Bound

- Windows kernel only has access to so much memory (4 GB)
- Typical applications written to run on single PC environment
- When application run in Citrix it runs out of kernel memory
- Adding more memory to the server does not help since the kernel can still only access 4 GB

Scaling Exchange

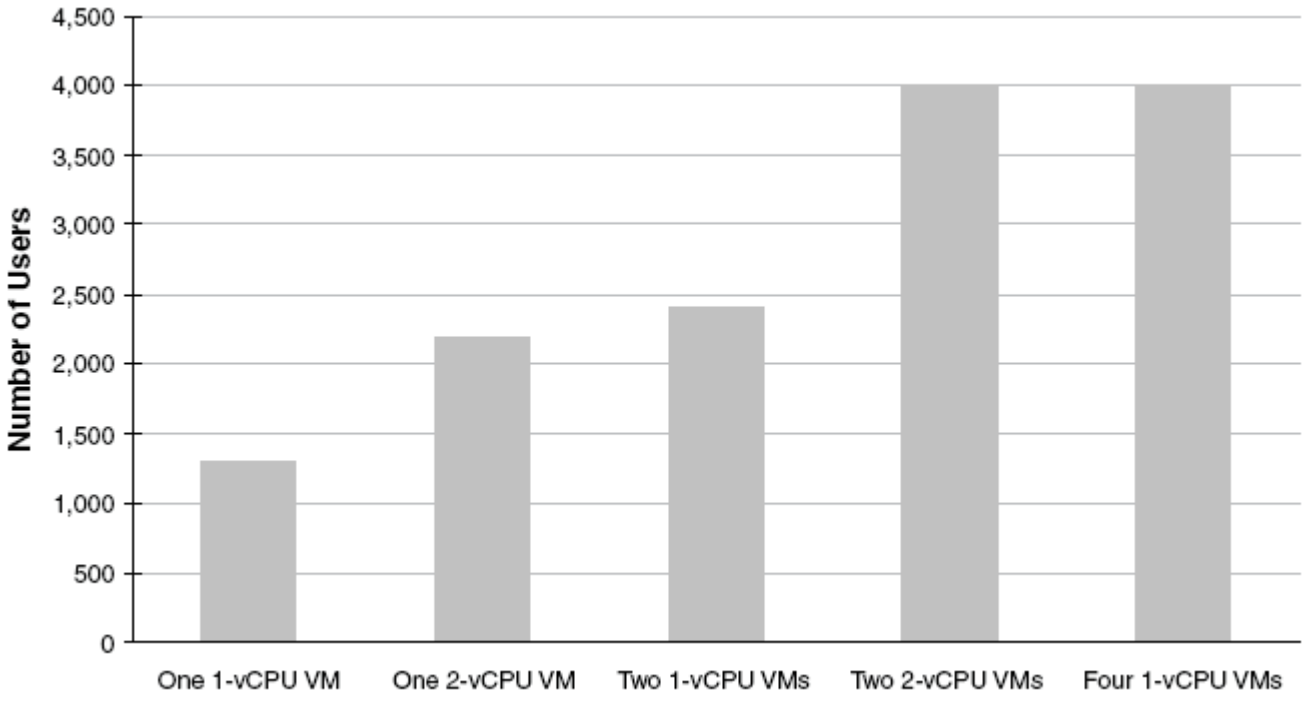

http://www.vmware.com/resources/techresources/978

Scaling Exchange

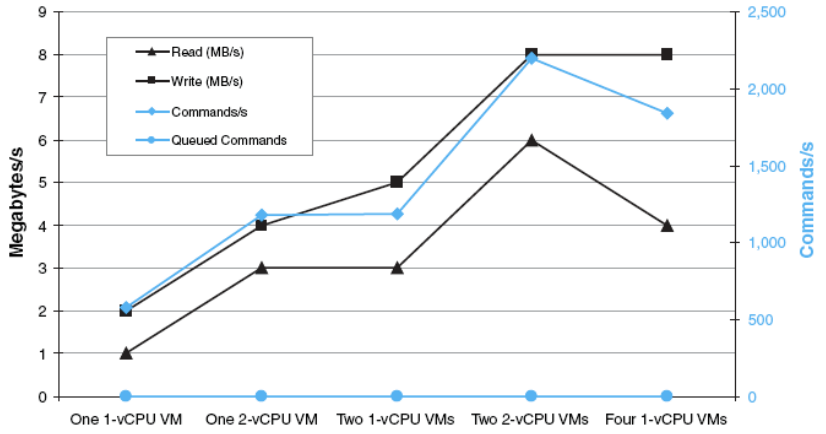

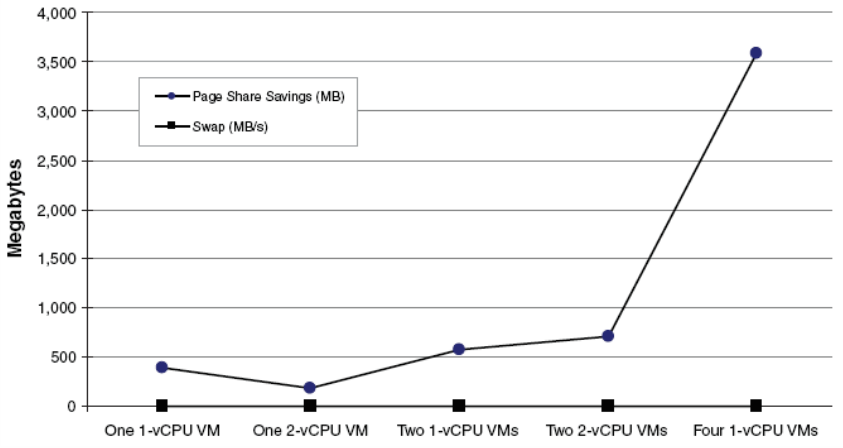

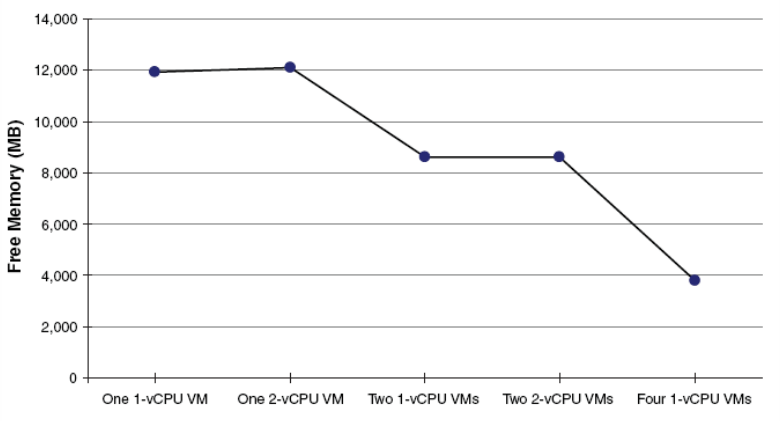

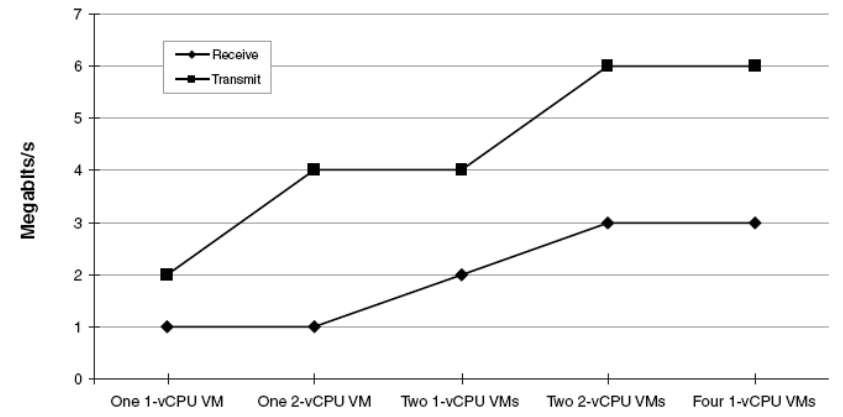

http://www.vmware.com/resources/techresources/978

Scaling Websphere

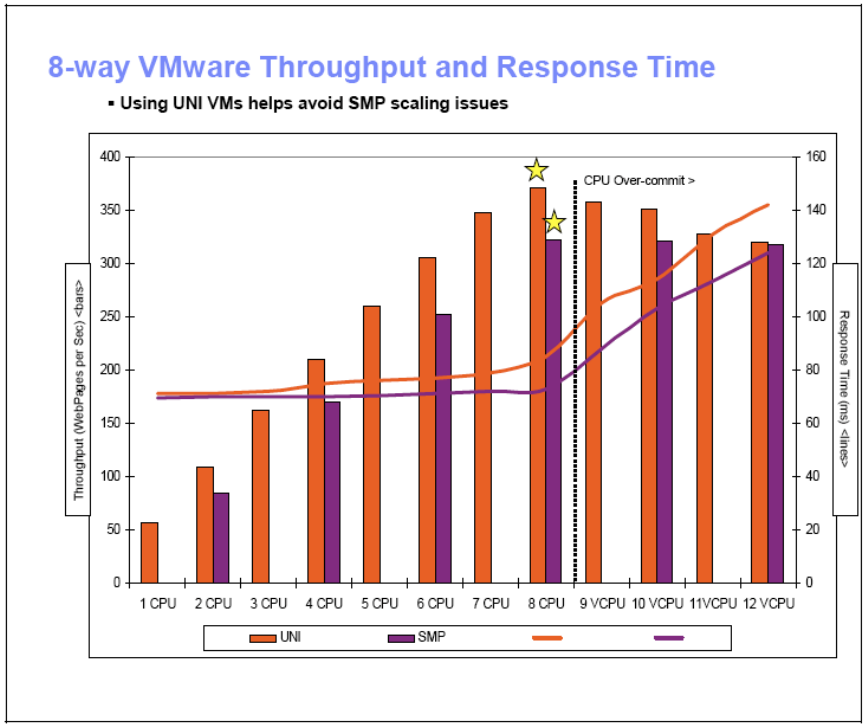

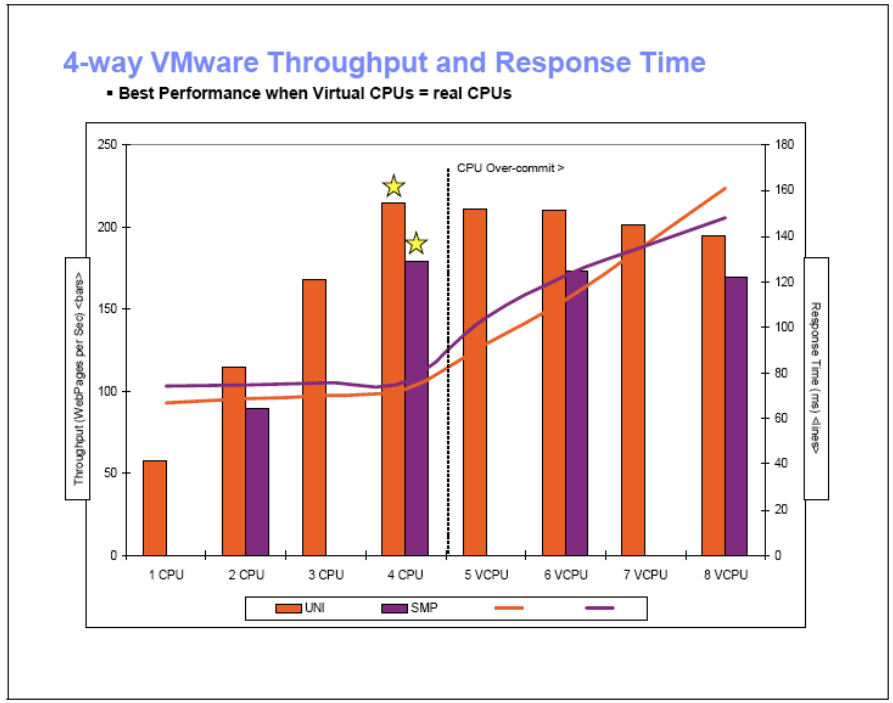

http://www.vmware.com/resources/techresources/527

Scaling Websphere

Best Throughput Comparisons

Note: Same Response Times Comparisons, Not Optimal Throughput Configurations

- ESX Server demonstrates a stable platform and good performance

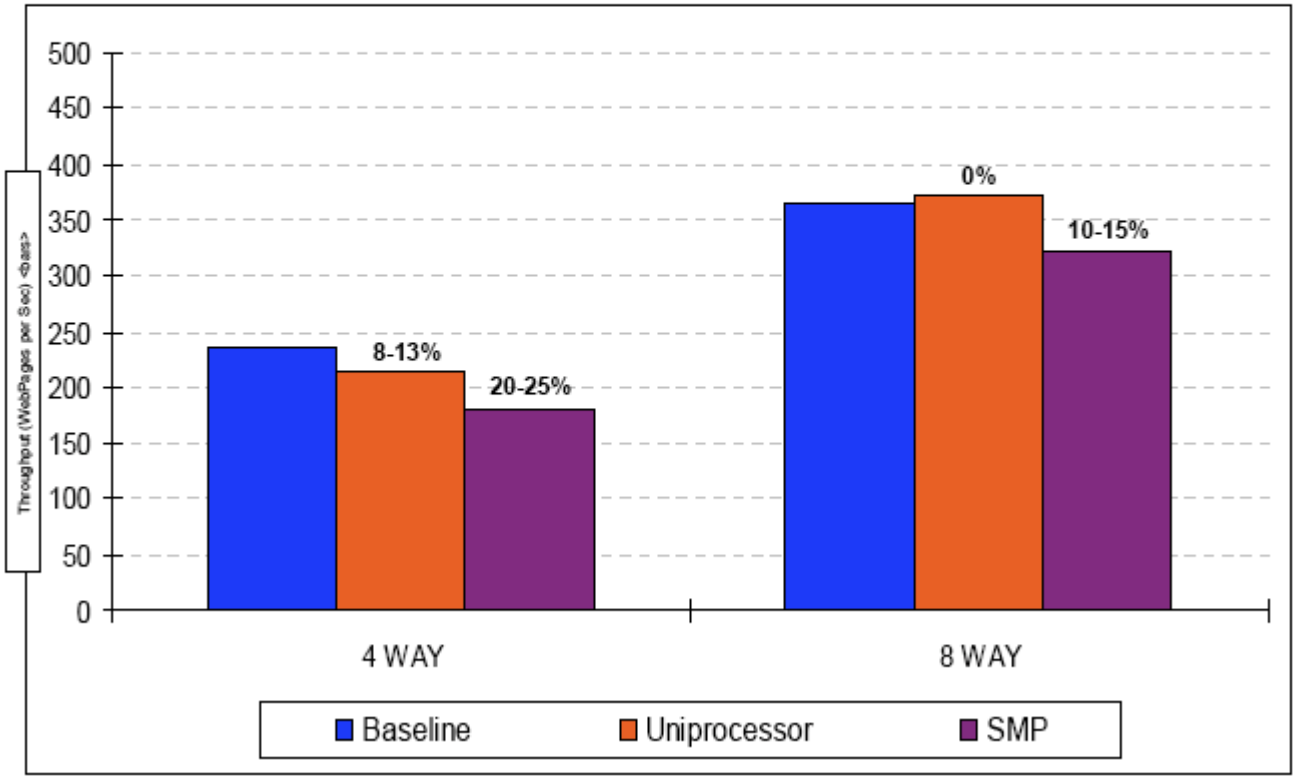

http://www.vmware.com/resources/techresources/527

Scaling Databases – A Study of Throughput

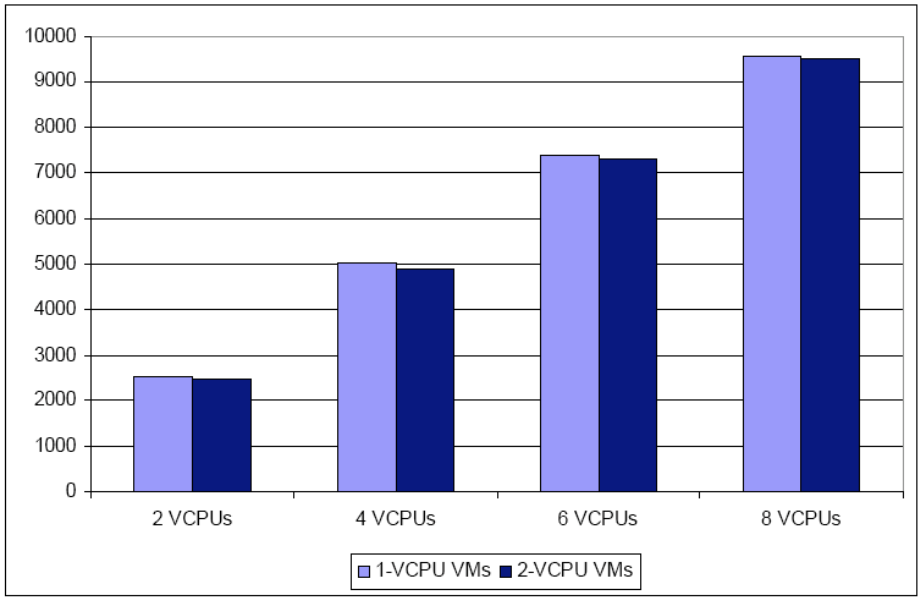

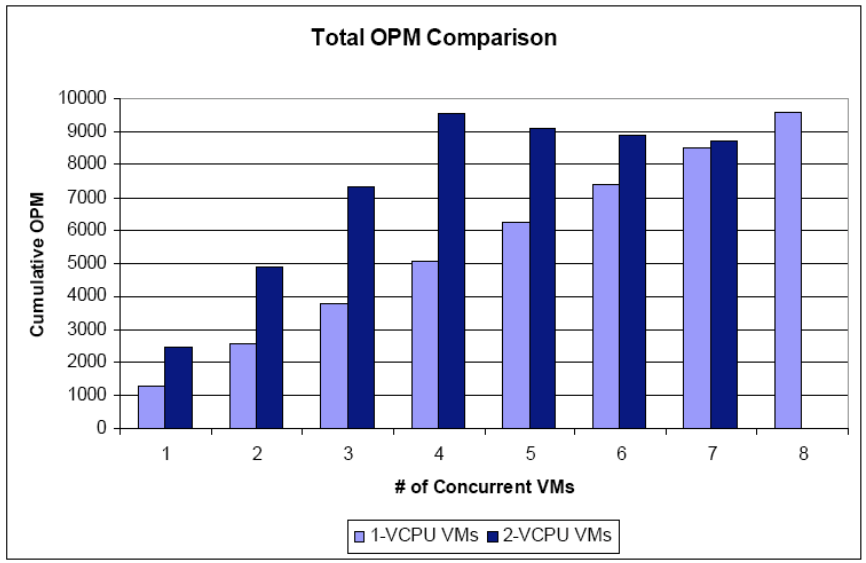

http://www.vmware.com/resources/techresources/643

Scaling Databases – A Study of Throughput

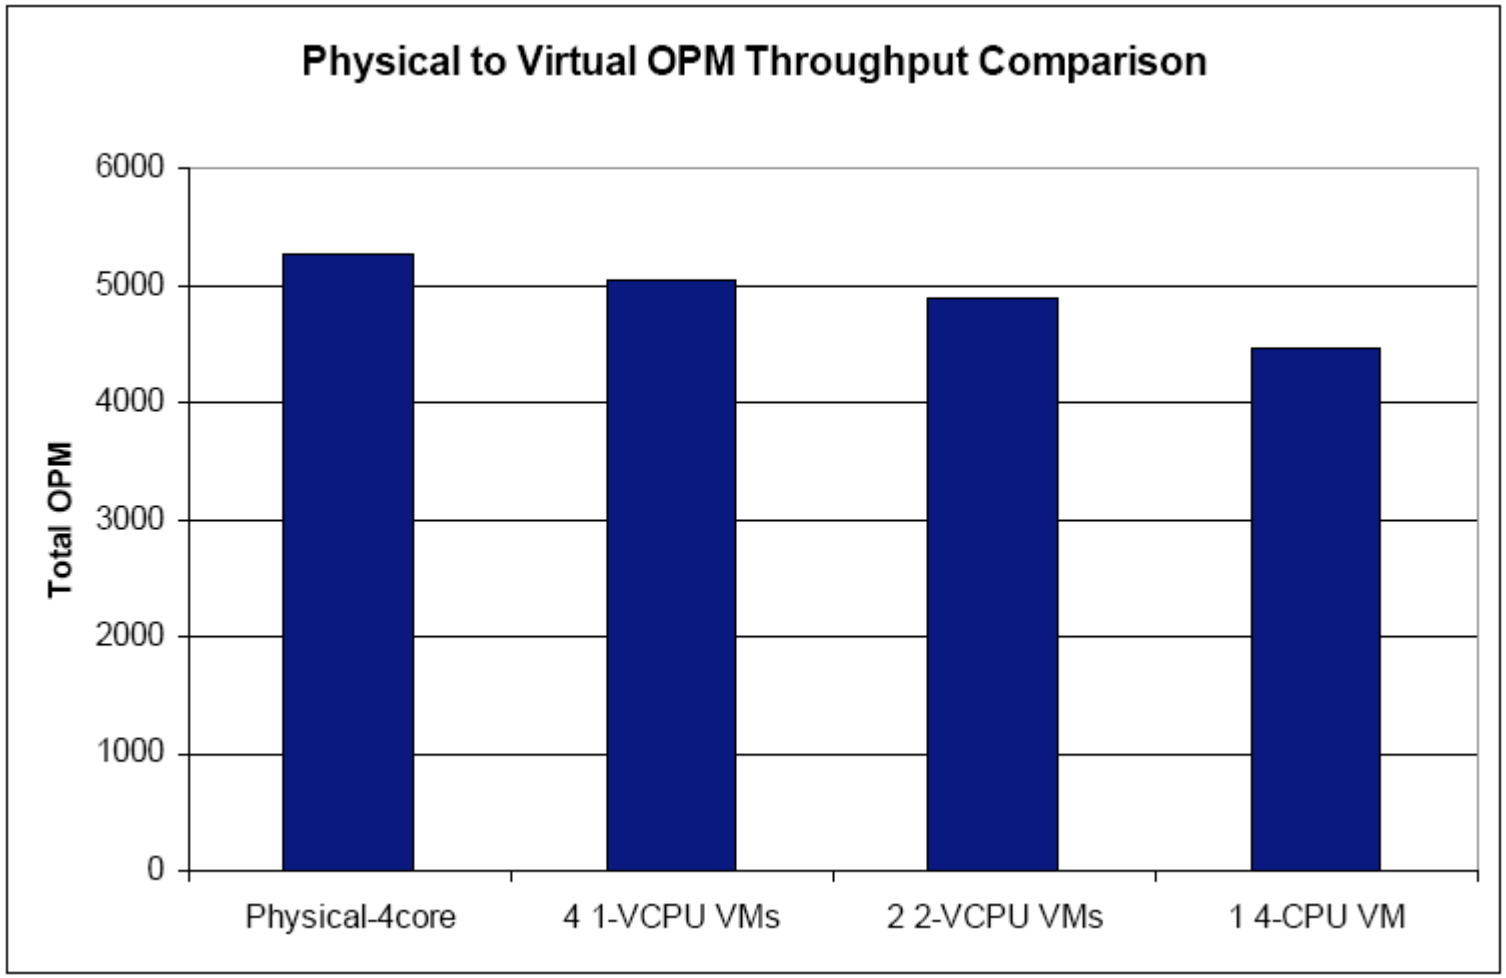

http://www.vmware.com/resources/techresources/643


Know Your Environment
Know Your Apps
Pick the Right Apps
Size VMs Correctly
Monitor the Progress

Use SMP Sparingly

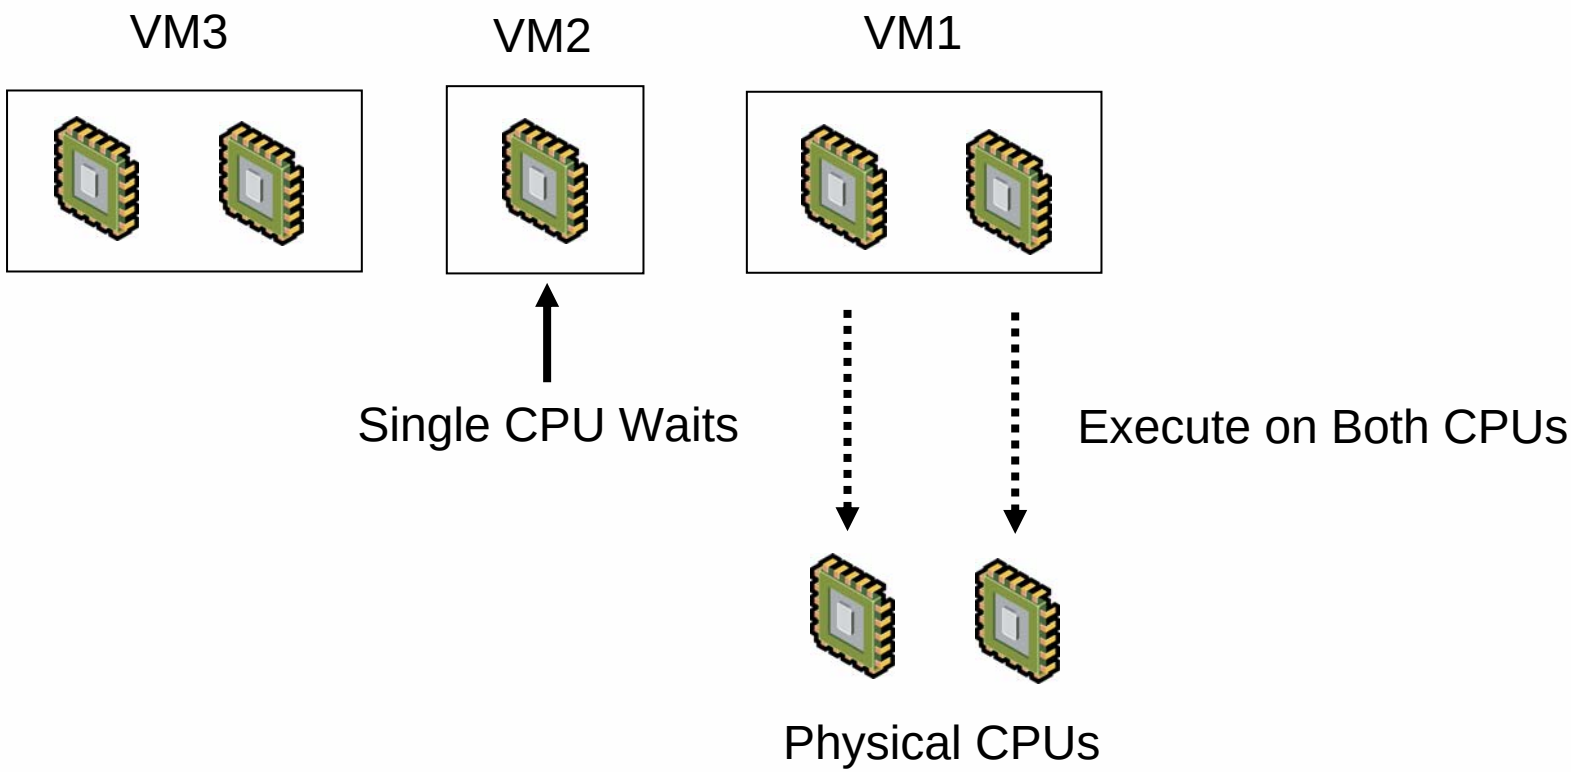

Use SMP Sparingly

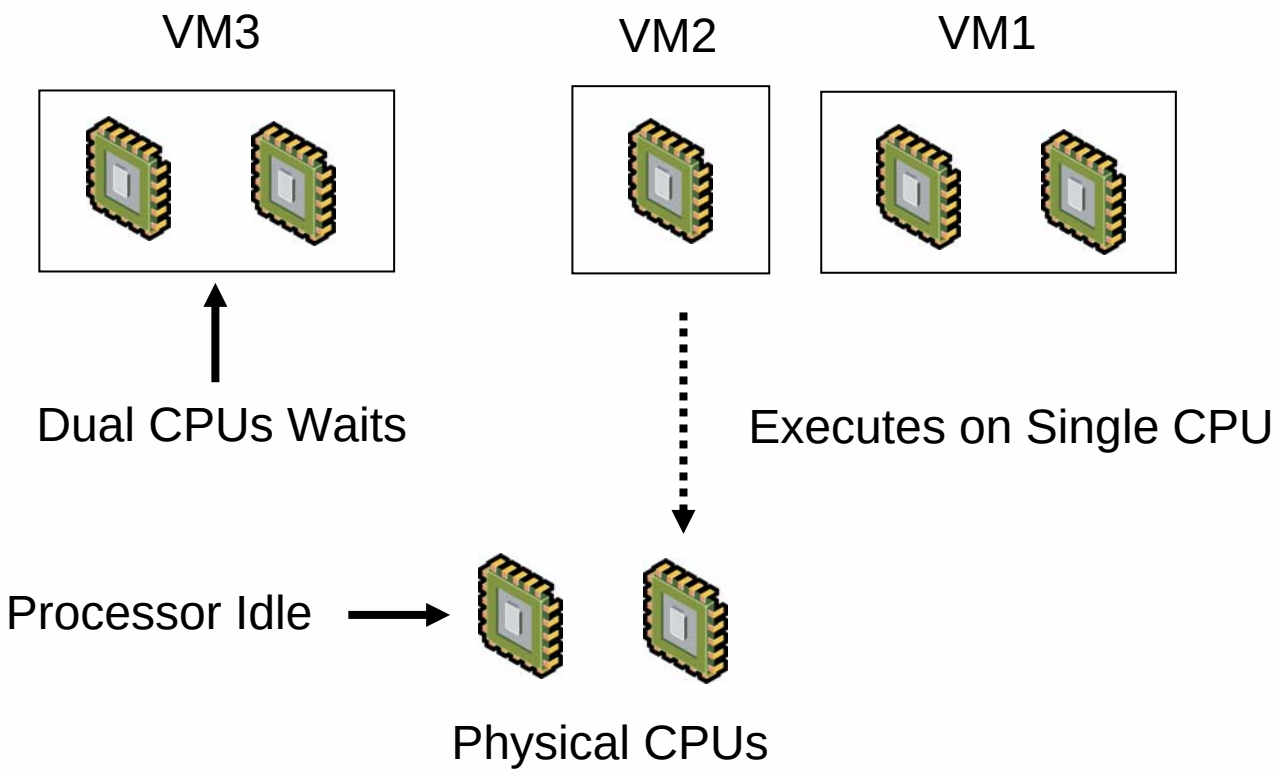

Use SMP Sparingly

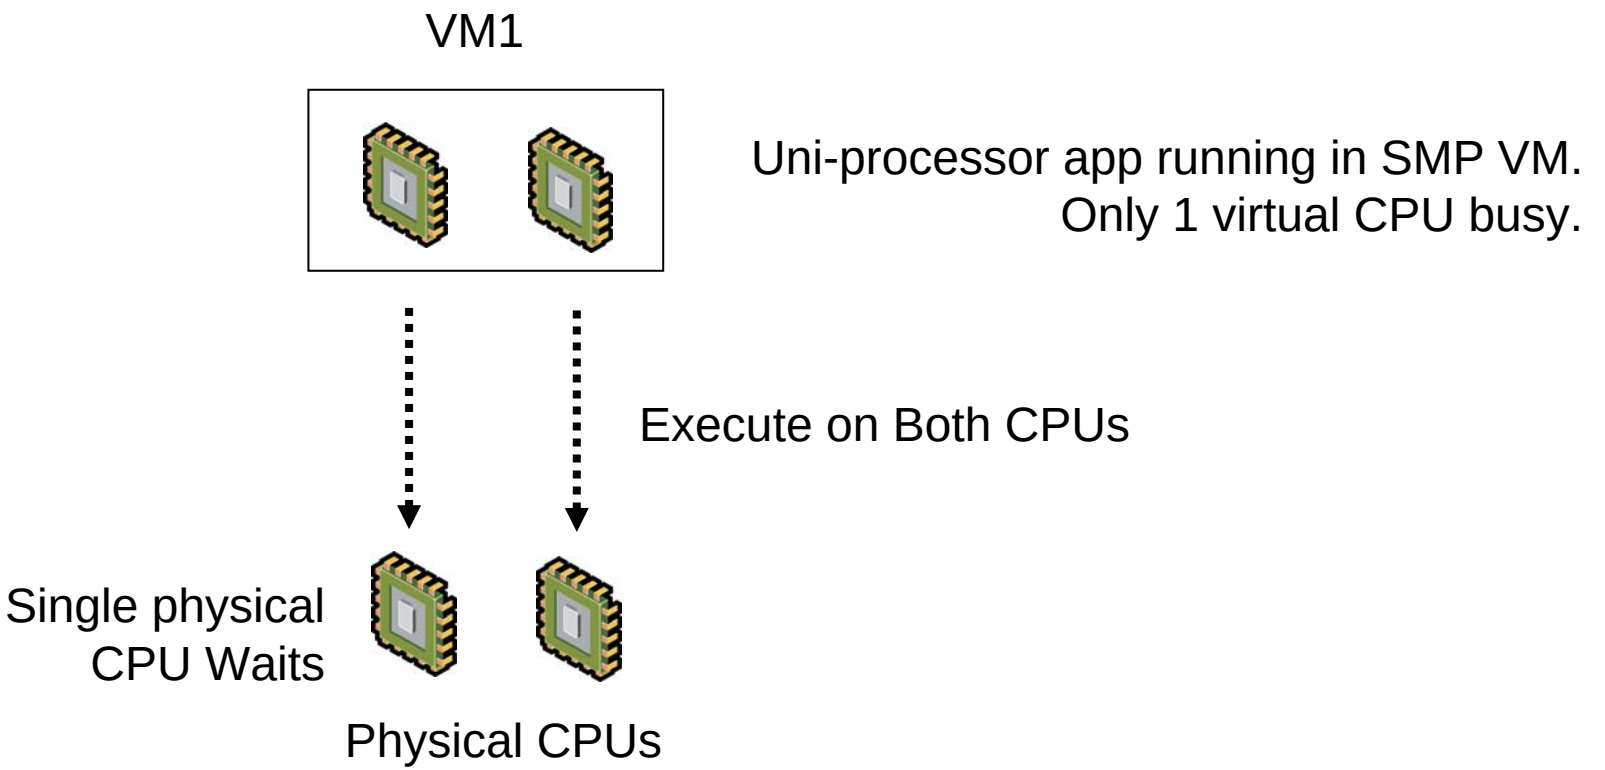

Use 64-bit Sparingly

An example using Citrix

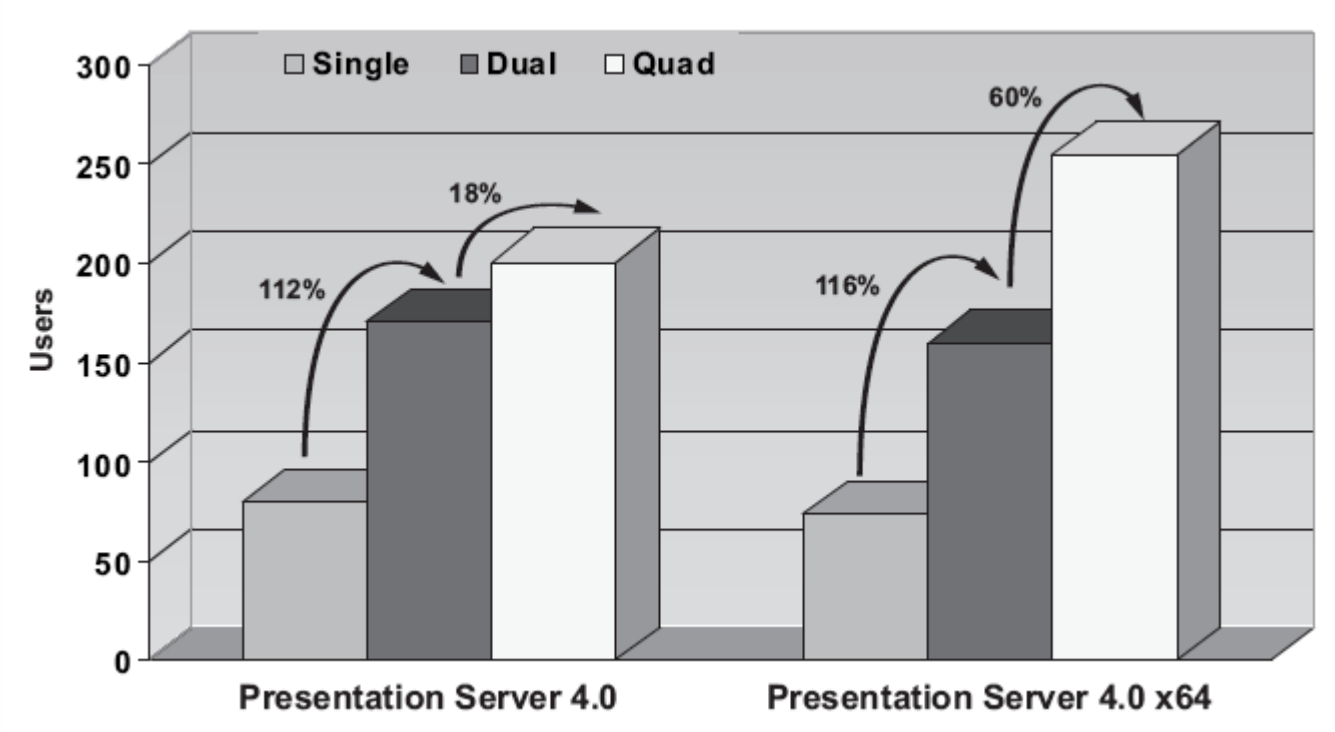

http://www.citrix.com/site/resources/dynamic/salesdocs/CPS4_x64_Performance_and_Scaling_Capabilities.pdf

Other Hints

- **CPU pinning can help you gain a few percentage points**
  - This impacts mobility
- **Monitor memory closely after a VM starts**
  - After a few minutes the VM will “calm down” and show true memory utilization
- **Look for new advances in the Linux kernel (VMI and paravirt-ops), timing adjustments (RHEL kernel), etc**
  - These can offer better performance for CPU hungry applications

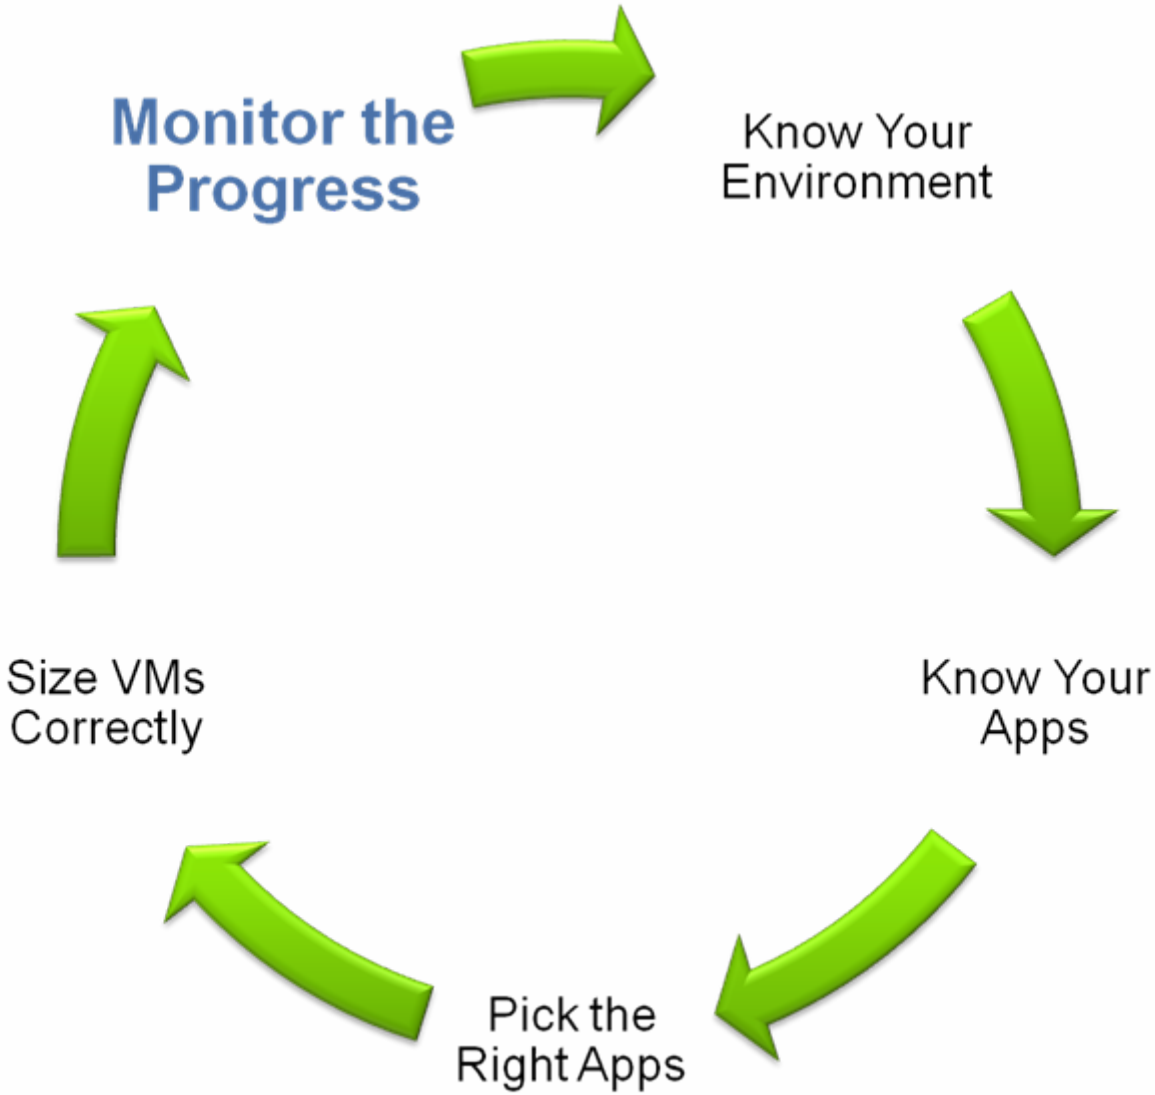
Monitor the Progress
Know Your Environment
Know Your Apps
Pick the Right Apps
Size VMs Correctly

Monitoring the Virtual World

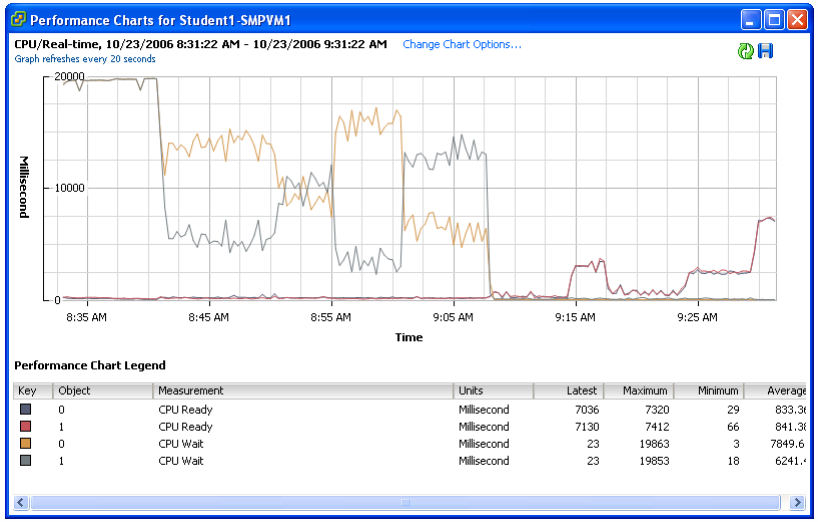

Use VMware Virtual Center to monitor the performance of the virtual machines.

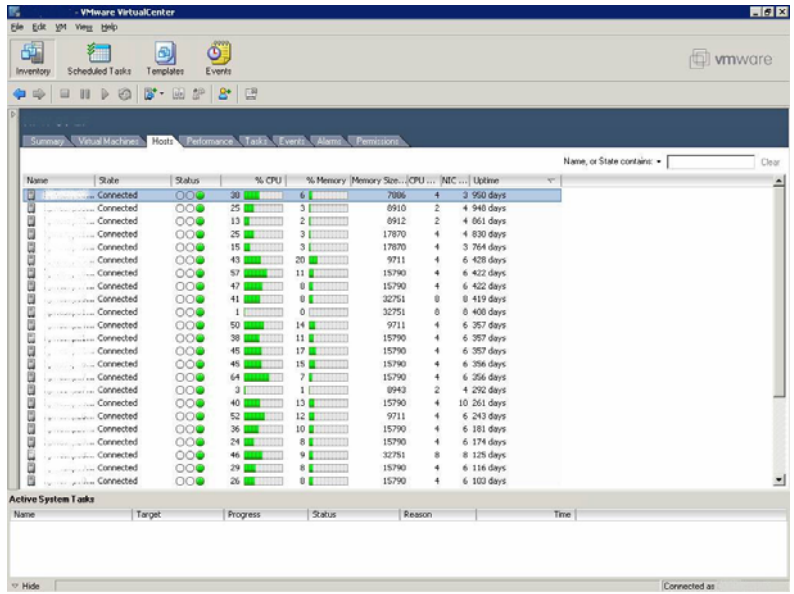

Monitoring the Physical World

Pay close attention to the relationship between physical and virtual

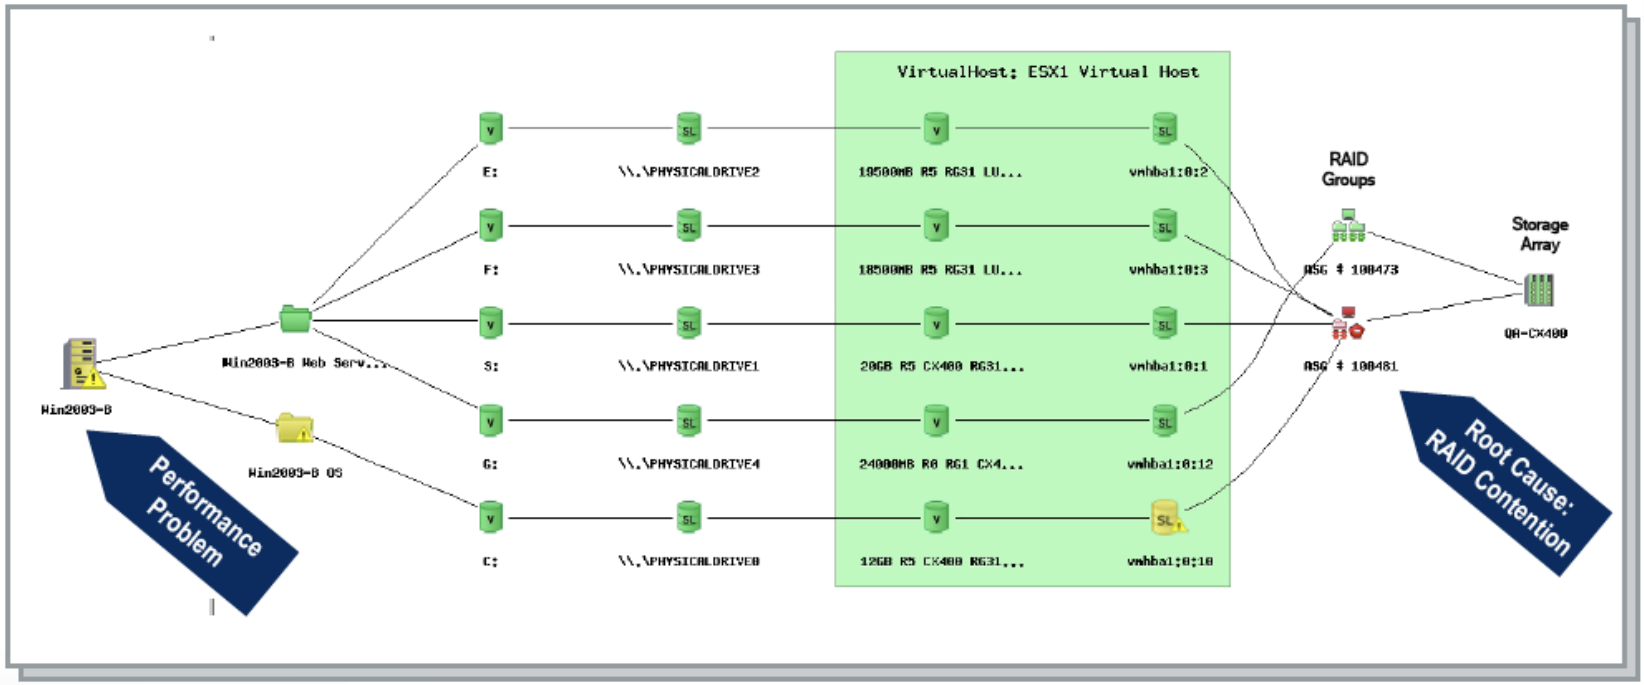

Scaling Apps Workflow

Questions?

TA30

Scaling Your Virtual Infrastructure for Larger Workloads

Mike DiPetrillo

VMware

For more information …

More workload sessions in the “Workloads and Virtual Appliances” track

Questions?

TA30

Scaling Your Virtual Infrastructure for Larger Workloads

Mike DiPetrillo

VMware

For more information …

More workload sessions in the “Workloads and Virtual Appliances” track

VMWORLD 2007

EMBRACING YOUR VIRTUAL WORLD

BREAKOUT SESSION