Virtualization : IBM Insights in Sizing Servers for Virtualization

Bob Zuber

WW Program Manager for Virtualization

VMWORLD 2006

Learning Objectives

- Title: Virtualization - IBM Insights in Sizing Servers for Virtualization

- Abstract: Come hear how IBM can help you build a better virtualized infrastructure. Learn how to use the new IBM Virtualization Sizing Guide, to help you understand the method and tools that IBM has used to assist their customers in many consolidation engagements. In addition, learn how IBM System x and Bladecenter servers can be deployed to provide the most efficient and reliable hardware platform for your VMware virtualization needs. Also see how IBM Director and Virtualization Manager provide a complete systems management solution, via a single pane of glass interface, for your virtual and physical servers

Agenda

- Sizing
  - The elements of sizing
  - IBM's insights into sizing
- IBM Portfolio
  - Portfolio Positioning
  - Competitive Advantage
- Systems Management – IBM Director
  - What is it
  - How it helps virtualization
- Summary

Long-term Focus On Virtualization Across Our Systems

System z9™
zSeries®

System p5 ™
pSeries®

iSeries ™

System x™ BladeCenter™

IBM System
Storage ™

OpenPower®

While virtualization sounds complex, it's really a simple idea. IBM Systems can provide virtualization capabilities that are unique in the marketplace.

- IBM mainframe virtualization – 40 year history of world-class innovation
- IBM X-Architecture designed for virtualization, shared cross platform
- CoolBlue™ - Power and Cooling designs that lead the industry
- Virtualization Management software that simplifies your environment
- Virtualization features do not require "rip and replace" upgrades
- VMware's number one OEM vendor(1)

Note (1) : Based on bundled license revenue

Why is Sizing Important

- There are so many sizing considerations, either the source or target
- Consider just the Source server
  - Average and Peak CPU utilization
  - Average and Peak memory utilization
  - Page Size and page extents
  - Disk I/O throughput
  - Network I/O throughput
- And now the Target server
  - CPU Utilization
  - Memory Utilization
  - Page size page extents
  - Disk I/O throughput

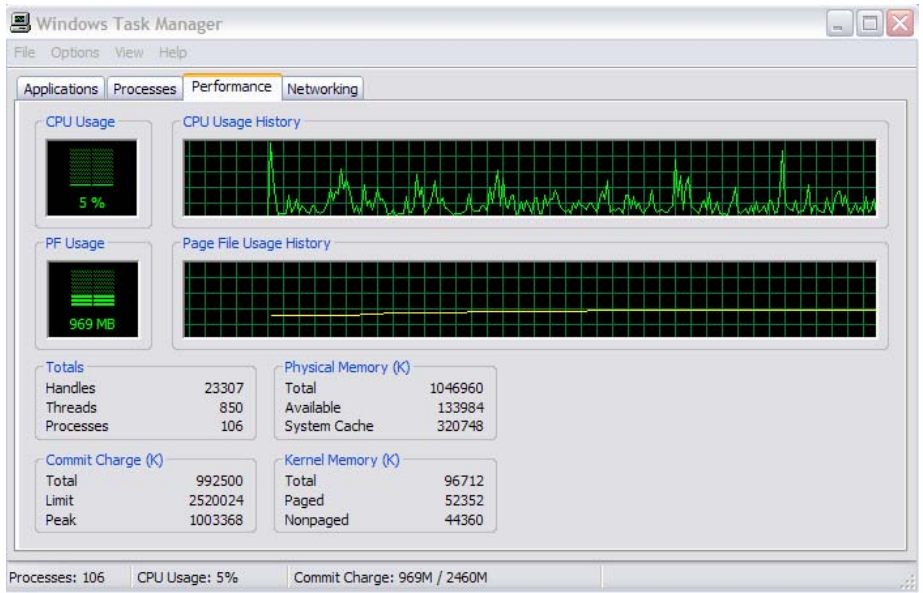

Lego SCON Analysis

Capacity uncertainties can drive costs up

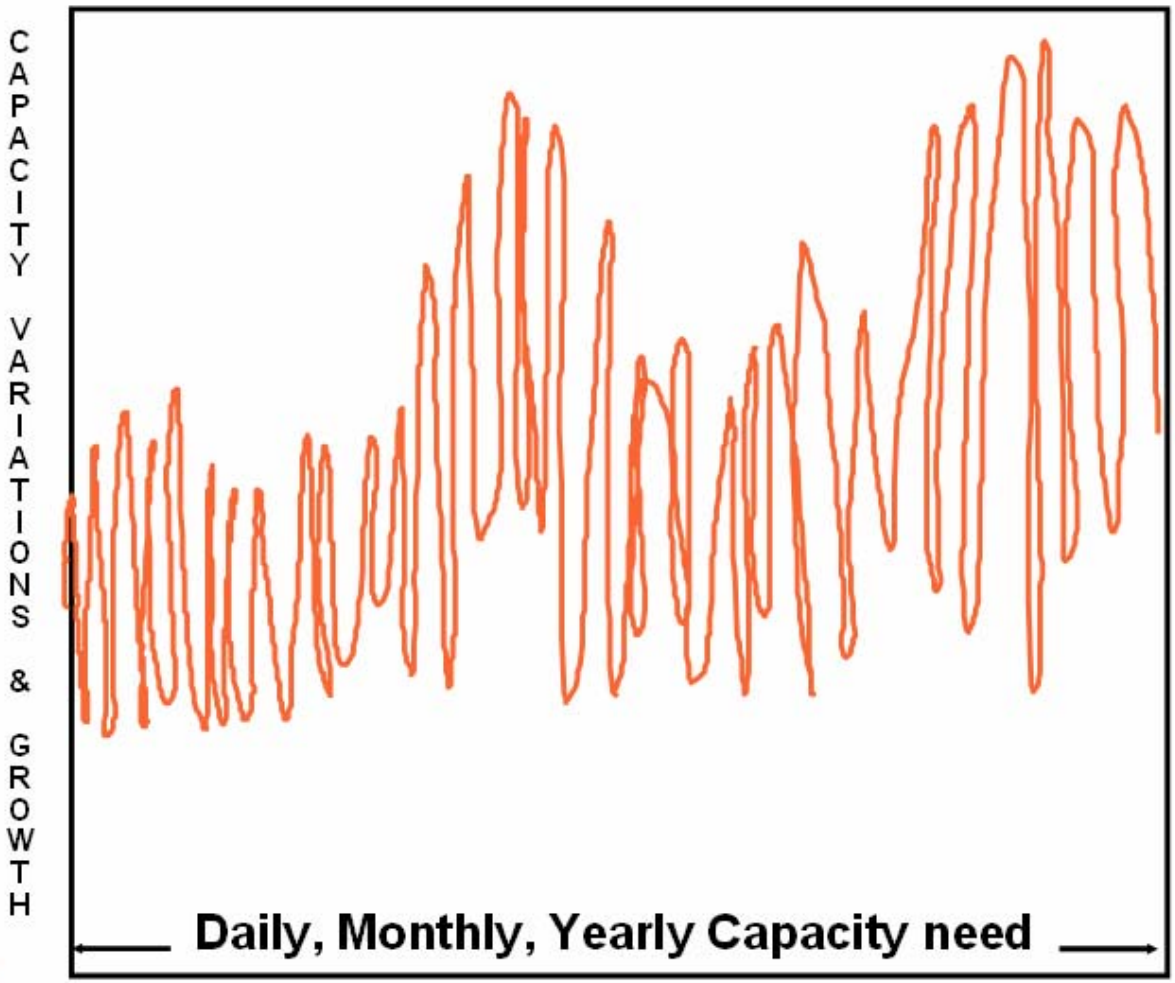

Source: IBM & Lego Joint SCON Study

Source: Lego Customer Study

Lego SCON Analysis

Capacity uncertainties can drive costs up

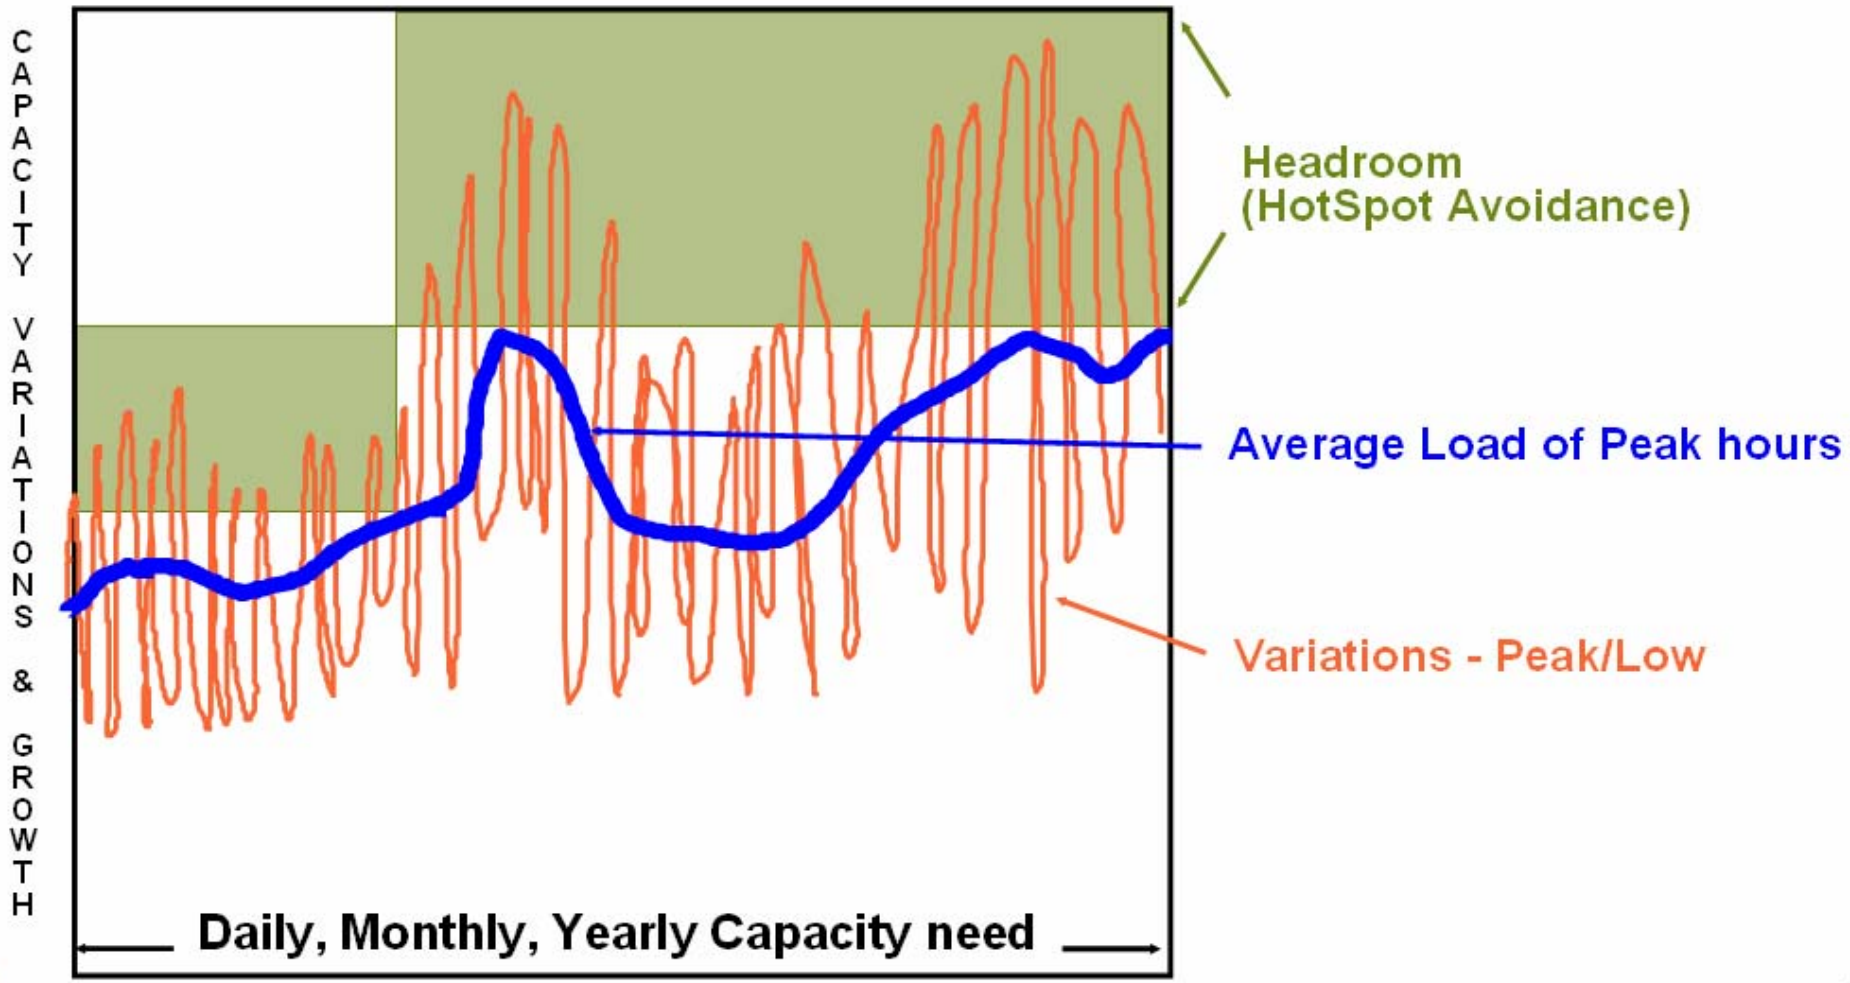

Source: IBM & Lego Joint SCON Study

Source: Lego Customer Study

Lego SCON Analysis

Capacity uncertainties can drive costs up

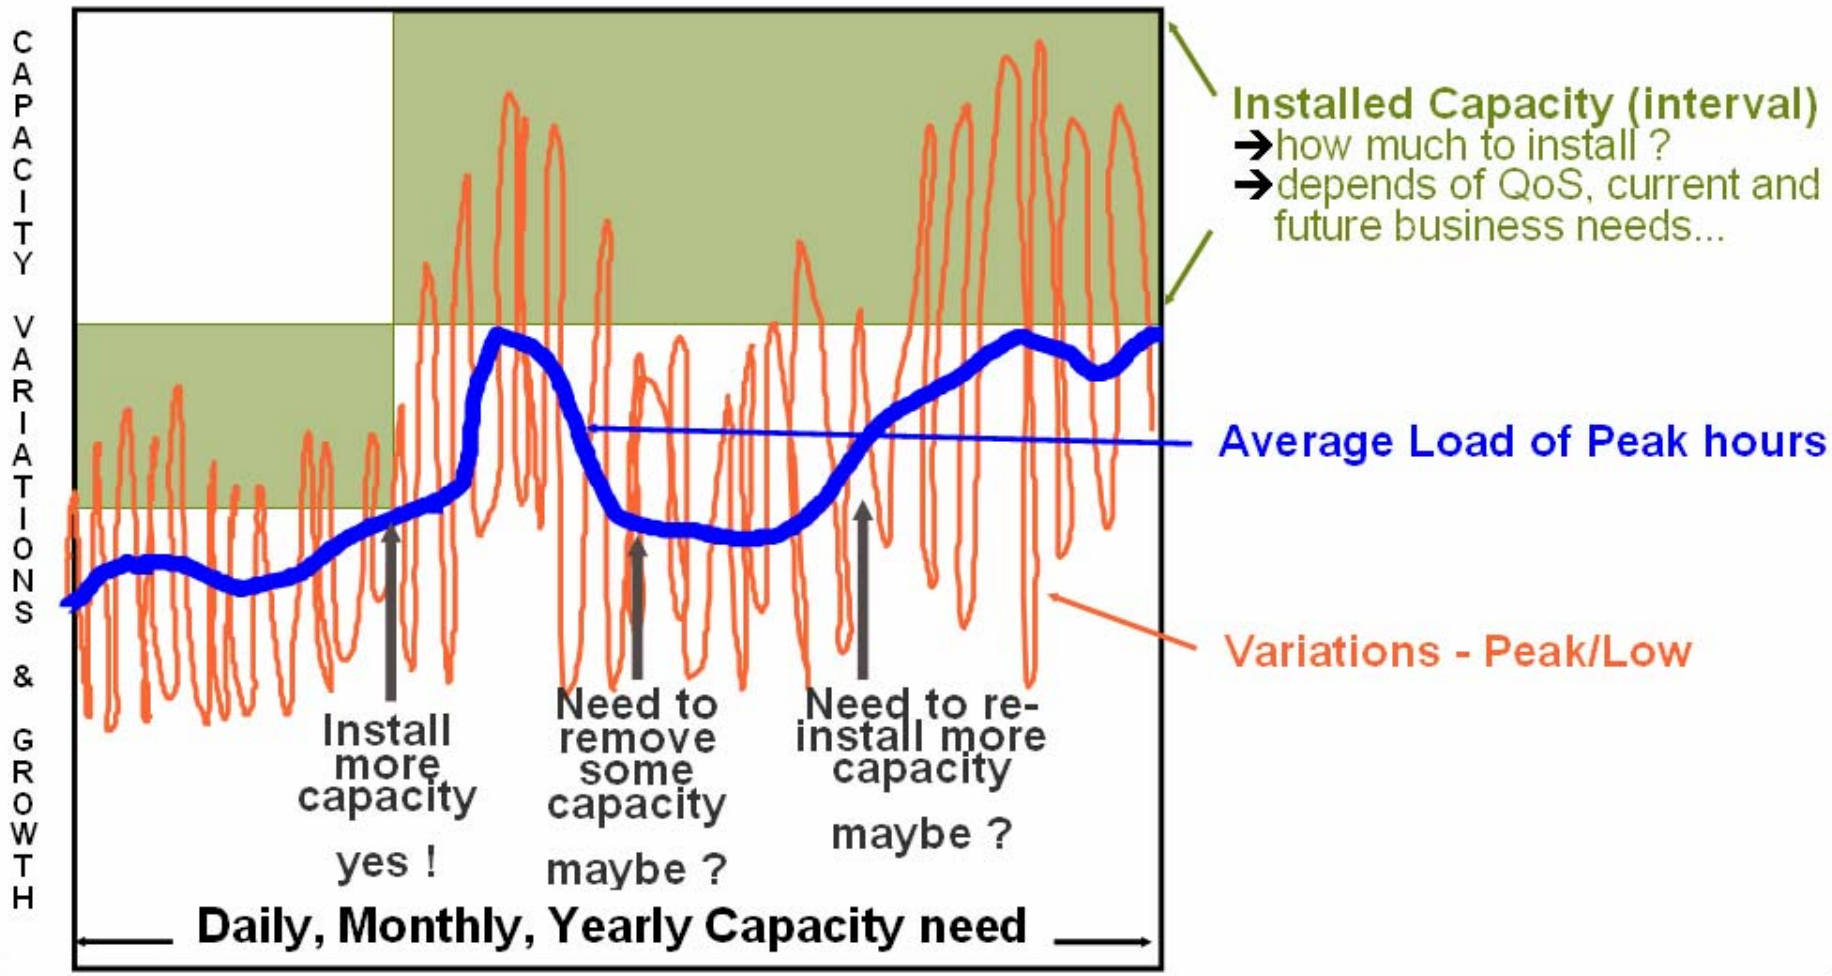

Source: IBM & Lego Joint SCON Study

Source: Lego Customer Study

Lego SCON Analysis

Capacity uncertainties can drive costs up

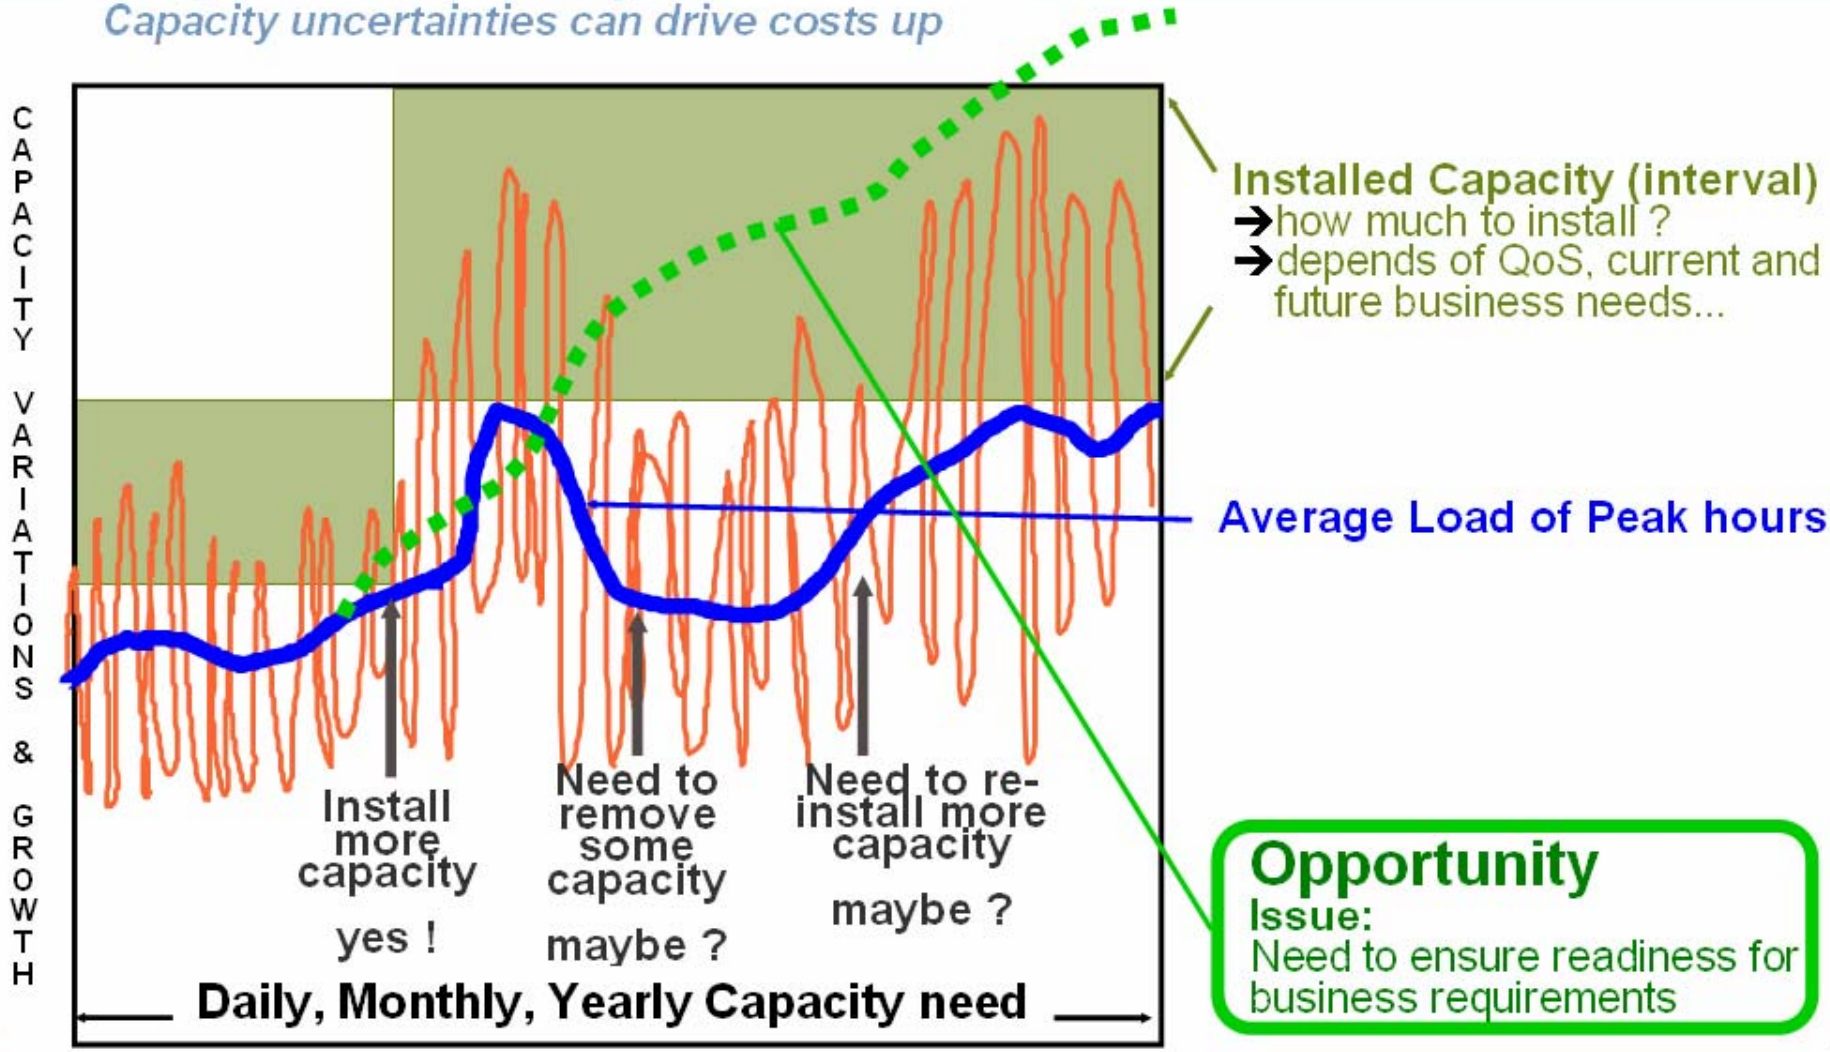

Source: IBM & Lego Joint SCON Study

Source: Lego Customer Study

Lego SCON Analysis

Capacity uncertainties can drive costs up

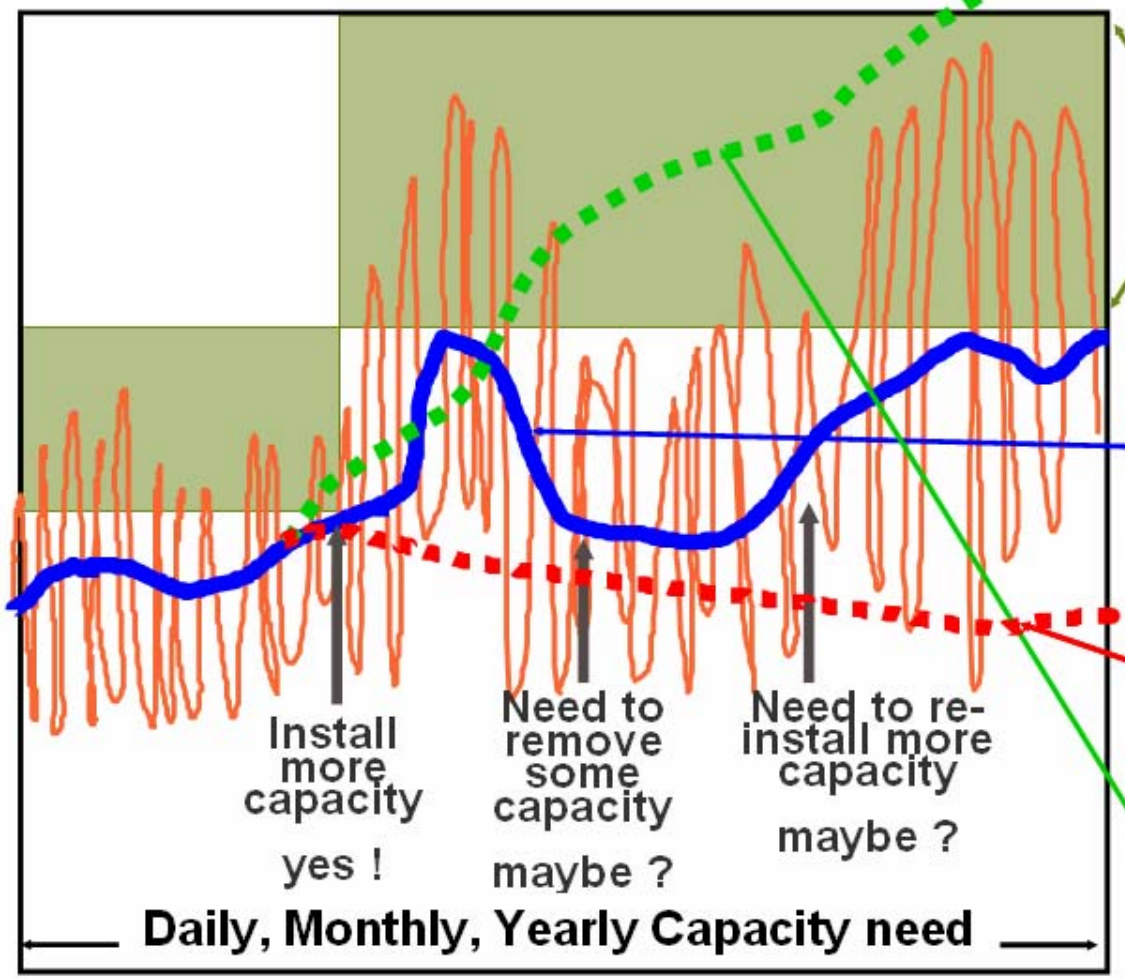

Installed Capacity (interval)
- ➔how much to install ?
- ➔depends of QoS, current and future business needs...

Average Load of Peak hours

Risk
Issue:
Need to align cost with business

Opportunity
Issue:
Need to ensure readiness for business requirements

Source: IBM & Lego Joint SCON Study

Source: Lego Customer Study

Rules for Virtualization

- Hot spot avoidance – Reduce risk of multiple workloads saturating given physical server
  - *If the virtual servers are spread over multiple small servers (e.g., four 2-ways) rather than hosted by a large server (e.g, a single 8-way x460) then surges in workload demand may saturate a given physical server while other physical servers are underutilized. A larger shared server can avoid artificial hot spots which impact workload QoS (throughput and response time).*
- Headroom for hard to predict workloads
  - *It is often the case in planning a new multi-system configuration that the resource needs of some systems/workloads are not known with confidence. In such cases, there is a big advantage in using a scalable server such as an x460 to host multiple virtual servers so that surprises in individual system/workload demand can be handled without acquiring additional server hardware or making physical changes.*
- Peak load Responsiveness
  - *With a large server rather than N smaller servers, there is a greater probability that a processor will be free at any point in time to handle spikes in any workloads demand.*

Source: J. Rymarczyk

IBM Fellow, Chief Virtualization Technologist

Gather The Data

- VMware® Capacity Planner
- IBM CDAT
- Both gather inventory and performance information for servers and clients.
- Data is needed to provide the best sizing for your customers workloads.
- Using Virtualization for Server Consolidation is only as good as the initial sizing information.

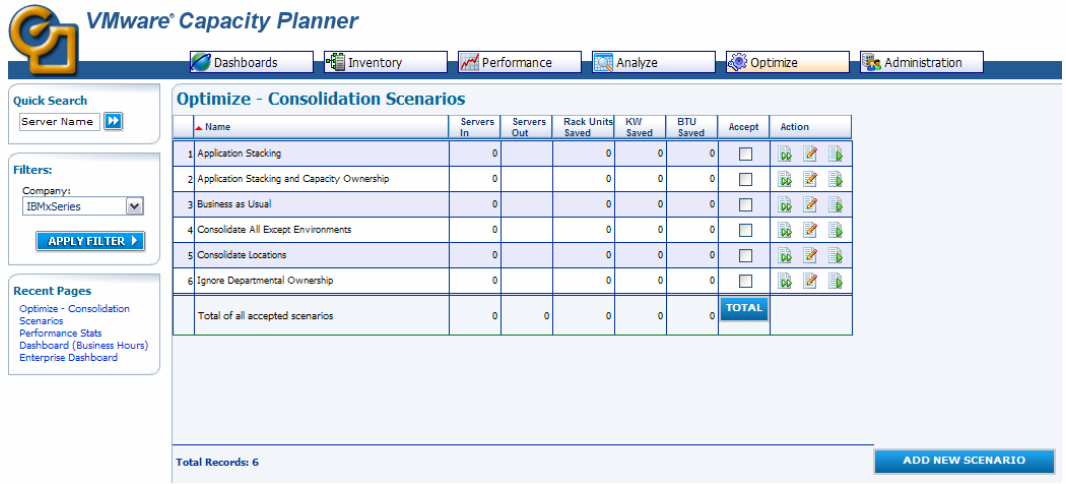

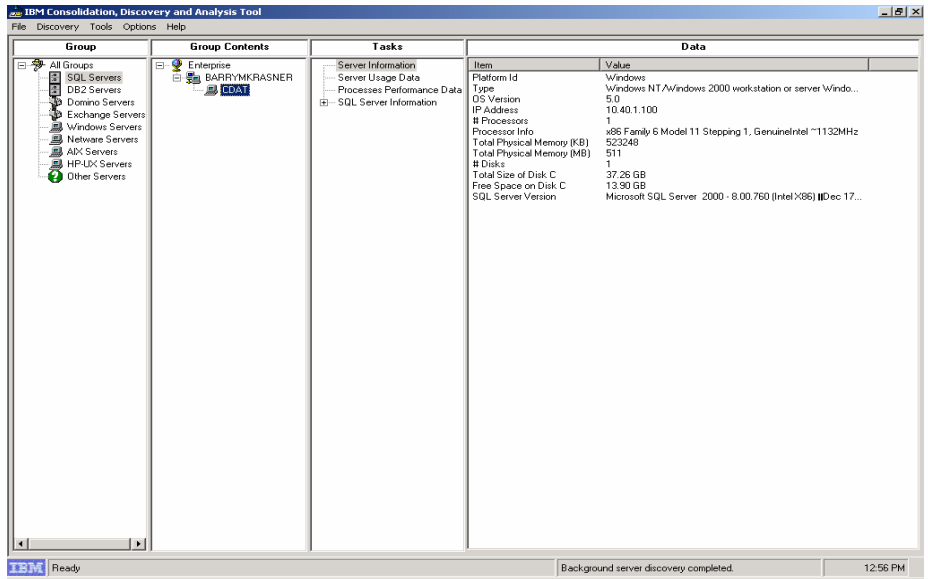

Real-World Workload Environments

Non-Steady State (Unpredictable)

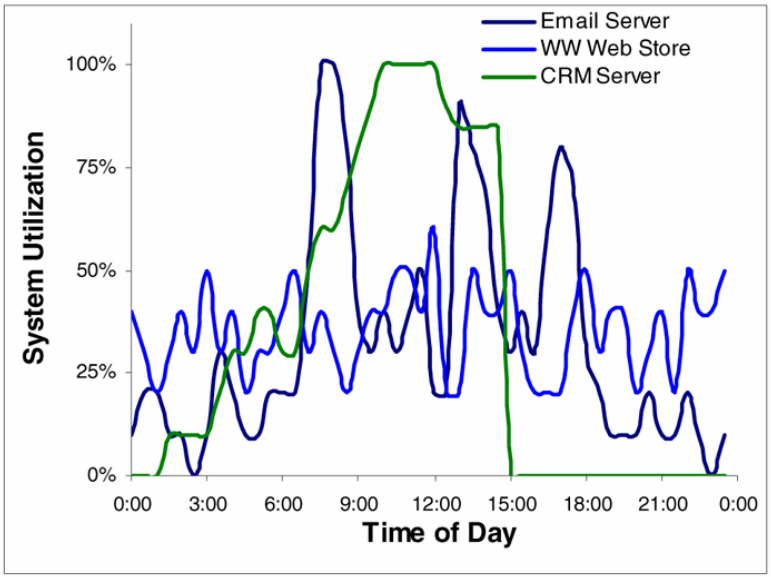

Steady State (Predictable)

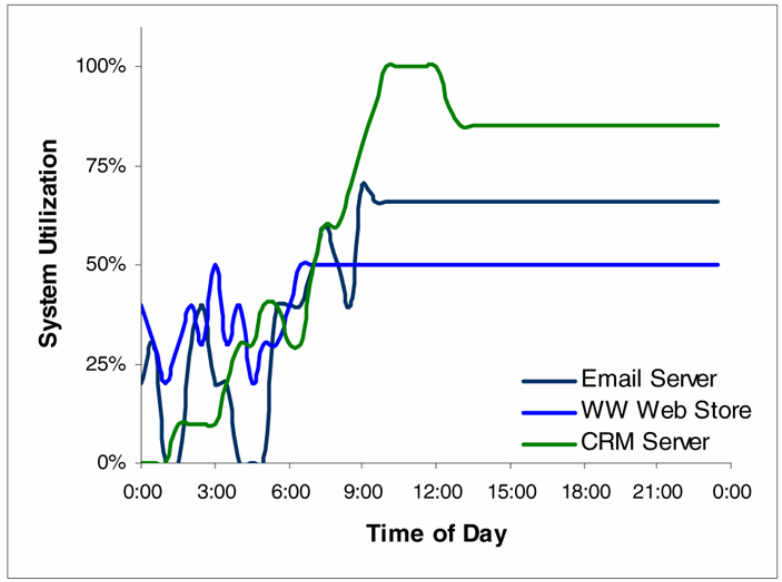

Non-Steady vs Steady State workloads

1.Peak utilization should be examined for all Non-Steady workloads. Email server peak periods may be entirely different from the CRM server, causing peaks and valleys to either negate each other, or doubling the amount of resources needed.

2.Average utilization can be examined for Steady state workloads, once they have achieved steady state. Measurements can be performed prior to Steady State, but this would not reflect the workloads true resource needs.

Source : Intel Technology Journal, Vol 10 Issue 3, published August 10, 2006 - "Redefining Server Performance Characterization for Virtualization Benchmarking"

What is VMware® Capacity Planner?

- Agent-less discovery, inventory, and performance collection
- Data Sources: WMI, Registry & Perfmon API calls on Windows and Remote SSH Sessions using UNIX and Linux utilities
- Hourly performance metrics for days/weeks
- CDAT compatible
- Web-based reporting
- Industry benchmarking / comparison
- Consolidation and virtualization scenario modeling
- Incorporated into BCE and Virtualization Assessment reports

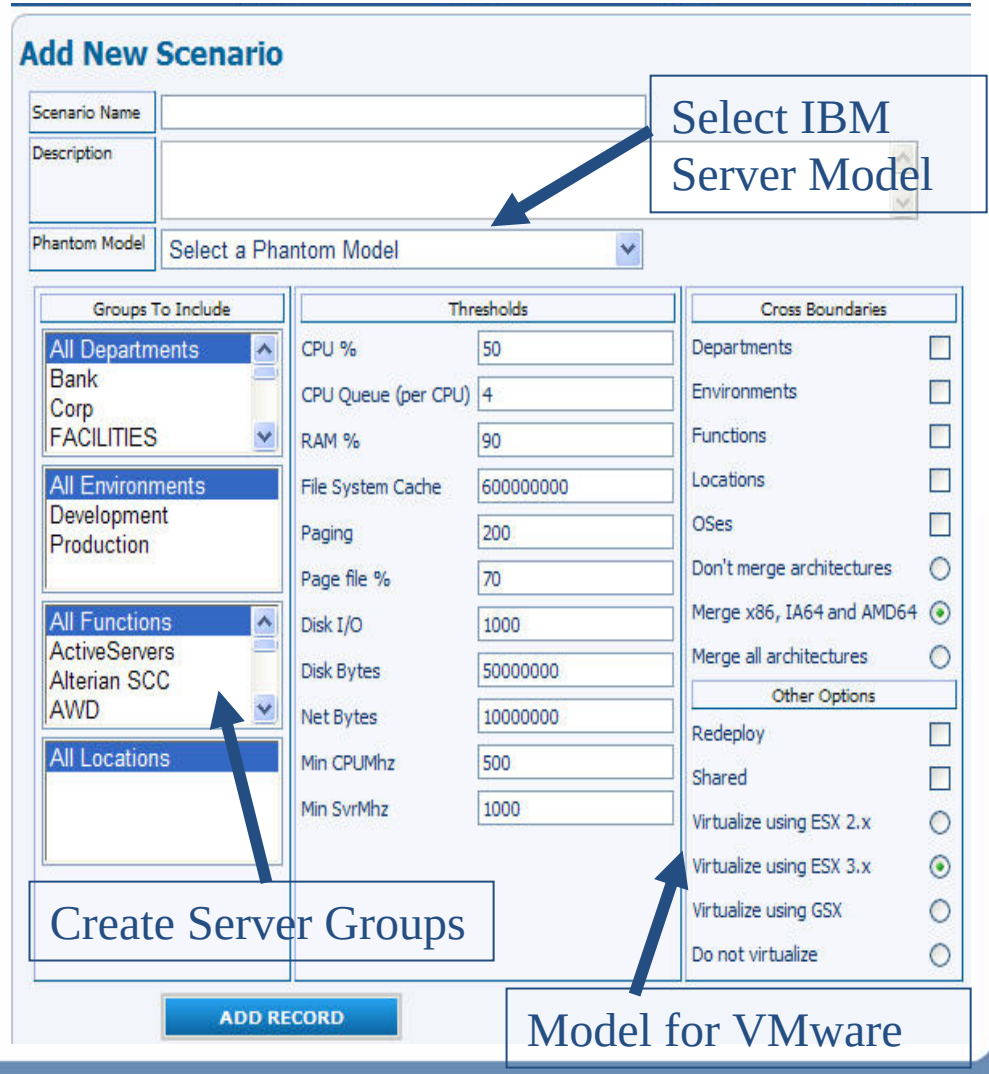

What is CDAT? – Consolidation, Discovery & Analysis

- Discovers servers on the network
- Establishes communication with the servers on the network:
  - SNMP for all servers except xSeries
  - Standard API calls for xSeries
- Takes a snap shot of the data that is collected and stored on the server by the OS and the performance counters
- Exports the data into Excel spread sheet which can be exported to any spread sheet of your choice

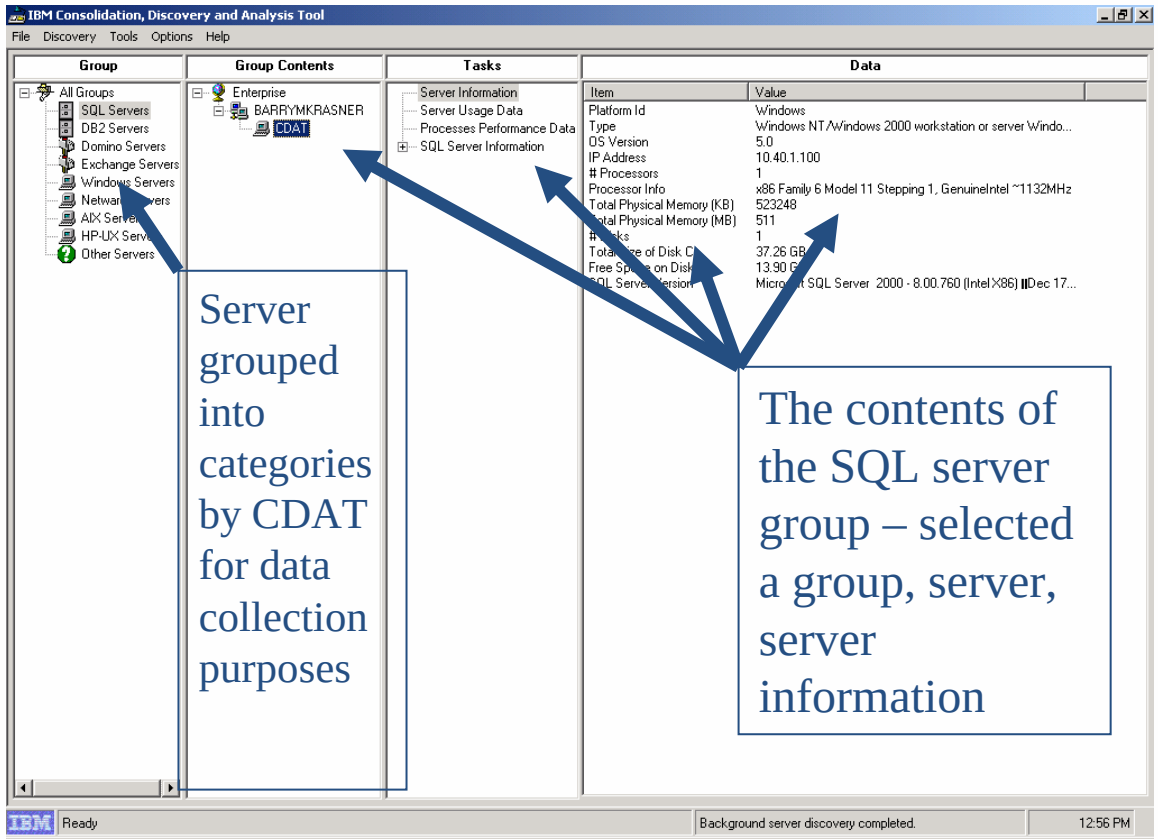

Customer Consolidation Study

- Methodology for the Virtualization Sizing Guide
  - Used customer survey data to define key workload classifications targeted for server consolidation
  - Analyzed recent CDAT(1) studies for 30 customers
    - Organized their legacy servers by application type into the key workload classifications
    - Entered more than 3,200 servers in the customer study database
    - Defined typical average and peak utilization statistics for each key workload classification
  - Input typical workloads from each workload classification into IBM's VISIAN(2) tool
    - VISIAN consolidated the workloads into virtual machines onto a target server according to the Headroom Rules.
    - VISIAN defined the limiting factor for adding additional workloads (CPU, Memory, Virtual CPU)

Notes :

1. IBM's CDAT is a stand-alone discovery tool that automatically gathers a significant amount of server inventory, configuration and performance data, generating an enterprise-wide view of server population and usage. IBM has performed over 2,000 customer server consolidation studies involving over 200,000 servers.
2. IBM's VISIAN is a tool that facilitates the calculation of the optimal number of machines that can be consolidated onto an IBM System x server running VMware ESX Server 2.5. VISIAN is an internal-use-only tool, and an patent is pending.

Virtualization Sizing Insights

- **Software Performance Variation**
  - Customer consolidation studies highlight a greater variation between average CPU utilization and peak CPU utilization on legacy 2-P source workloads than on 4-P and 8-P source workloads.
  - Variations in the CPU utilization will determine the number of virtual machines possible on a server configuration.
- **Server Headroom**
  - Size for excess capacity to ensure that workloads have headroom for peak periods of operation when they require greater server resources than their average or steady state period of operation.
  - To calculate the correct amount of headroom, consider the variation between average and peak server resource requirements for each consolidated workload.
  - The Headroom Rules account for this variation.
- **Peak Utilization versus Average Utilization**
  - To correctly size the number of virtual machines per platform, consider the average and peak utilization of the processor, memory, and disk.
  - The chart on the next page shows the recommended number of VMs based on both the peak and average utilizations. The first lettered indicator aligns with the number of recommended VMs based on the peak utilization values. The second lettered indicator aligns with the number of recommended VMs based on the average utilization values.
  - The letter represents the constraining factor which is the resource on the server that limits the recommended number of VMs for that server configuration.
- **Memory versus CPU constraints**
  - Physical CPU utilization is not always the primary constraint for consolidated workloads.
  - Workloads that are consolidated based on peak utilizations tend to be CPU-constrained and those based on average utilizations tend to be memory-constrained or vCPU-constrained.
- **VMware Scheduler Contention**
  - The amount of work that can be scheduled has a non-linear relationship to the number of CPUs and the number of vCPUS defined for a given virtual machine.
  - VMs with more than 1 vCPU have lower scheduler latency on SMP servers with more CPU cores compared to servers with less CPU cores.

Server Headroom
Server Headroom
Peak and Average Utilization
Scheduler Contention

Server Headroom Calculation

- Data collected from 3,000 servers via CDAT consolidation studies
  - Windows and Linux non-virtualized environments
  - Average seven different application workloads
  - Larger servers with more resources are less volatile
  - 2005 & 2006 studies adds another 12K servers to our data
  - VISIAN consolidated the workloads into virtual machines onto a target server.
  - VISIAN defined the limiting factor for adding additional workloads (CPU, Memory, Virtual CPU)

Virtualized Workloads

Server	Avg CPU Utilization	Peak CPU Utilization	Server Headroom
2-P	44%	90%	37%
4-P	60%	90%	55%
8-P	68%	90%	65%

- Virtualization increases server utilization, but proper configuration must allow for application usage spikes
  - 2-P Add 37% headroom to 7% average to achieve **44%** avg and 90% peak virtual capacity
  - 4-P Add 55% headroom to 5% average to achieve **60%** avg and 90% peak virtual capacity
  - 8-P Add 65% to 3% average to achieve **68%** avg and 90% peak virtual capacity

Rule of thumb for virtualized System x Servers:

2P = 44% 4P = 60% 8P = 68%

NOTE : The information in the following table represents the conclusions of IBM from testing of systems in a controlled environment. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. IBM makes no representation or warranty that an individual user will achieve results equivalent to the levels stated in this document.

Virtualization Sizing Insights

- Server Headroom
- Peak and Average Utilization
- Scheduler Contention

x86 Server Utilization Observations

- The CDAT information was broken into two types of workloads, 2-P and 4-P.
- The sizing was based on both the average and peak for CPU, memory, disk, and network I/O.
  - The disk and network I/O was not included in the print document, it was not a constraint that was seen in the sizing table.
- This information was used to size the target servers.
  - Each cell of the sizing guide is represented by these 6 workloads, and added in a round robin fashion to achieve the needed workloads.

Consolidation Parameters for Source Workloads					
Legacy 2-P Workloads	**Typical Processor**	**Avg CPU Utilization**	**Peak CPU Utilization**	**Avg Memory Used**	**Peak Memory Used**
Infrastructure	Xeon 2.0GHz	8%	48%	568	768
Web	Xeon 1.8GHz	5%	47%	592	768
Application	Xeon 1.8GHz	8%	52%	611	768
Database	Xeon 1.8GHz	9%	60%	1,199	1,536
Terminal Server	PIII 1.4GHz	9%	70%	603	1,024
Email	Xeon 2.0GHz	6%	50%	994	1,280

Legacy 4-P Workloads	**Typical Processor**	**Avg CPU Utilization**	**Peak CPU Utilization**	**Avg Memory Used**	**Peak Memory Used**
Infrastructure	Xeon MP 2.5GHz	6%	35%	841	1,024
Web	Xeon MP 2.5GHz	4%	24%	737	1,024
Application	Xeon MP 2.7GHz	4%	34%	935	1,280
Database	Xeon MP 2.5GHz	5%	37%	1,553	1,792
Terminal Server	Xeon MP 2.7GHz	6%	45%	882	1,536
Email	Xeon MP 2.8GHz	4%	34%	1,295	1,536

Table C - Consolidation Parameters for Source Workloads

NOTE : The information in the following table represents the conclusions of IBM from testing of systems in a controlled environment. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. IBM makes no representation or warranty that an individual user will achieve results equivalent to the levels stated in this document.

Virtualization Sizing Guidance

Legacy 2-P Workloads virtualized to a VM defined as a Single processor (1vCPU)

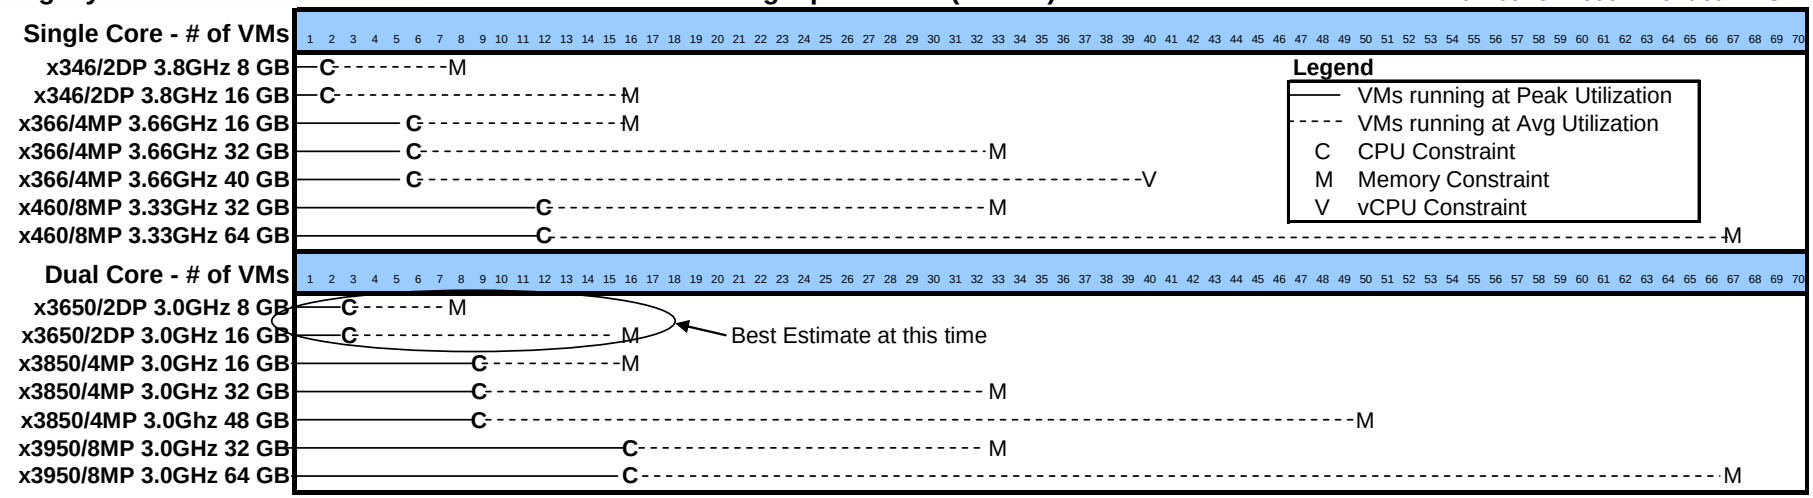

Legacy 4-P Workloads virtualized to a VM defined as a Dual processor (2vCPU)

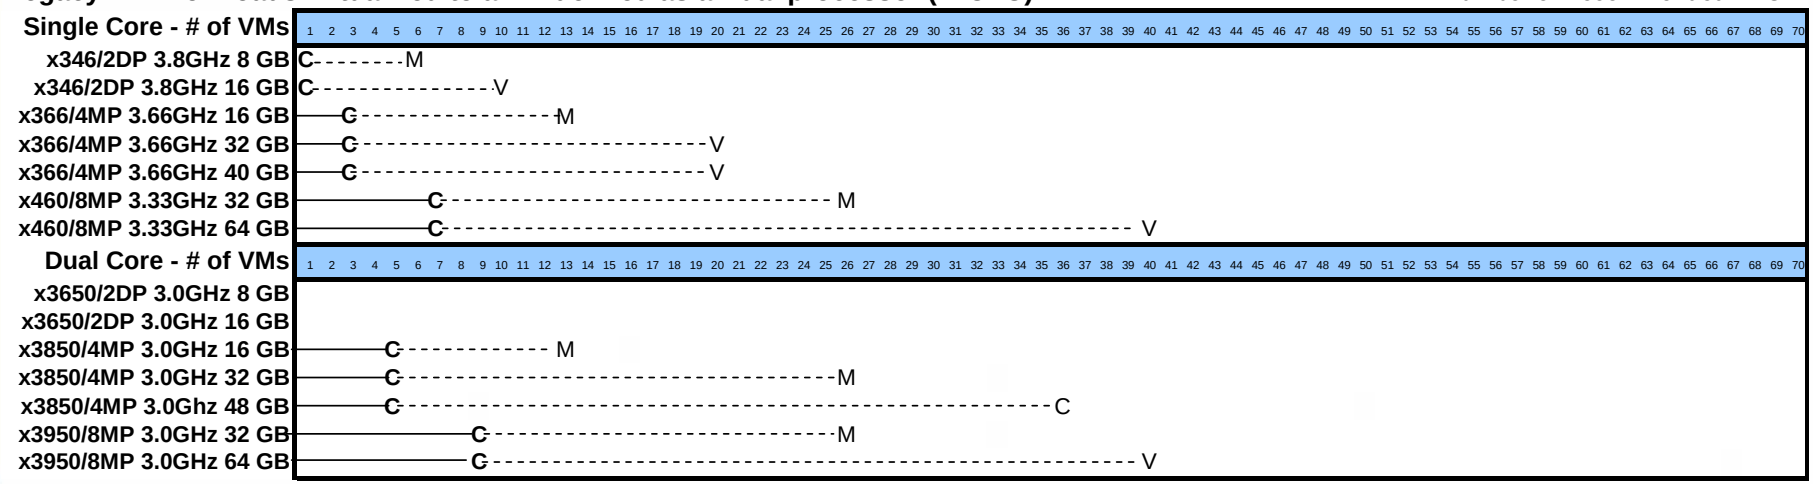

NOTE : The information in the following table represents the conclusions of IBM from testing of systems in a controlled environment. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. IBM makes no representation or warranty that an individual user will achieve results equivalent to the levels stated in this document.

Selection Criteria

Legacy 2-P Workloads virtualized to a VM defined as a Single processor (1vCPU)

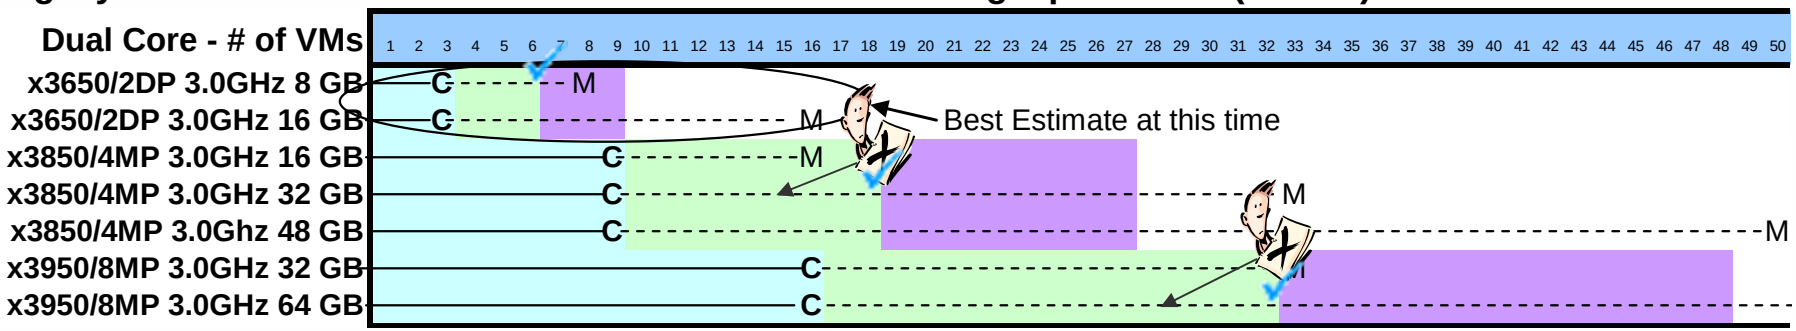

	2-P	4-P	8-P
Zone 1 (VMs/Memory)	**3 / 8GB**	**9 / 16GB**	16 / 24GB
Zone 2	6 / 8GB	18 / 24GB	32 / 40GB
Zone 3	**9 / 16GB**	**27 / 32GB**	48 / 52GB

- Zone 1 *(Conservative)*
  - Aggressive application (ie. Database)
  - Unpredictable workload
  - No tolerance for performance degradation
- Zone 2 *(Moderate)*
  - Moderately aggressive application (ie. Collaboration)
  - Predictable workload
  - Some tolerance for performance degradation
- Zone 3 *(Aggressive)*
  - Lightly aggressive application (ie. Web)
  - Predictable Workload
  - High tolerance for performance degradation

NOTE : The information in the following table represents the conclusions of IBM from testing of systems in a controlled environment. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. IBM makes no representation or warranty that an individual user will achieve results equivalent to the levels stated in this document.

Virtualization Sizing Insights

- Server Headroom
- Peak and Average Utilization
- Scheduler Contention

Guidance from VMware®

The Benefits of CPU Dense ESX Server Hosts:

"The chance that the scheduler can find room for a particular workload without much reshuffling of virtual machines will always be better when the scheduler has more CPUs across which it can search for idle time. For this reason, it will generally be better to purchase two four-way ESX Server licenses than to purchase four two-way machines. Similarly, two eight-way servers will provide more scheduling flexibility than four four-way servers."

Tips and Techniques for Implementing Infrastructure Services on ESX Server

VMware® Scheduler

Single Core 2-P

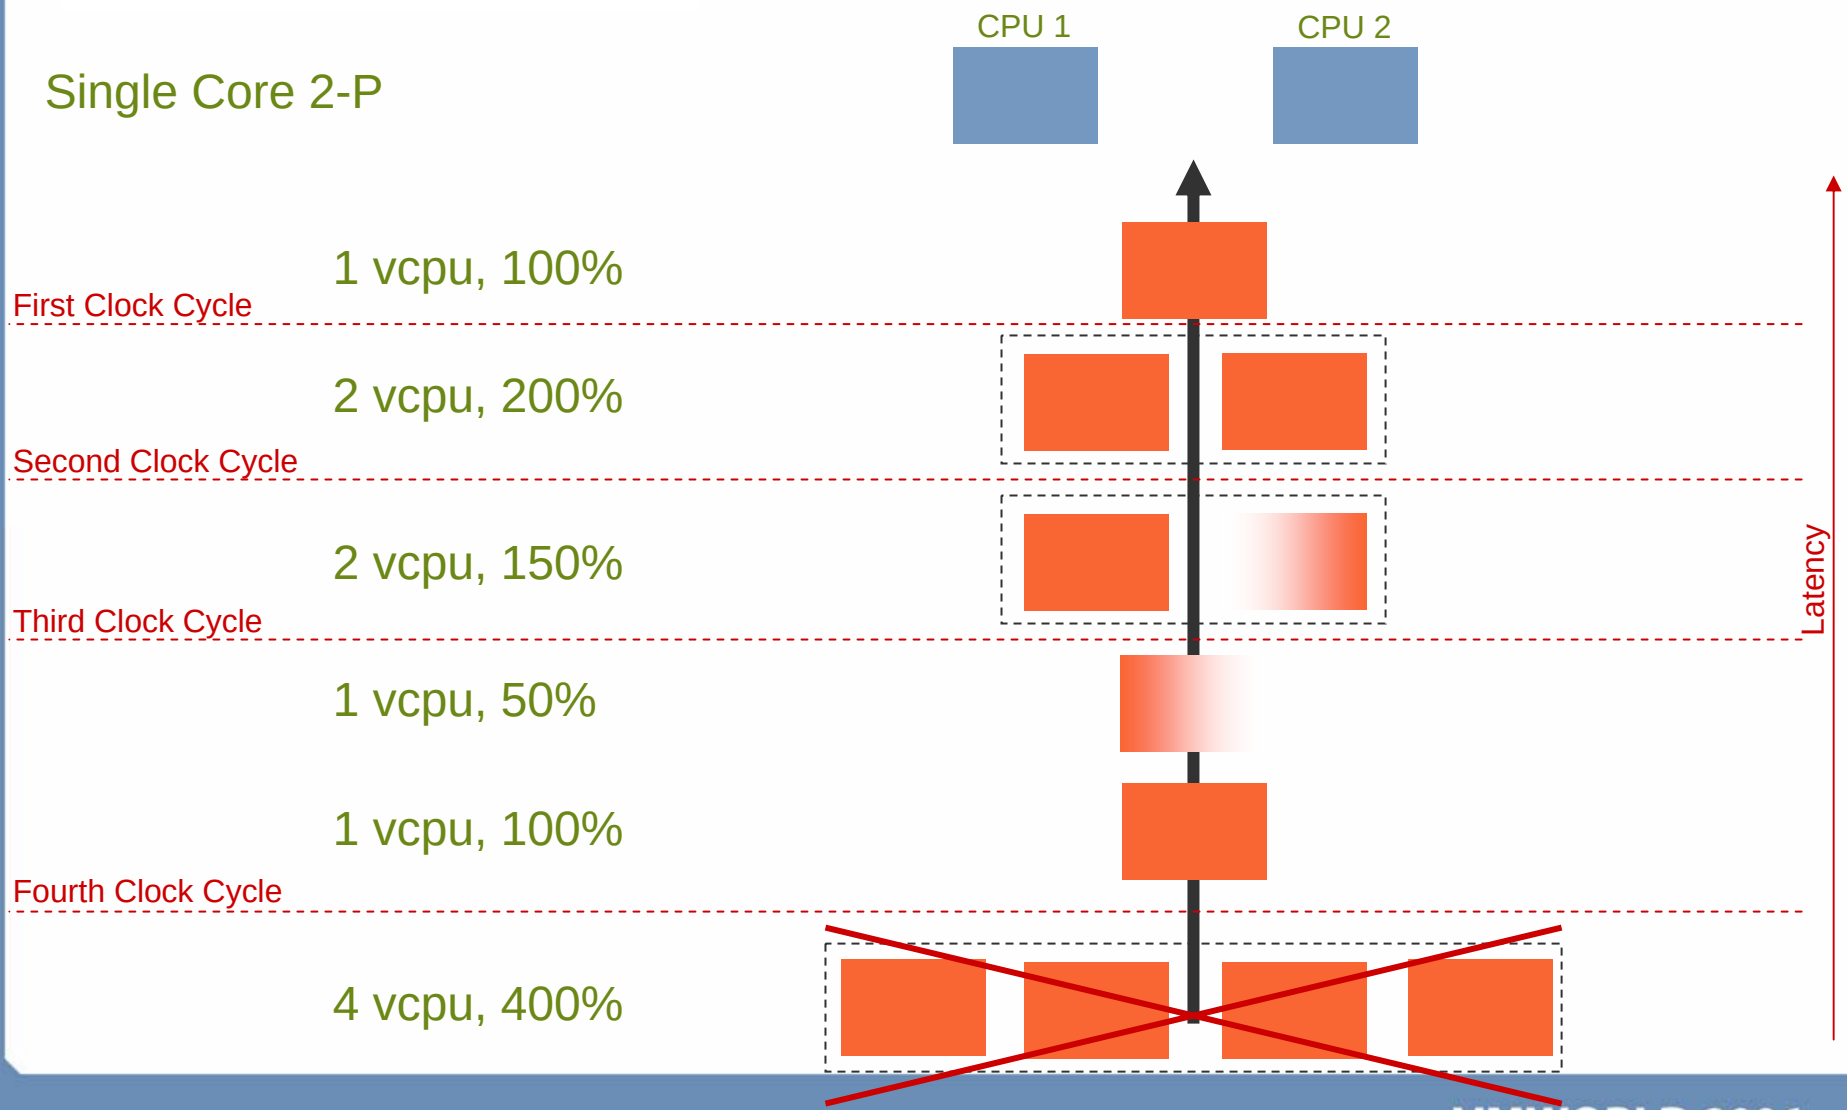

VMware® Scheduler

Dual Core 2-P
&
Single Core 4-P

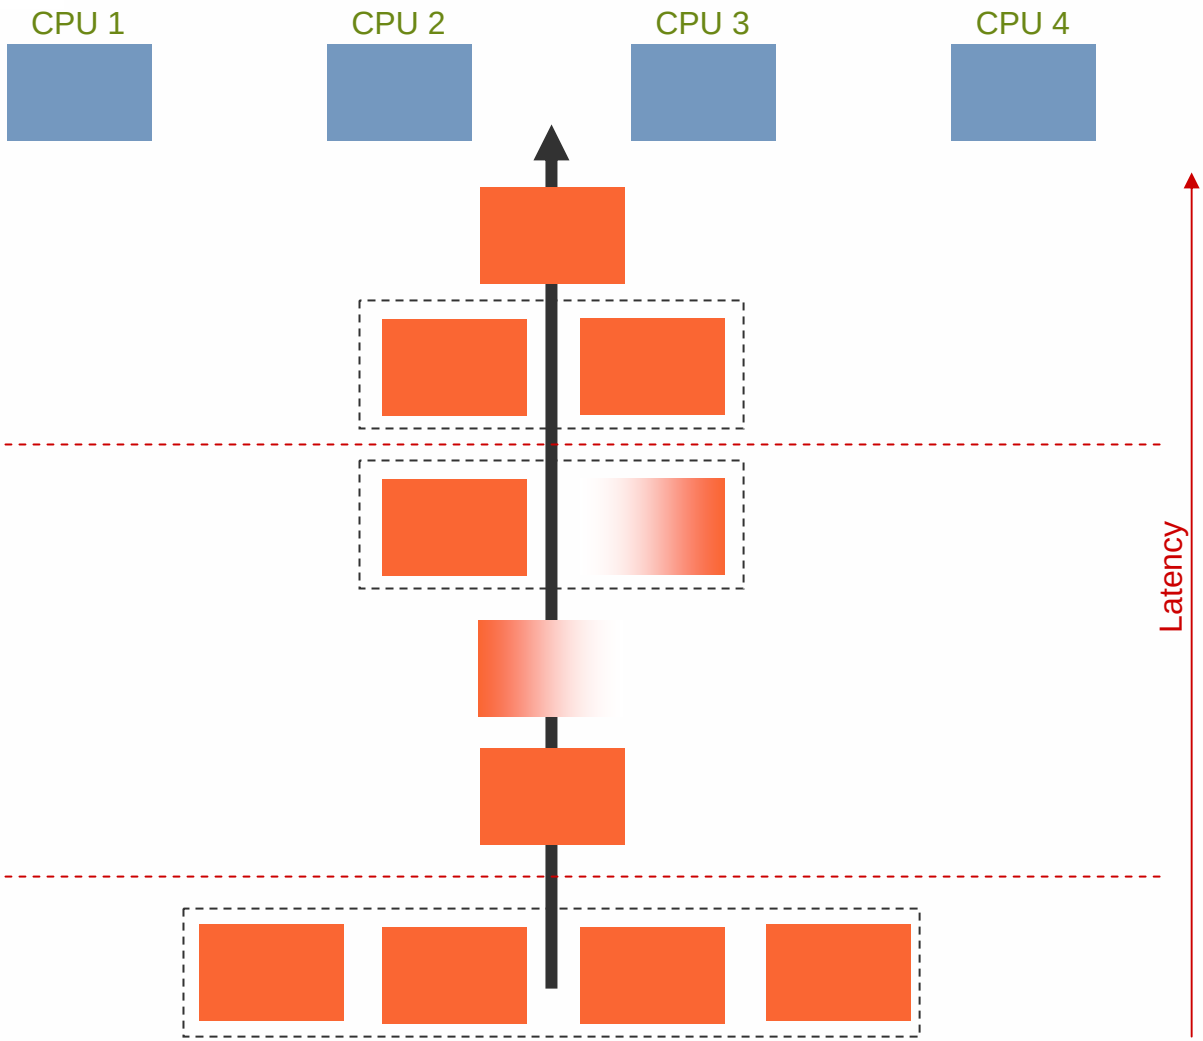

VMware® Scheduler

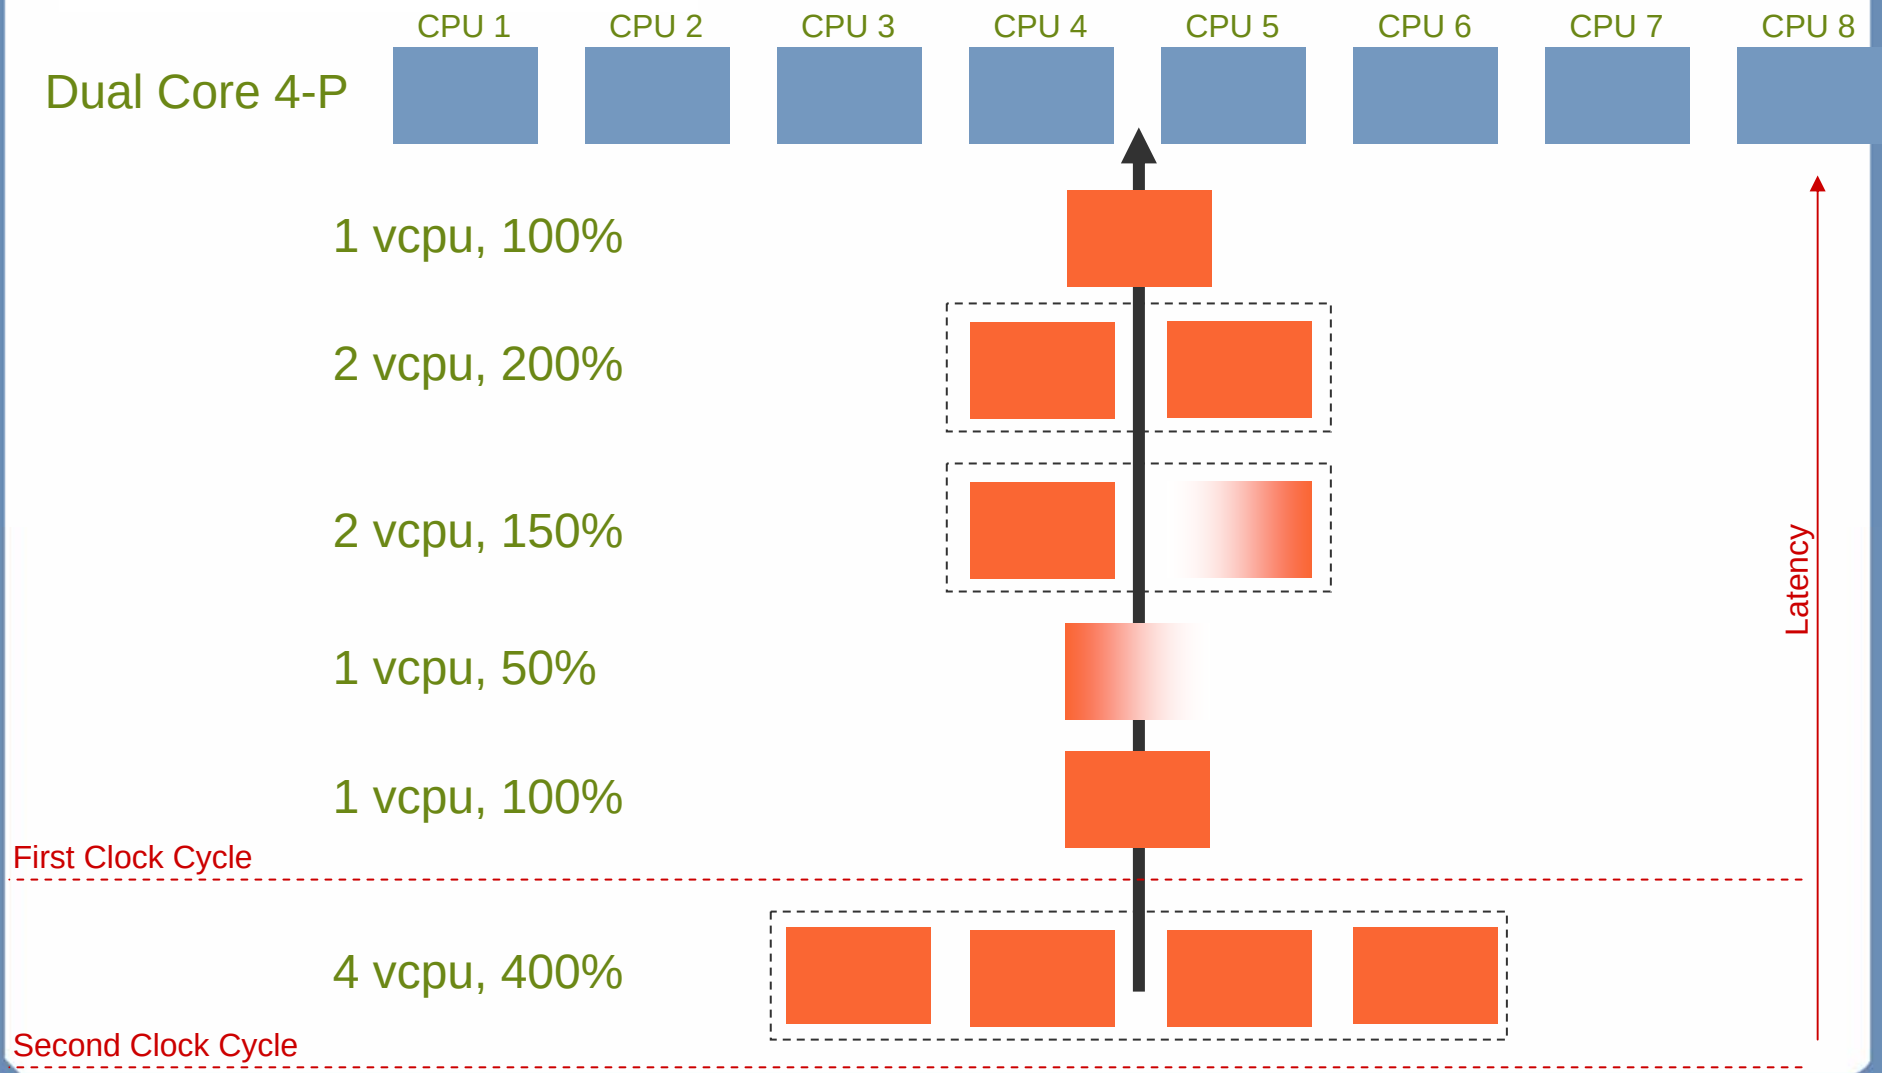

Agenda

- Sizing
  - The elements of sizing
  - IBM's insights into sizing
- IBM Portfolio
  - Portfolio Positioning
  - Competitive Advantage
- Systems Management – IBM Director
  - What is it
  - How it helps virtualization
- Summary

x86 Servers Made Better With Virtualization

	Infrastructure simplification	Application serving	Server consolidation
Platform	**IBM BladeCenter®**	**IBM System x™ Two-processor rack & tower**	**IBM System x™ Four-processor rack**
Description	**Integrating server, networking and storage resources**	**Combining a few applications on a single server for greater utilization**	**Consolidating large numbers of underutilized servers for greatest TCO**
Key attributes	■ **Hardware usability** ■ **Packaging density** ■ **Unified management** ■ **Power/cooling savings** ■ Server mobility ■ High availability/disaster recovery capability	■ **Industry-standard design** ■ **Price/performance** ■ **Compatibility** ■ Multiple apps per server ■ High utilization	■ **Performance** ■ **Scalability** ■ **Strong reliability features** ■ Application isolation ■ Reduction of physical systems ■ Rapid system deployment

IBM differentiates itself from the competition

- Compatibility across chassis
- Comprehensive ecosystem
- Power management
- Two redundant high-speed fabrics

- Innovative design with standard parts
- More memory DIMMS per processor than competitors
- More I/O slots per processor than competitors

- IBM X3 chipset
- Up to 32-socket scalability
- Mainframe-inspired reliability
- PFA on more components than competitors

Common across the portfolio

- IBM CoolBlue™ innovations

- Largest vmware® OEM vendor(1)

- IBM Director
- IBM Virtualization Manager

Note (1) : Based on bundled license revenue

IBM and VMware®: A Relationship Of Firsts!

- First VMware® system vendor
- First joint development partner
- First to leverage VMware® SDK
- First to offer comprehensive support
  - Microsoft Windows, Linux and IBM software in a VMware virtual machine
- First blade offering to include VMware®
- First to integrate VMware® into a Virtual Client Solution
- First system vendor to announce VMware® Infrastructure 3

BladeCenter ® Competitive Advantage

- Up to 30% better power efficiency
- Fully redundant configurations for higher availability
- Largest ecosystem provides more flexibility and choice
- Innovative ‘snap-in’ scalable blade server that scales from 2P to 4P in seconds
- Open architecture allows 3rd party hw and sw vendors to provide more BladeCenter® solutions
- Blade compatibility across chassis provides investment protection

Common across the portfolio

- IBM CoolBlue™ innovations
- Largest vmware® OEM vendor(1)

- IBM Director

- IBM Virtualization Manager

Note (1) : Based on bundled license revenue

When do I deploy IBM BladeCenter® Systems?

When you want to:

- Consolidate infrastructure and centralize management
- Reduce the complexity of ‘scale out’
- Reduce power consumption
- Make more efficient use of datacenter space
  - Reduce floor space consumed
  - Reduce weight
- Achieve maximum processor density

It’s that Simple ….. With IBM

I/O Virtualization Solution

Building Blocks

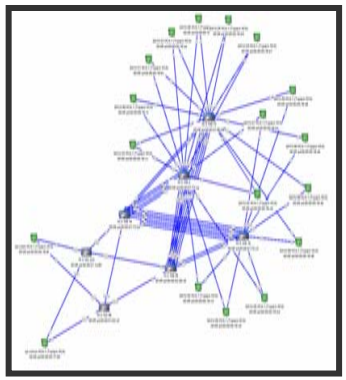

Network-embedded InfiniBand Subnet Manager

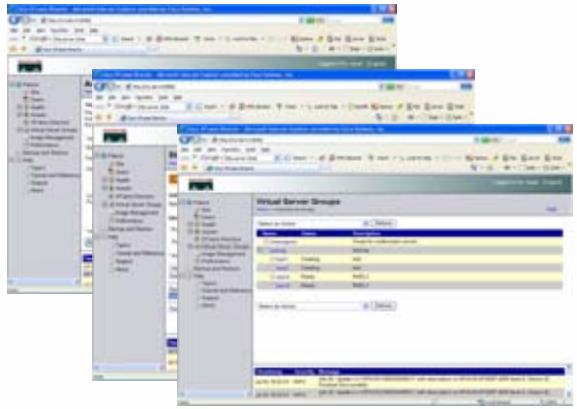

Cisco® VFrame Server Virtualization Software

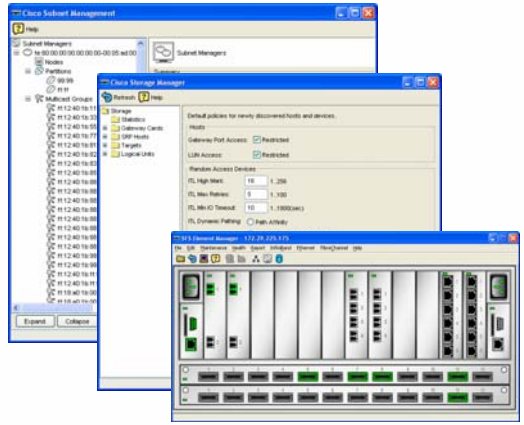

Embedded system and fabric management

Host channel adapter (HCA) With upper layer protocols

- MPI
- IPoIB
- SDP
- SRP

Linux, Windows, and Solaris driver support

Cisco InfiniBand Switch Module for IBM BladeCenter

Cisco SFS 3012

- InfiniBand Switching
- InfiniBand to Ethernet
- InfiniBand to Fibre Channel

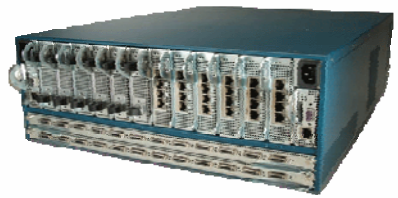

Infrastructure Simplification Benefit

Before Virtualization ….

14 x 2P Blades

- Bladecenter® chassis = $21.9K
- HS20 Blade 3.8GHz w/4GB
- Blade cost = $5.6K x 14 = $78.4K
- Windows SE cost = $800 x 14 = $11.2K
- Average CPU Utilization 8%
- Space Consumption = 7U
- Power Consumption = 2000W x 2 = 4000W
- Total Cost ~ $111.5K
- Cost/Application ~ $7,964

Virtualization could provide up to a ~ 2.25x cost improvement.

After Virtualization ….

14 x 2P Blade
56 Applications

- Bladecenter® chassis = $21.9K
- HS20 Blade 3.8Ghz w/4GB
- Blade cost = $5.6K x 14 = $78.4K
- Windows SE cost = $800 x 56 = $44.8K
- VMware® cost = $3,750 x 14 = $52.5K
- Supports 4 avg and 2 peak utilization VMs(1)
- Recommend 4 VMs/system (Zone 2)
- 84 Applications
- Space Consumption = 7U
- Power Consumption = 2000W x 2 = 4000W
- Total Cost ~ $197.6K
- Cost/Application ~ $3,528

Note (1) : Configuration performed after publication of Sizing Guide, 4 avg vs 2 peak for 4GB configuration

Web list prices found on www.ibm.com as of 8/1/06

BladeCenter® Competitive Advantage

	HP BL465c	LS21	LS41 scalable blade
Dual core AMD Opteron Processors	2	2	4
DIMMs/Max RAM	8 DIMMs 16GB	8 DIMMs 32GB	16 DIMMs 64MB
HDDs	2	1(4)*	2(5)*
I/O expansion slots	2	2	3
High speed (10Gb) support	In future	now	now

Note : Specifications for the server hardware can be found on the following web sites : www.ibm.com, www.dell.com, and www.hp.com

* With optional SIO expansion blade

IBM differentiates itself from the competition

- Innovative design with standard parts
- More memory DIMMS per processor than competitors
- More I/O slots per processor than competitors

Common across the portfolio

- IBM CoolBlue™ innovations

- Largest vmware® OEM vendor(1)

- IBM Director
- IBM Virtualization Manager

Note (1) : Based on bundled license revenue

When do I deploy IBM Two Socket Systems?

When you want to:

- Optimize highly demanding business applications
- Run applications that require high-speed memory subsystems
- Need broad and low-latency I/O choices for Ethernet and fibre channel
- Actively manage energy consumption
- Consolidate legacy 2 Socket servers using virtualization
  - When consolidating several traditional 2 way workloads
  - When consolidating less than 18 traditional 2 way workloads

It's that Simple ….. With IBM

Application Serving Benefit

Before Virtualization ….

6 x 2P Server

- x3650 3.0GHz DC w/4GB
- Server cost = $4.9K x 6 = $29.4K
- Windows SE cost = $800 x 6 = $4.8K
- Average CPU Utilization 8%
- Space Consumption = 12U
- Power Consumption = 242Wx6 = 1452W
- Total Cost ~ $34.2K
- Cost/Application ~ $5,700

After Virtualization ….

2P Server

- x3650 3.0GHz DC w/8GB
- Server cost = $9.5K
- Windows SE cost = $800 x 6 = $4.8K
- VMware® cost = $3,750 x 2 = $7.5K
- Supports 8 avg and 3 peak utilization VMs
- Recommend 6 VMs/system (Zone 2)
- Space Consumption = 2U
- Power Consumption = 242W
- Total Cost ~ $21.8K
- Cost/Application ~ $3,633

Virtualization could provide up to a ~ 1.5x cost improvement.

Web list prices found on www.ibm.com as of 8/1/06

How about 12 legacy 2 way workloads

Before Virtualization ….

After Virtualization ….

12 Servers

Serving from….

2P Servers

4P Server

- x3650 3.0GHz DC w/8GB
- Server cost = $9.5K x 2 = $19K
- Windows DC cost = $2K x 2 x 2 = $8K*
- VMware® cost = $3,750 x 2 x 2 = $15K
- Supports 8 avg and 3 peak utilization VMs
- Recommend 6 VMs/System (Zone 2)
- Space Consumption = 18U
- Power Consumption = 246Wx2 = 492W
- Total Cost ~ $42K
- Cost/VM ~ $3,500

- x3850 3.0GHz DC w/24GB
- Server cost = $24.4K = $24.4k
- Windows DC cost = $2K x 4 = $8K*
- VMware® cost = $3,750 x 4 = $15K
- Supports 24 avg and 9 peak utilization VMs
- Recommend 18 VMs/system (Zone 2)
- Space Consumption = 3U
- Power Consumption = 799W
- Total Cost ~ $47.4K
- Cost/VM ~ $3,950

Cost per VM could be up to 22% less on the 2P Platform

Note : MS Windows Datacenter licensing – Estimated $2,000 processor, unlimited virtualization client. GA 9/28/06.

Web list prices found on www.ibm.com as of 8/1/06

IBM 2 Socket Competitive Advantage

	Dell	HP DL385 G2	IBM x3655
Dual core AMD Opteron Processors	?	2	2
DIMMs/Max RAM	?	8 DIMMs 32GB	**16** DIMMs 64MB
HDDs	?	2(8)*	2(8)*
I/O expansion slots	?	4	4
High speed (10Gb) support	?	now	now

Dell has NO mainstream 2U two-socket AMD offering!!

Note : Specifications for the server hardware can be found on the following web sites : www.ibm.com, www.dell.com, and www.hp.com

IBM differentiates itself from the competition

- IBM X3 chipset
- Up to 32-socket scalability
- Mainframe-inspired reliability
- PFA on more components than competitors

Common across the portfolio

- IBM CoolBlue™ innovations
- Largest vmware® OEM vendor(1)

- IBM Director

- IBM Virtualization Manager

Note (1) : Based on bundled license revenue

When do I deploy IBM Scaleable Systems?

When you want to:

- Run Applications that Multi-Thread
  - SAP, SAS, DB2, Oracle, SQL, Cognos
- Run Applications that have large local storage requirements
  - Database, Large Collaboration and or Messaging
- Consolidate legacy servers using virtualization
  - When consolidating more than 18 traditional 2 way workloads
  - When consolidating more than 10 traditional 4 way workloads
  - When partitions needs 2 or more virtual CPU's

It's that Simple ….. With IBM

Datacenter Consolidation Benefit

Before Virtualization ….

18 Servers

- x3650 3.0GHz DC w/4GB
- Server cost = $4.9K x 18 = $88.2K
- Windows SE cost = $800 x 18 = $14.4K
- Average Utilization 5%
- Peak Utilization 47%
- Space Consumption = 36U = 1 Rack
- Power Consumption = 242W x 18 = 4356W
- Total Cost ~ $102.6K
- Cost/Application ~ $5,700

After Virtualization ….

4P Server

- x3850 3.0GHz DC w/24GB
- Server cost = $24.4K
- Windows DC cost = $2K x 4 = $8K*
- VMware® cost = $3,750 x 4 = $15K
- Supports 24 avg and 9 peak utilization VMs
- Recommend 18 VMs/system (Zone 2)
- Space Consumption = 3U
- Power Consumption = 835W
- Total Cost ~ $47.4K
- Cost/Application ~ $2,633

Virtualization could provide up to a 2.2x cost improvement

Note : MS Windows Datacenter licensing – Estimated $2,000 processor, unlimited virtualization client. GA 9/28/06.

Web list prices found on www.ibm.com as of 8/1/06

How about 36 legacy 2 way workloads

Before Virtualization …. **36 Servers** *After Virtualization ….*

Consolidated to ….

2P Servers

- x3650 3.0GHz DC w/8GB
- Server cost = $9.5K x 6 = $57K
- Windows DC cost = $2K x 2 x 6 = $24K*
- VMware® cost = $3,750 x 2 x 6 = $45K
- Supports 8 avg and 3 peak utilization VMs
- Recommend 6 VMs/System (Zone 2)
- Space Consumption = 18U
- Power Consumption = 246Wx6 = 1476W
- Total Cost ~ $126K
- Cost/VM ~ $3,500

4P Server

- x3850 3.0GHz DC w/24GB
- Server cost = $24.4K x 2 = $48.8k
- Windows DC cost = $2K x 4 x 2 = $16K*
- VMware® cost = $3,750 x 4 x 2 = $30K
- Supports 24 avg and 9 peak utilization VMs
- Recommend 18 VMs/system (Zone 2)
- Space Consumption = 6U
- Power Consumption = 799Wx2 = 1598W
- Total Cost ~ $94.8K
- Cost/VM ~ $2,633

Cost per VM could be up to 25% less on the 4P Platform

Plus Lower Software & Management Costs

Note : MS Windows Datacenter licensing – Estimated $2,000 processor, unlimited virtualization client. GA 9/28/06.

Web list prices found on www.ibm.com as of 8/1/06

X Architecture Competitive Advantage

	Dell 6950	HP DL585	IBM x3755
Dual core AMD Opteron Processors	4	4	4
DIMMs/Max RAM	16 DIMMs 64GB	32 DIMMs 64GB	32 DIMMs 128MB
HDDs	4 X 3.5” SAS	8 x 2.5” SAS	5 x 3.5” SAS
I/O expansion slots	PCI-X 0 PC-E 2(x8), 5(x4)	PCI-X 2(100mhz) PCI-E 3(x8), 4(x4)	PCI-X 2(66mhz) PC-E 1(x16), 3(x8)
HTX Support	No	No	Yes – 1 slot

Note : Specifications for the server hardware can be found on the following web sites : www.ibm.com, www.dell.com, and www.hp.com

Application Serving versus Server Consolidation

2P Server

- Workgroup Consolidation
- Drive server utilization
- Multiple applications per server

- Power Reduction
- Management Features

4P Server

- Datacenter Consolidation
- Smallest number of servers
- Application Isolation

- Power Reduction
- Management features

Platform selection driven by size, and number of virtual machines

Agenda

- Sizing
  - The elements of sizing
  - IBM's insights into sizing
- IBM Portfolio
  - Portfolio Positioning
  - Competitive Advantage
- Systems Management
  - IBM Virtualization Manager
  - IBM Director
- Summary

IBM Director

From Deployment Through Maintenance …

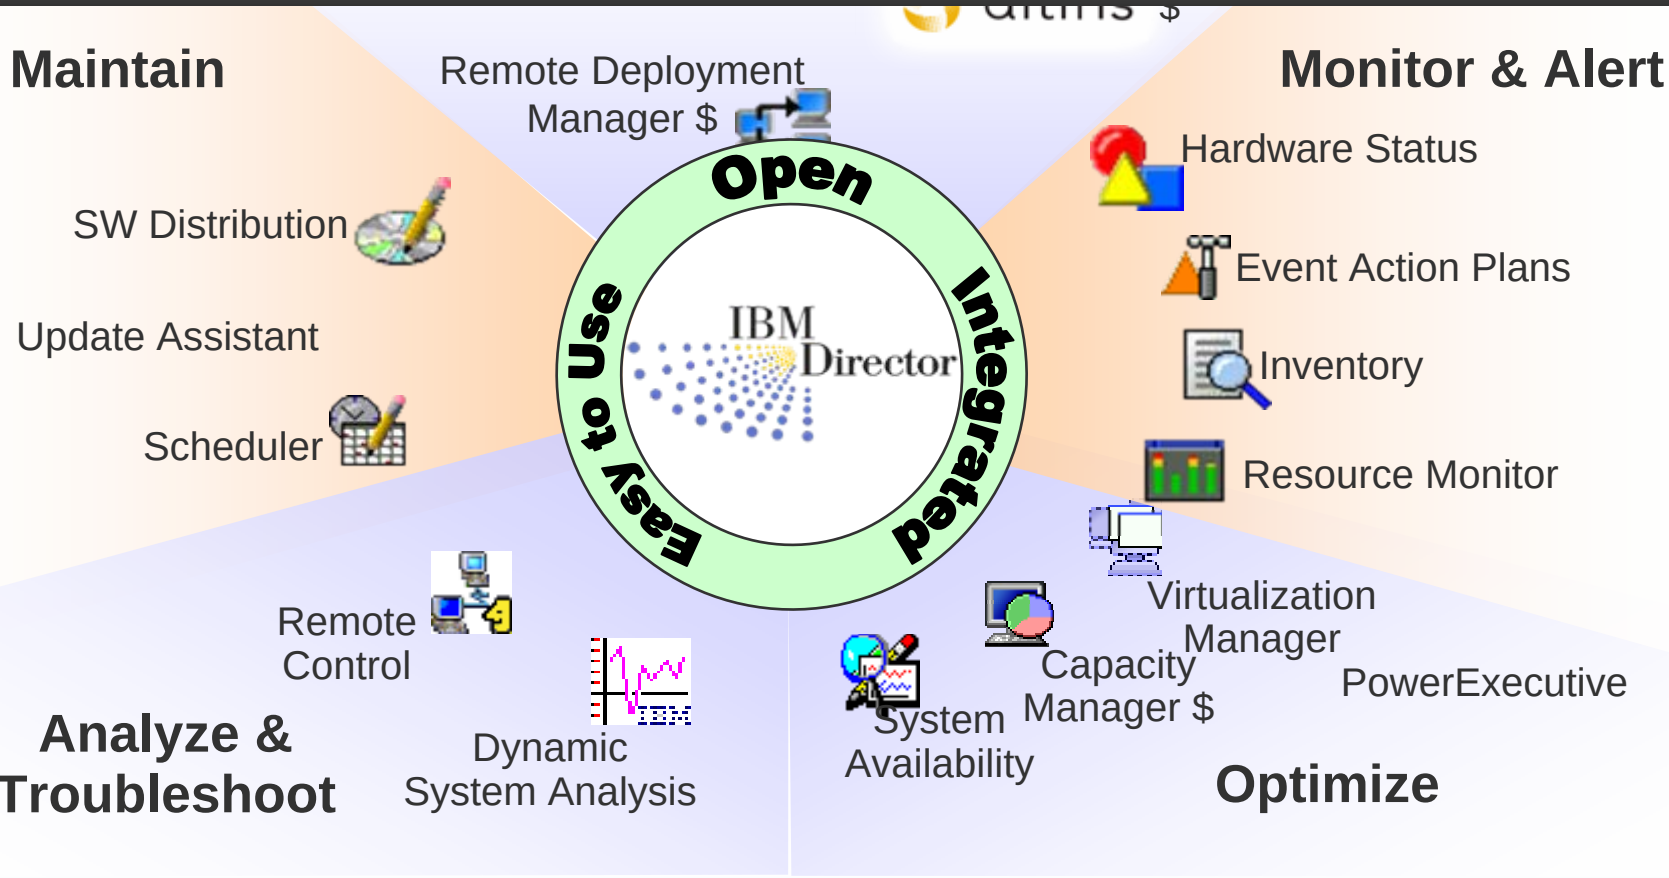

IBM Virtualization Manager

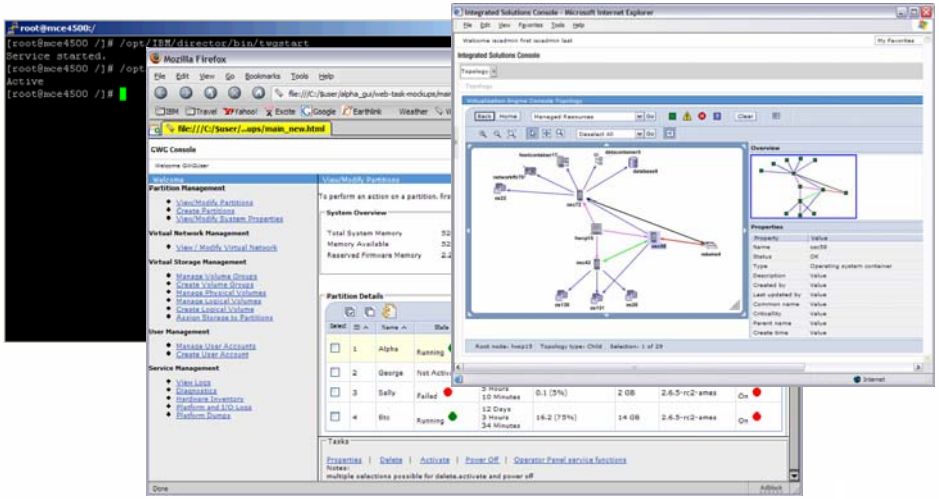

Simplified management of virtualized systems via standard interfaces regardless of platform or virtualization technology

- ✓ Remove barriers to virtualization
- ✓ Remove pain points of virtualization

Industry standards-based virtualization mgmt interfaces (DMTF)

CIM in DMTF

SuSE redhat Windows

Xen vmware

intel inside XEON AMD

CIM in DMTF

AIX AIX SuSE redhat i5/OS SuSE i5/OS

PHyP MEIOSYS

Power changes everything

CIM in DMTF

z/OS SuSE SuSE redhat redhat

VM/ESA z/VM

zSeries makes it possible

IBM Virtualization Manager

Key Design Points

- Discover and visualize resources and relationships
  - Out-of-the-box discovery, easily find resources and relationships, analyze environment, etc.
- Show health and tasks from all resources and relationships
  - Define and monitor health, drill down on problems quickly to find root cause
  - Provide common tasks that work across all resources
  - Provide detailed tasks to perform platform-specific tasks in common ways
- Grow existing workload by expanding or migrating
  - Increase virtual server's memory, storage, processing, networking capabilities
  - Growing virtual server workload can be moved to more powerful physical server
- Use virtualization to help with repair and upgrade of servers
- Add new resources and work into virtual environment
  - Easily allocate, configure, and manage virtual servers
  - Make changes transparently without "change windows"

IBM Virtualization Manager and VMware VirtualCenter

- VMware® VirtualCenter provides management of homogeneous VMware environments
  - Dynamic (live) migration via VMotion™
  - Supports cloning and P2V
  - Single point of management for multiple virtual hosts
  - Focuses on virtualization management, not physical systems
- IBM Virtualization Manager complements VMware® VirtualCenter, providing one management interface for both the virtual and physical resources
- IBM Virtualization Manager provides common interface for hardware and virtual machine management
  - VMware®, with or without VirtualCenter
  - Microsoft Virtual Server
  - Xen hypervisor support in SLES 10
  - pHype endpoints managed by HMC (System p)
- Availability and failover with Event Action Plans
  - Actual or predicted system failure
  - Static migration of all supported hypervisor products
  - Dynamic migration via VMotion™ for VMware® environments

Resource Navigator

- The Resource Navigator page is used to determine the status of virtualization components, as well as their relationships to each other. In addition to the list view shown here, other views are available.

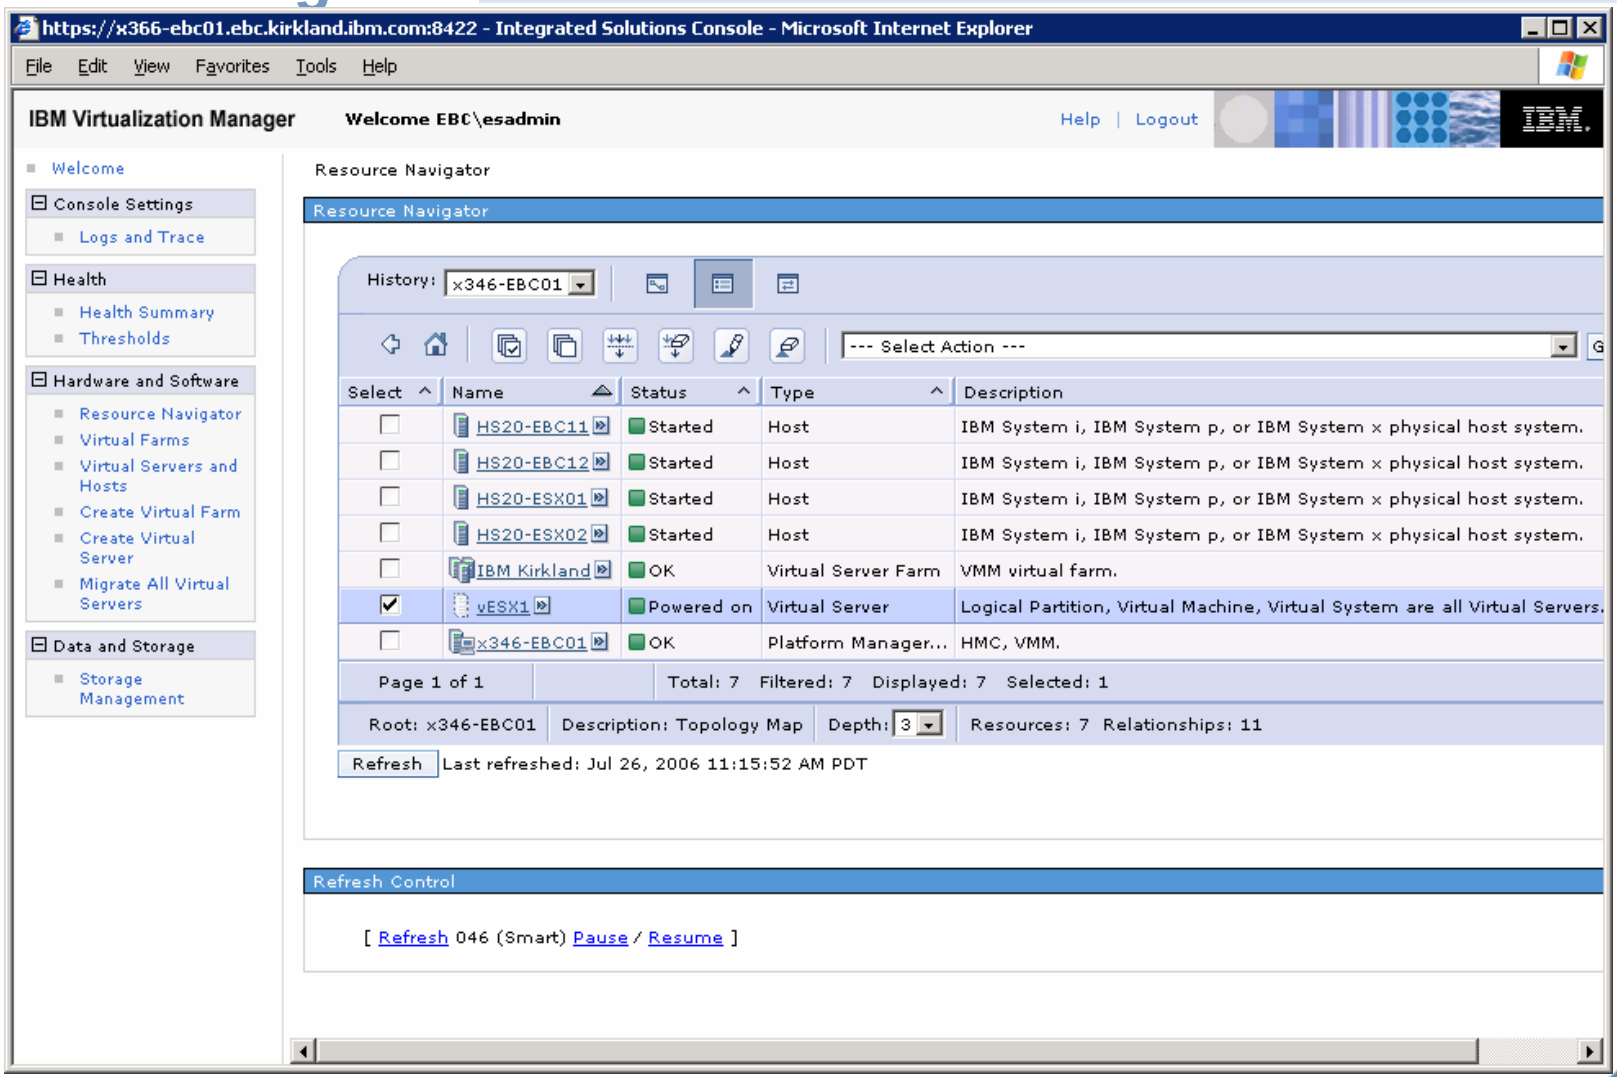

Resource Navigator

- This page shows all virtual servers and their physical hosts, regardless of platform or virtualization technology. You can dynamically track overall health and CPU/memory utilization, as well as run tasks. Some tasks launch IBM Director or the HMC in context.

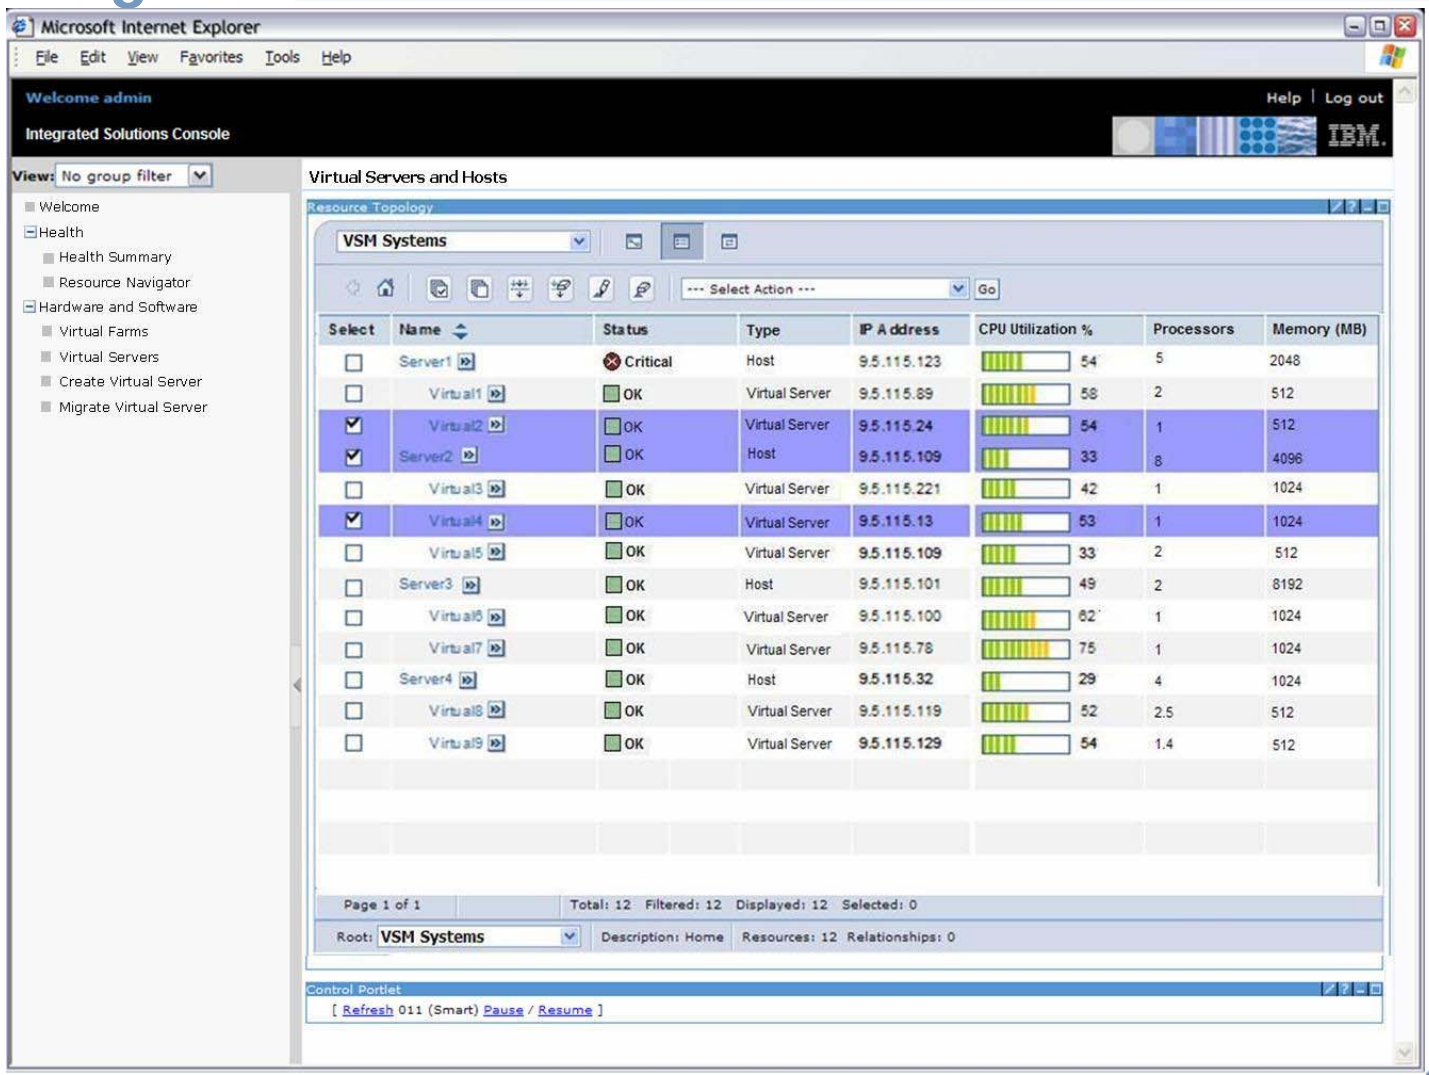

Resource Navigator

- The topology view, a network diagram is displayed. Relationships between components are indicated by colored arrows between systems.

.

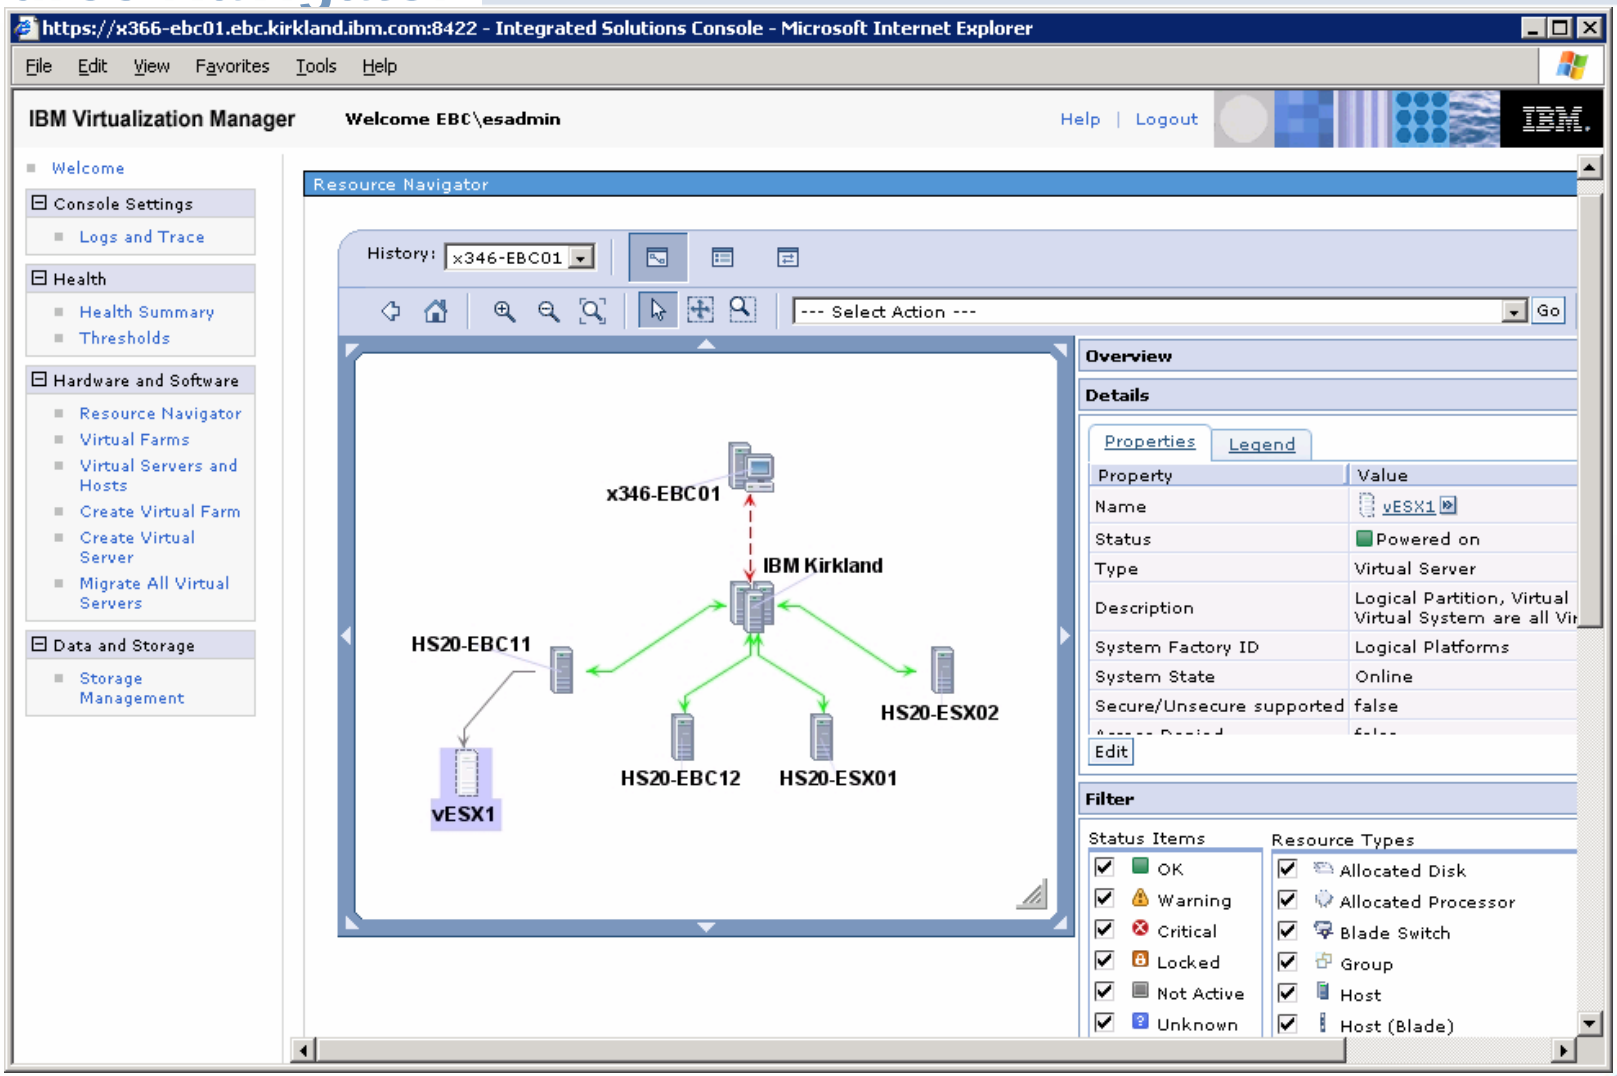

IBM Director Event Filters

- The VSM tree adds new Event Filters for managing virtual environments. Many events unique to these environments can now be detected and acted upon by Event Action Plans.

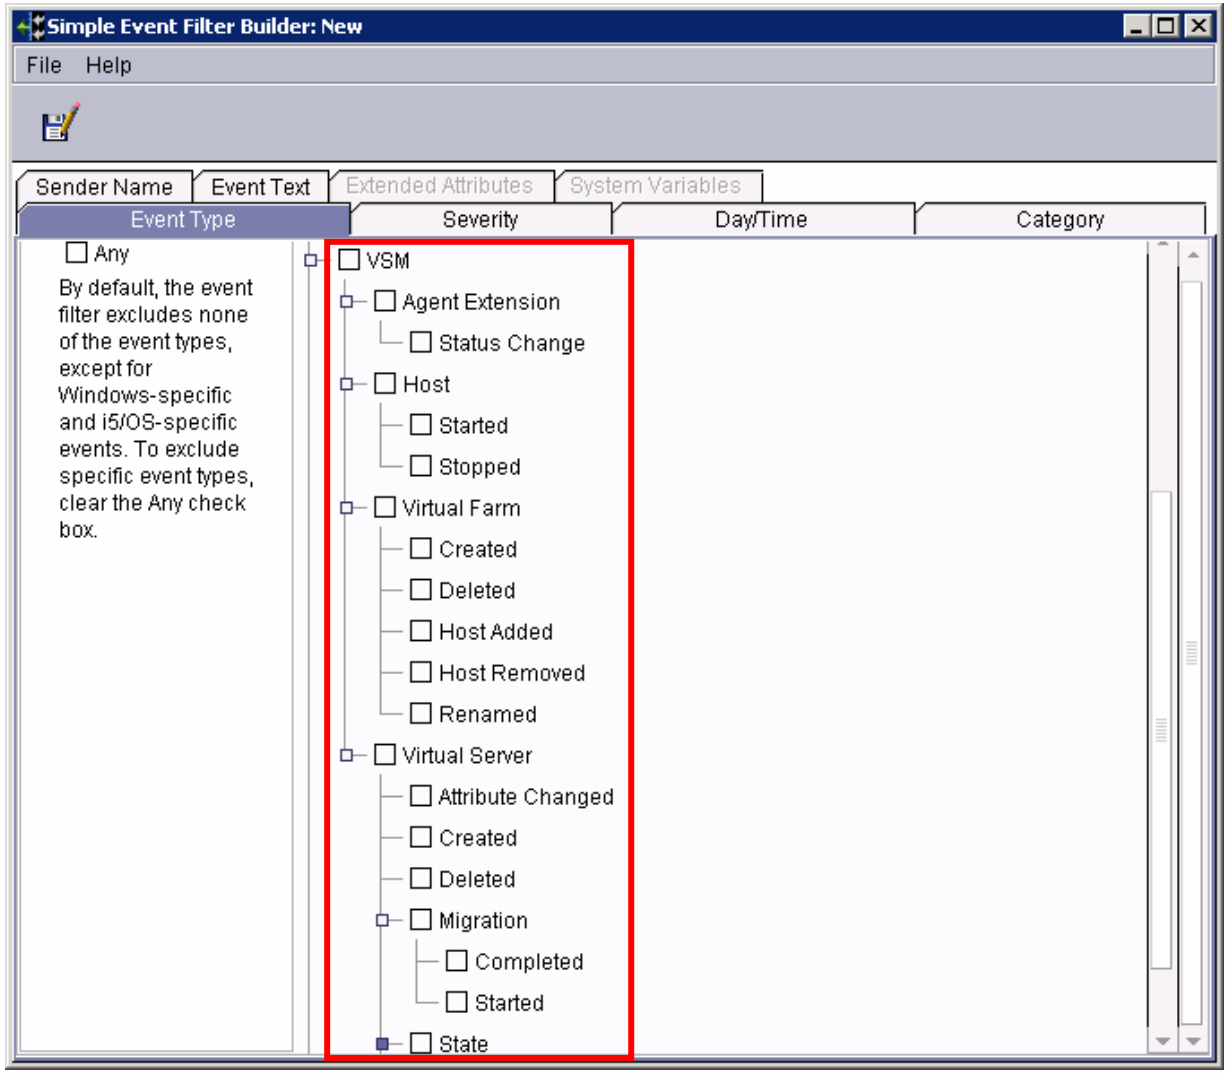

Agenda

- Sizing
  - The elements of sizing
  - IBM's insights into sizing
- IBM Portfolio
  - Portfolio Positioning
  - Competitive Advantage
- Systems Management
  - IBM Virtualization Manager
  - IBM Director
- Summary

IBM has more x86 virtualization experience

- IBM has more experience in virtualization providing a competitive advantage
  - Intelligent virtualization with IBM X-Architecture® technology
  - More Industry Experience for Proper Sizing
  - Outstanding IBM CoolBlue™ Power and Cooling designs
  - Leadership Management software
  - VMware's Largest OEM Vendor(1)

It's that Simple With IBM

Note (1) : Based on bundled license revenue

Trademarks & Legal References

IBM, BladeCenter, ClusterProven, DB2, ServerProven, System x, Virtualization Engine, WebSphere, XpandOnDemand, X-Architecture, Memory ProteXion and xSeries are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml.

Intel and Xeon are registered trademarks of Intel Corporation.

Linux is a trademark of Linus Torvalds in the United States, and/or other countries.

Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation.

VMware and ESX Server are trademarks of VMware, Inc.

All other products may be trademarks or registered trademarks of their respective companies.

The workloads and the associated hardware utilization data and the consolidation data described are representative of the server population that was analyzed in the IBM server consolidation study. The workloads and data are not intended to represent a specific customer environment. To arrange a customized IBM server consolidation study, contact your IBM Sales representative.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven® or ClusterProven®.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

MB, GB, and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology. Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

The following are trademarks or registered trademarks of other companies

NOTES:

(c) 2005 International Business Machines Corporation. All rights reserved.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

MB, GB, and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology. Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Appendix – Backup

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
http://www.vmware.com/vmtn/vmworld/sessions/

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

VMWORLD 2006