



Il Mondo dei Partner
INNOVARE E CRESCERE. INSIEME 2006
Milano 19 - 20 Ottobre **Technical World**

La virtualizzazione su piattaforma x86

- *Overview*
- *VMware Infrastructure 3*

November 2,

2006

Salvatore Morsello
IBM System x™ Field Technical Support Specialist

IBM Systems & Technology Group

Agenda


- **Overview**
 - What is virtualization
 - Benefits
 - Drawbacks
- **VMware Infrastructure 3 – What's new**
 - ESX 3 new features
 - VC 2.0
 - DRS, HA
 - Consolidated backup

Overview

Virtualization Definition

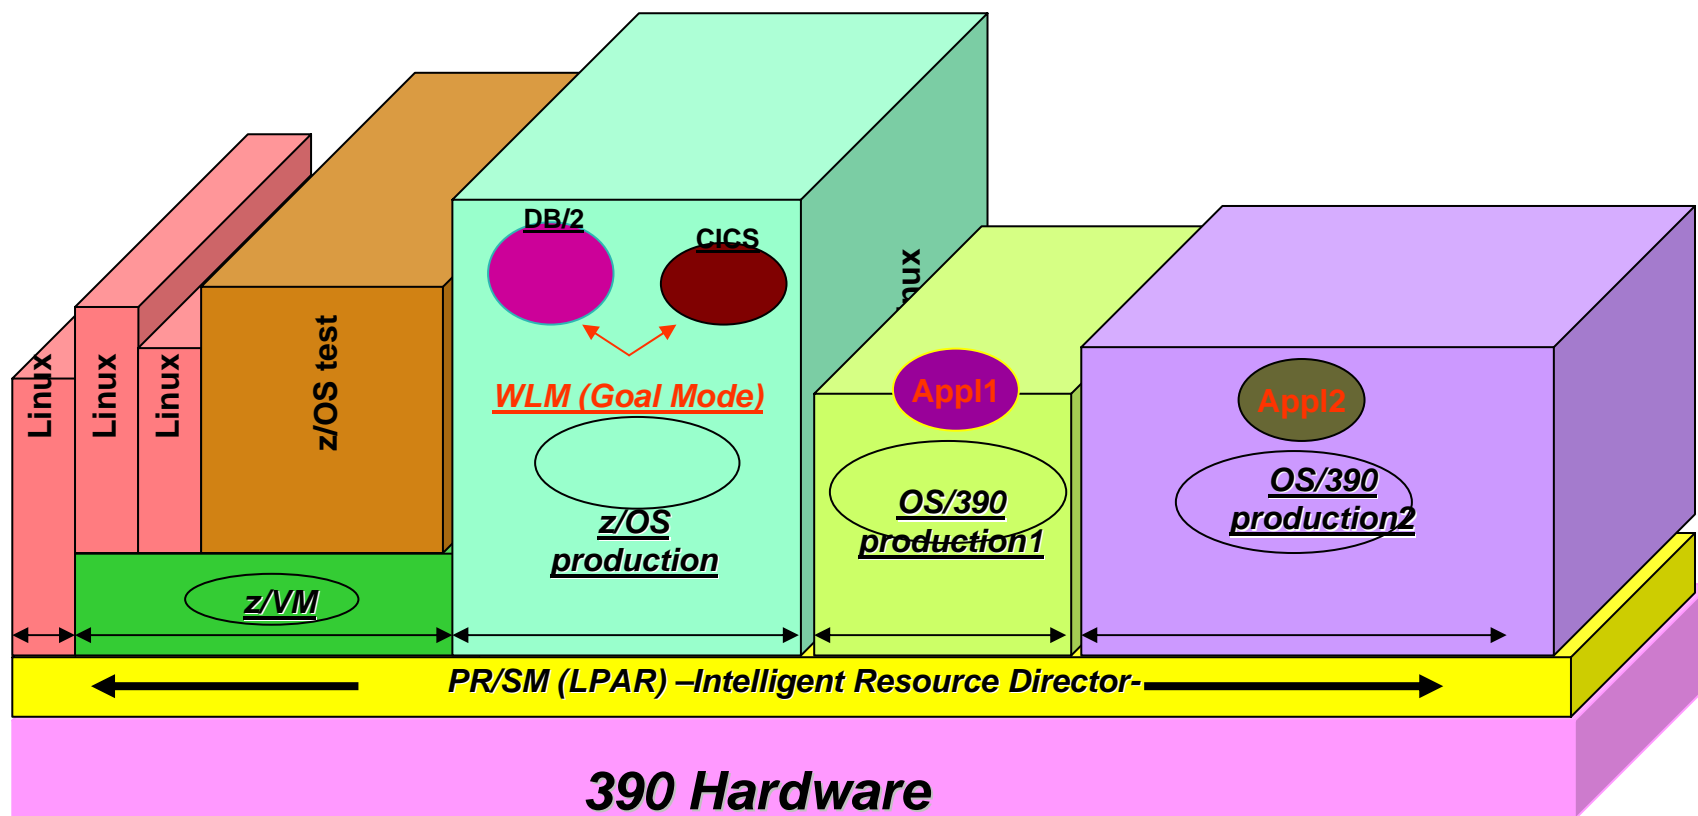
“Virtualization is the process of presenting computing resources in ways that users and applications can easily get value out of them, rather than presenting them in a way dictated by their implementation, geographic location, or physical packaging. In other words, **it provides a logical rather than physical view of data, computing power, storage capacity, and other resources**”

Jonathan Eunice, Illuminata Inc.

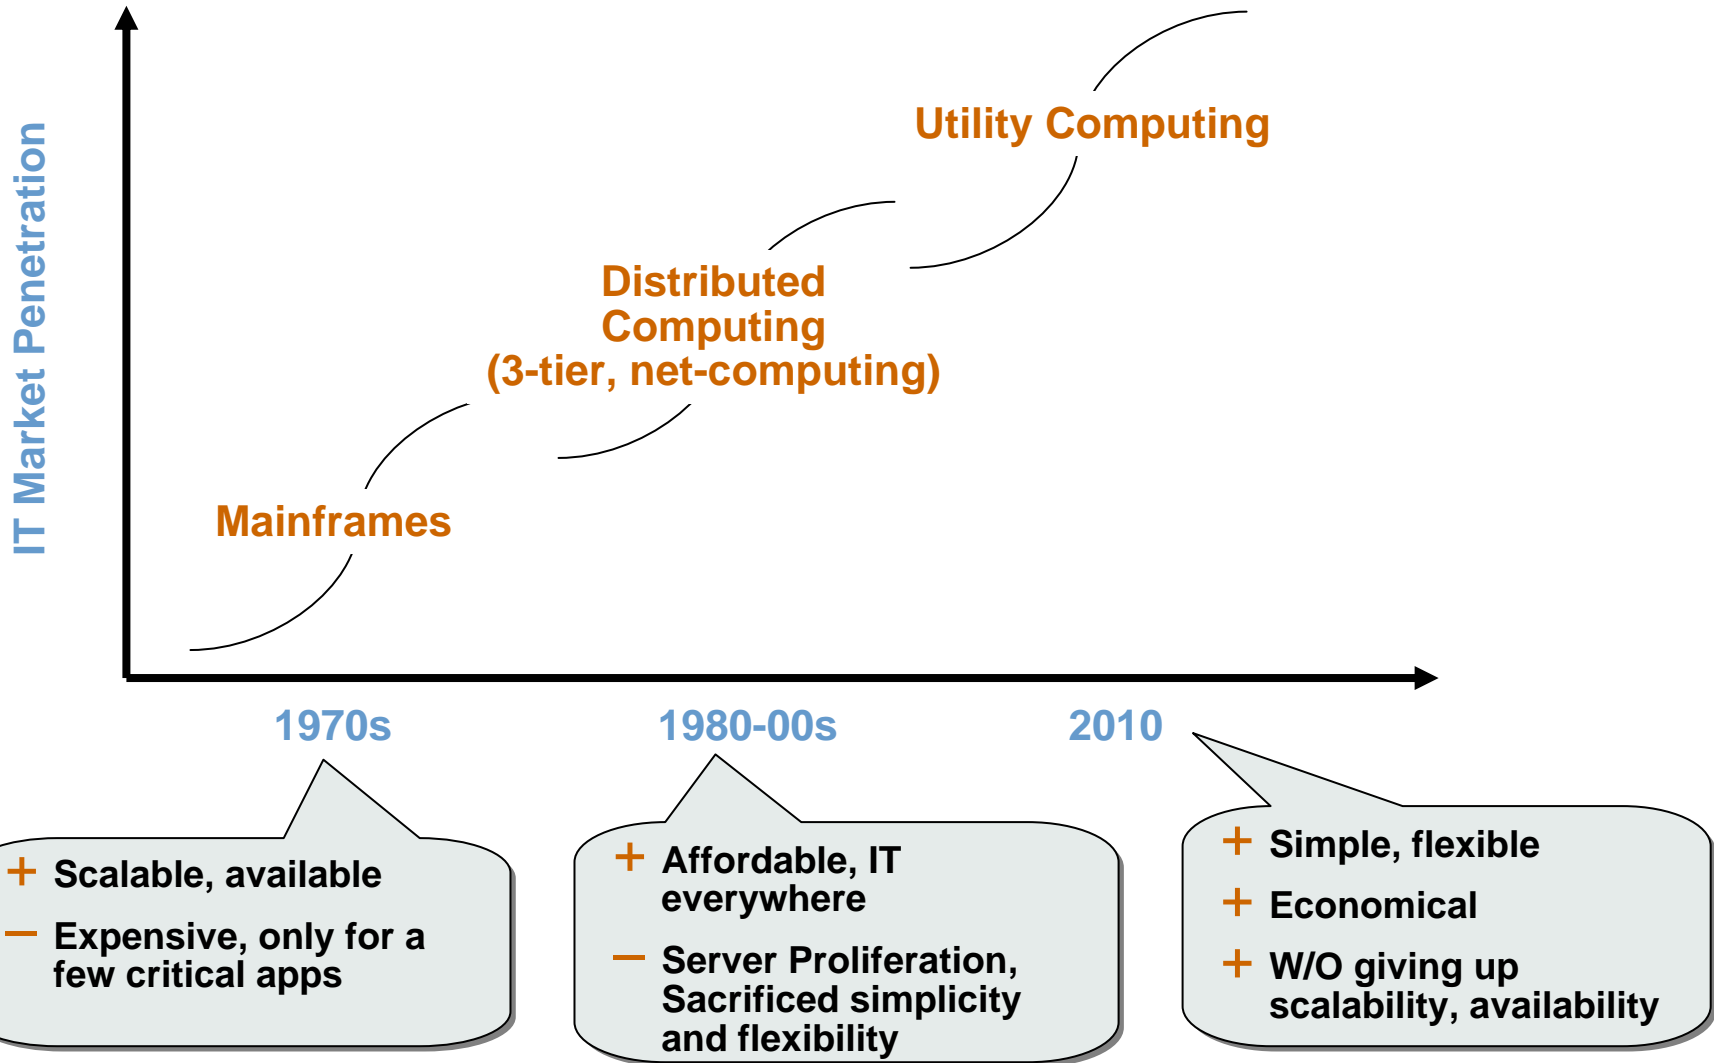


Server Virtualisation – Basic Definition

*“Server Virtualisation - enables **multiple** operating systems and applications to run isolated, concurrently in virtual machines (partitions) on a **single** server “*



Industry Trend



Industry trends – x86 Server Virtualisation becomes pervasive

x86 Virtualisation - Competitive Landscape

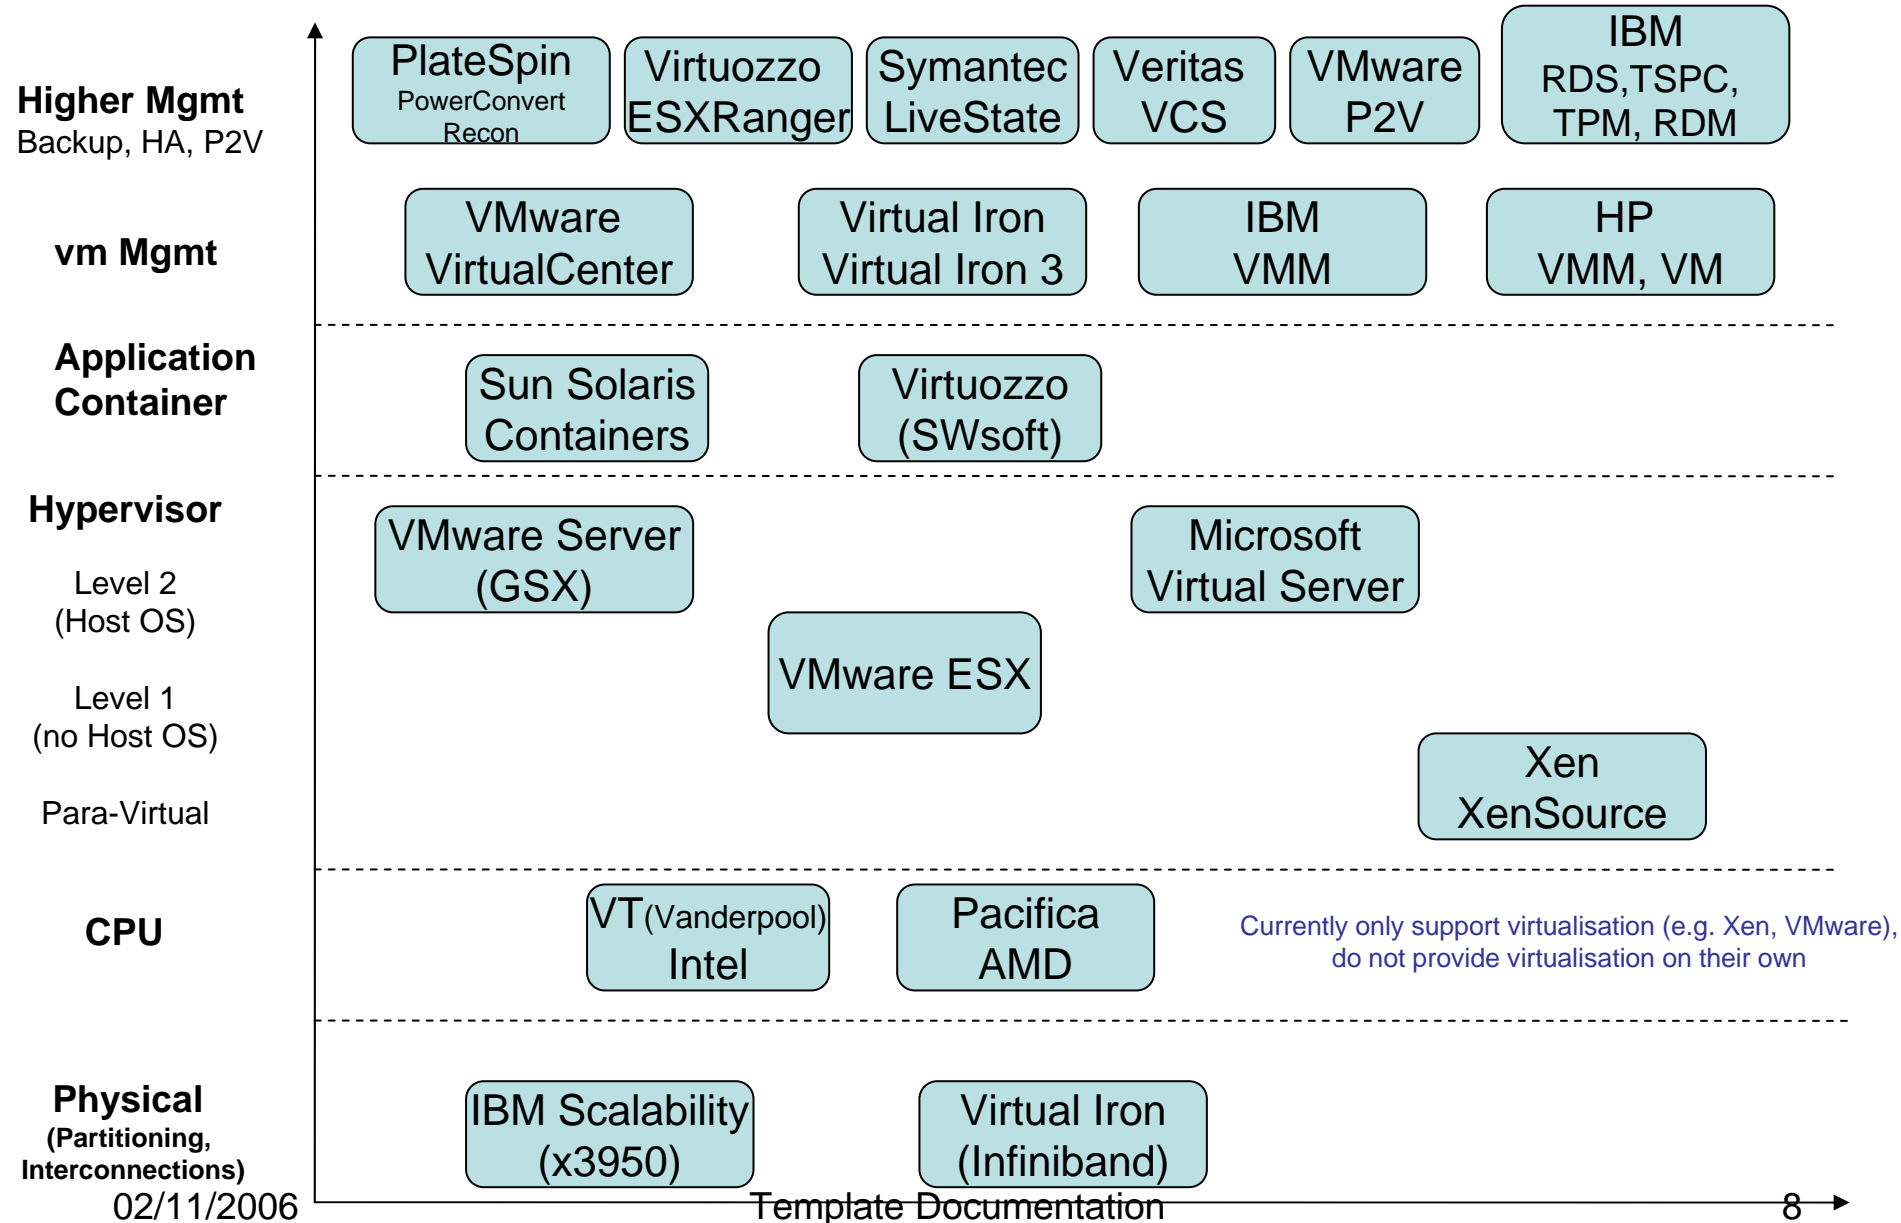
(Hypervisor + Management)



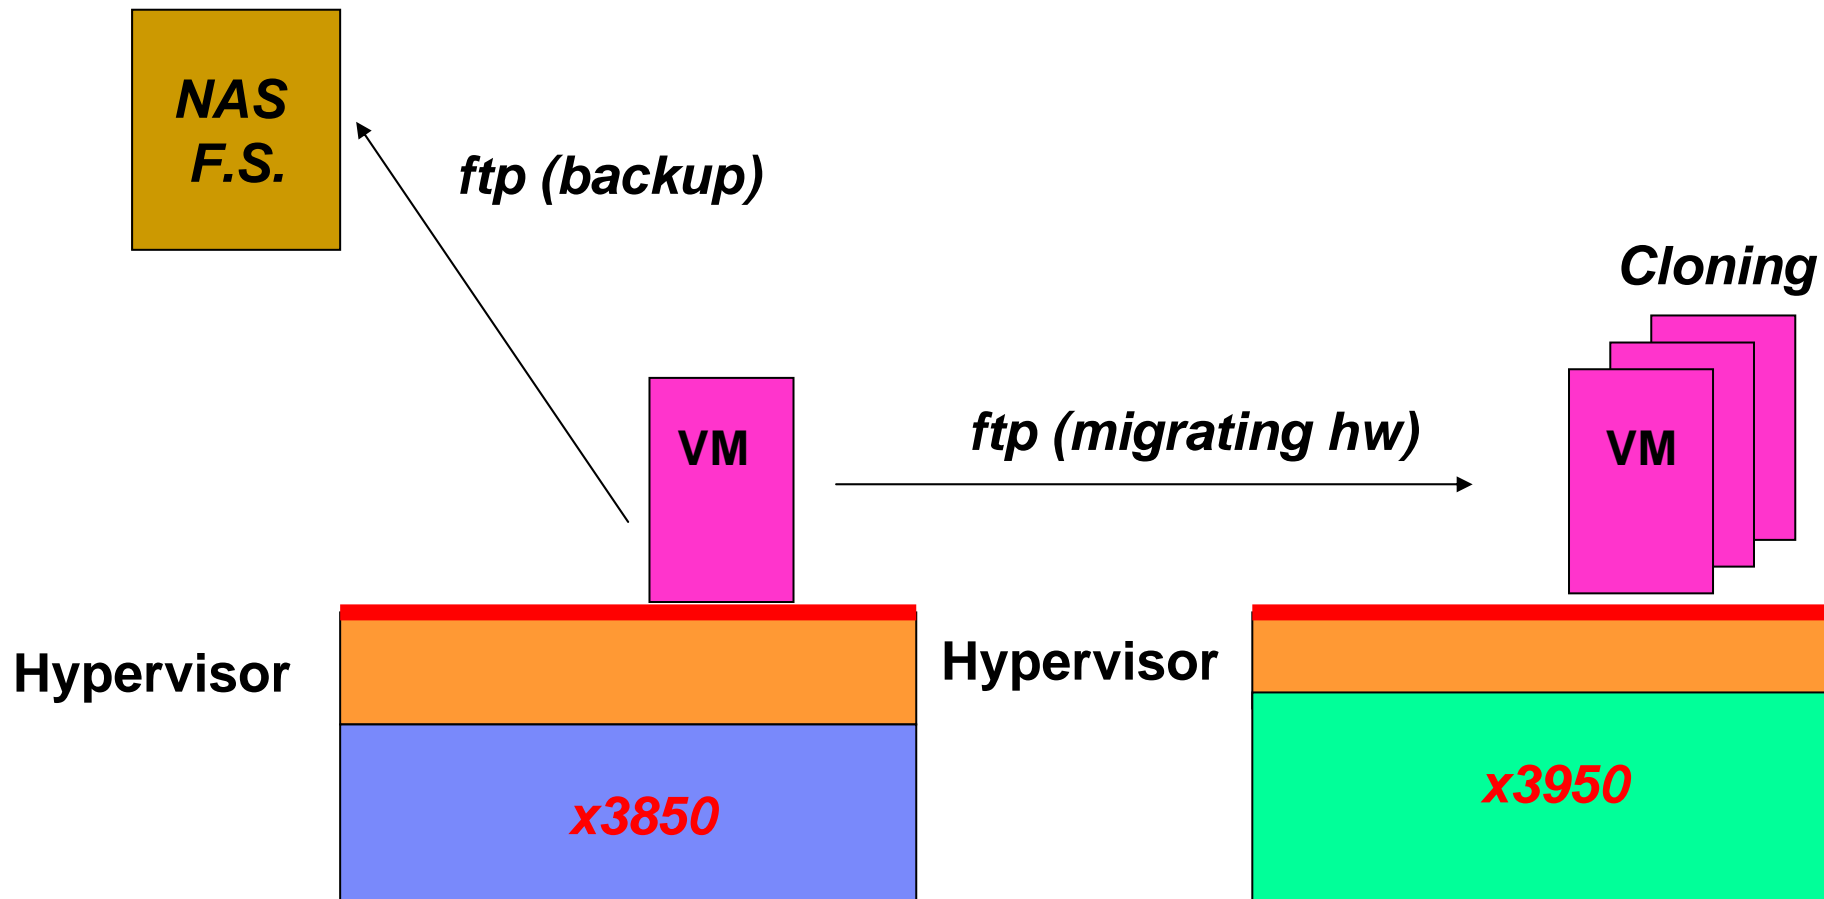
Are they all “doing” the same....?

Layers of the Virtualisation market for x86 Servers

High Level Overview



Benefits - Flexibility



- *Disaster Recovery*
- *Simplified Windows/Linux management*

Benefits – HW Utilization



MS Exchange – W2K



MS SQL – W2K



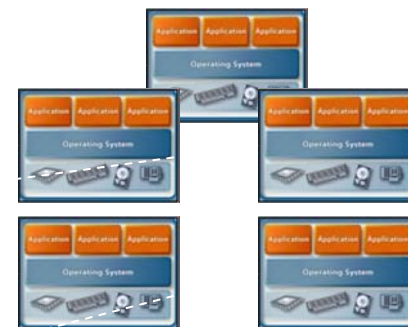
Linux






NT4 SP6



NT4-SP3 +
in-house application requiring SP3



VMware
Virtualization Layer



Benefits – HW Utilization

Virtualization as a Resource Multiplier

All physical resources are shared by virtual machines resulting in a resource multiplier effect



Physical

Virtualized


Physical

Virtualized

4-way Server



Run an average of 32 Virtual Machines !




2 HBAs (paired)



Create and allocate up to 32 Virtual Disks !




6GB Memory



Allocate up to 12 GB Virtual Machine Memory !



2 NICs (paired)



Create and allocate up to 16 Virtual NICs !



x86 Server Utilization Observations – We Know More!

■ 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 CPU Server Workloads

Server	Avg CPU Utilization	Peak CPU Utilization	Server Headroom
2-P	44%	90%	37%
4-P	60%	96%	55%
8-P	68%	96%	65%

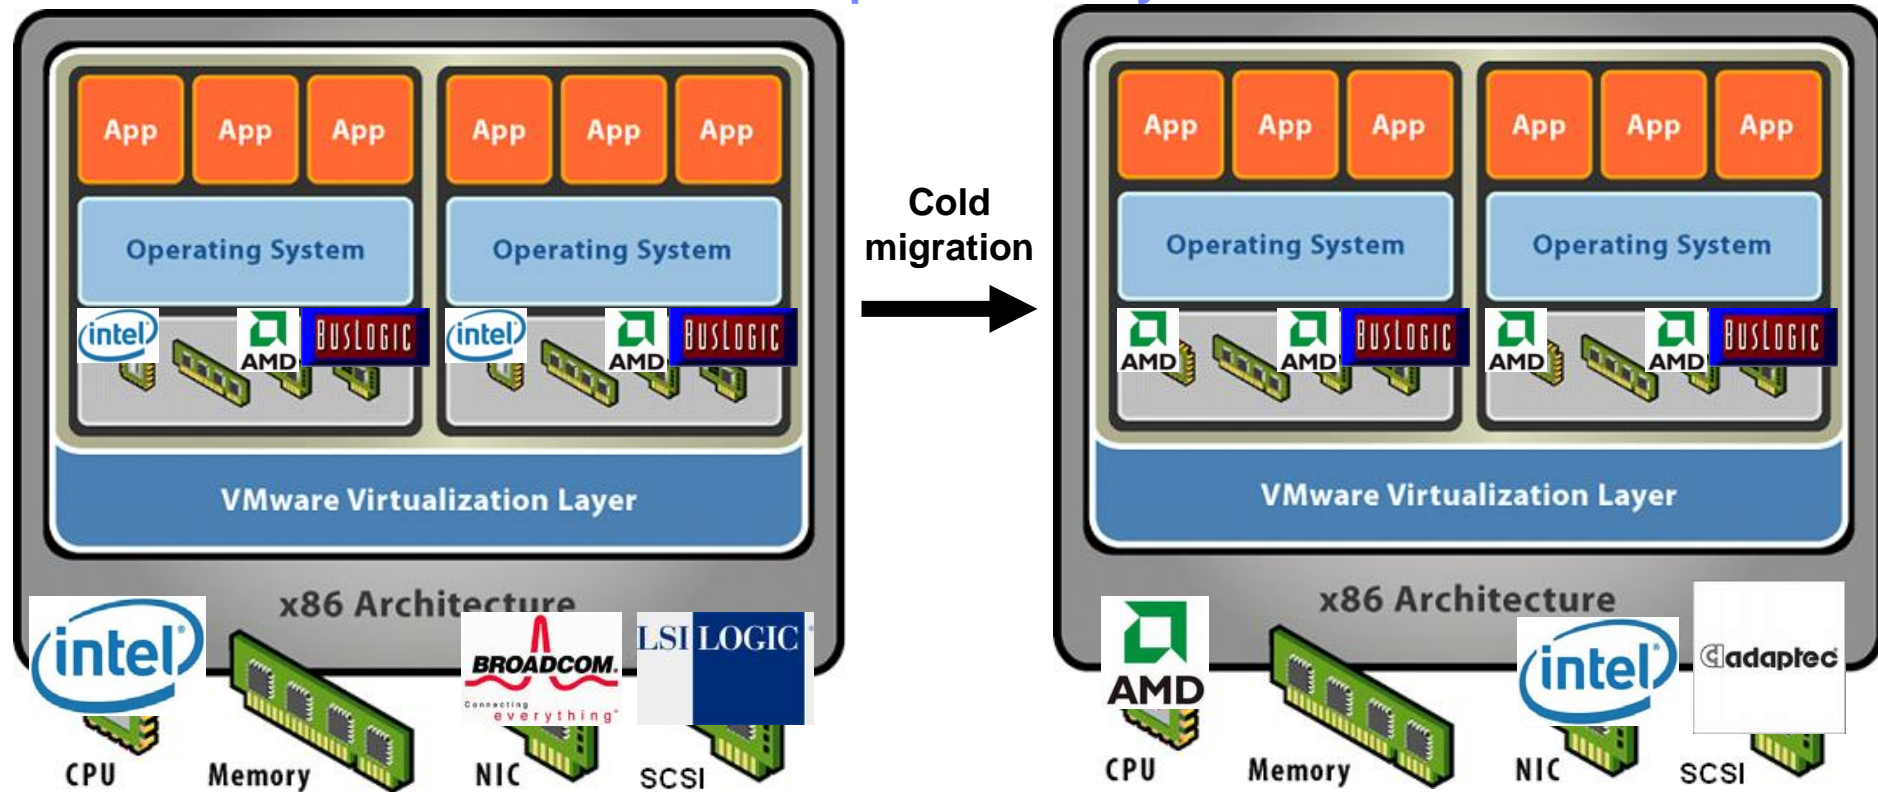
■ Virtualization increases server utilization, but proper configuration must allow for application usage spikes

- 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%

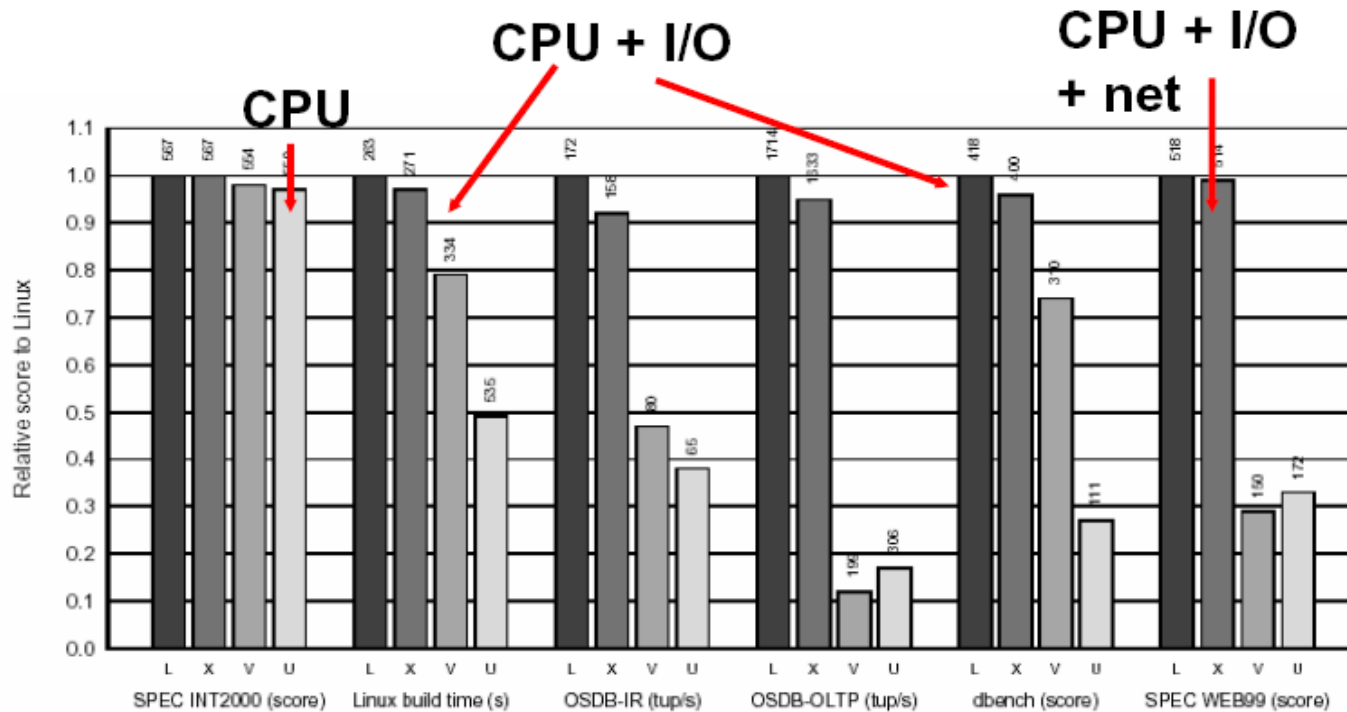
Benefits - Hardware independency



- Hardware changes don't affect Applications and O.S. compatibility

Drawbacks - Overhead


- The virtualization overhead becomes problematic, particularly in the areas of disk and network I/O.
- And it is variable. For sizing it's used a virtualization overhead of 25%, but this is an average.



VMware Infrastructure 3


Virtualization Everywhere!

Data
Center



+

Branch
Office



+

Dev &
Test



- **NAS and iSCSI storage**
- **Expanded hardware compatibility list**
- **4-way Virtual SMP**
- **16GB guest memory**
- **Hot-add virtual disks**
- **Red Hat Enterprise Linux 4 guests**
- **Multiple snapshots**
- **Enhanced performance**
- **Updated Service Console (Red Hat Enterprise Linux 3)**
- **More flexible networking**
- **64-bit guest technology preview**

New – Faster, Bigger Virtual Machines

Run the most resource intensive enterprise applications such as databases, CRM and ERP applications in a virtual machines



ORACLE®



WebSphere software

DB2. Information Management Software










CITRIX®



- **Expanded memory, storage, network limits; up to 4 virtual CPUs per virtual machines and 16GB virtual memory**

Heterogeneous Operating System Support

Freedom to choose the most appropriate OS for any application

	Windows Server 2003 Standard, Enterprise, Web Editions, and Small Business Server
	Windows 2000 Server and Advanced Server
	Windows NT : 4.0 Server
	Windows XP Professional
	Red Hat Linux 7.2, 7.3, 8.0, & 9.0 Red Hat Enterprise Linux 2.1 & 3
	Solaris 10 (on x86) 
	SUSE Linux 8.2, 9.0 and 9.1 SUSE Linux Enterprise Server 8
	Novell NetWare 5.1, 6.0 and 6.5
	FreeBSD 4.9

- **Rigorously tested to run 28 versions of all major operating systems**
- **Experimental 64-bit operating system support**

NAS & iSCSI Details

Lower Cost Access to the Benefits of Virtual Infrastructure

- **NAS**

- NFS client built into ESX Server
- NAS can be used in place of VMFS as the file system for virtual disks

- **iSCSI**

- Use hardware iSCSI card or built-in software iSCSI
- ESX Server boot from iSCSI (for hardware iSCSI only)

- **Both**

- VMotion, Distributed Resource Scheduling, and VMware High Availability all extend to non-fibre channel shared storage
- Storage option transparent to guests

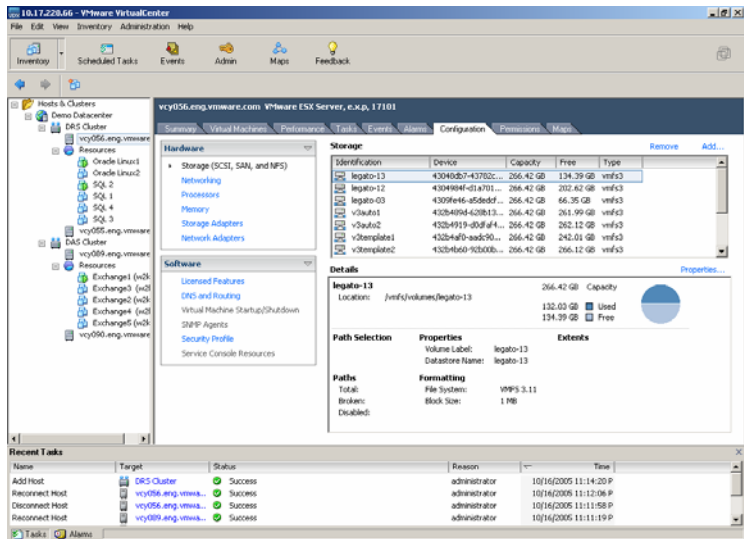
How is NAS used with ESX Server?

- **The VMkernel only supports NFS**
 - More specifically NFS version 3, carried over TCP
- **NFS volumes are treated just like VMFS volumes in Fibre Channel or iSCSI storage**
 - Any can hold VMs' running virtual disks
 - Any can hold ISO images
- **Virtual machines with virtual disks on NAS storage can be VMotioned, subject to the usual constraints**
 - Compatible CPUs
 - All needed networks and storage must be visible at destination

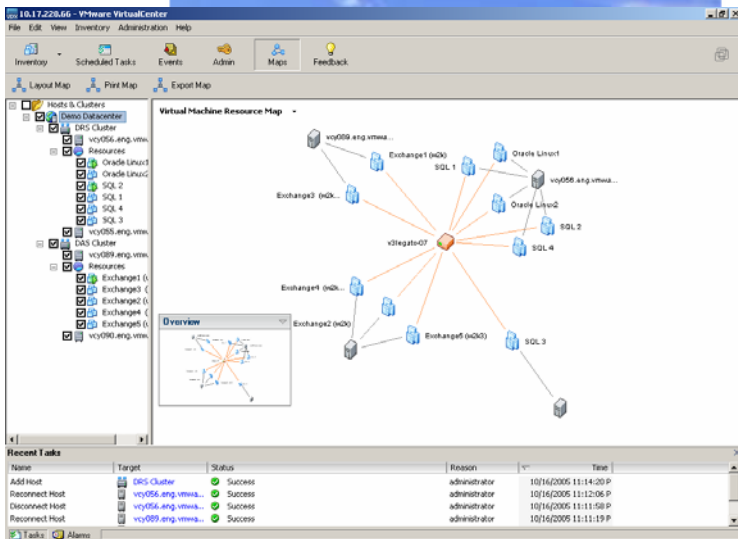
VMware's Implementation of iSCSI

- **Software initiator**
 - Use existing NICs
 - Use native vmkernel stack
 - Used when performance IS NOT an issue, as server and application performance can degrade significantly!
 - Used when cost IS an issue
 - Used when no PCI slots are available
 - Used for simple connectivity to storage or tape backup
- **iSCSI storage adapters (hardware initiator)**
 - Uses less ESX Server resources, especially CPU
 - Initially supported adapter-Qlogic qla4010

Simplified Management



- Single Windows and browser client for ESX Server and VirtualCenter
- ESX Server configuration through VirtualCenter



- Remote devices
- Topology maps
- Centralized licensing
- All VM files (vmx, nvram...) on VMFS

Host Configuration

10.17.228.66 - VMware VirtualCenter

File Edit View Inventory Administration Help

Inventory Scheduled Tasks Events Admin Maps Feedback

Hosts & Clusters

- Demo Datacenter
 - DRS Cluster
 - vcy056.eng.vmware
 - Resources
 - Oracle Linux1
 - Oracle Linux2
 - SQL 2
 - SQL 1
 - SQL 4
 - SQL 3
 - vcy055.eng.vmware
 - DAS Cluster
 - vcy089.eng.vmware
 - Resources
 - Exchange1 (w2k)
 - Exchange3 (w2k)
 - Exchange2 (w2k)
 - Exchange4 (w2k)
 - Exchange5 (w2k)
 - vcy090.eng.vmware

vcy056.eng.vmware.com VMware ESX Server, e.x.p, 17101

Summary Virtual Machines Performance Tasks Events Alarms Configuration Permissions Maps

Hardware

- Storage (SCSI, SAN, and NFS)
- Networking
- Processors
- Memory
- Storage Adapters
- Network Adapters

Software

- Licensed Features
- DNS and Routing
- Virtual Machine Startup/Shutdown
- SNMP Agents
- Security Profile
- Service Console Resources

Storage Remove Add...

Identification	Device	Capacity	Free	Type
legato-13	43048db7-43782c...	266.42 GB	134.39 GB	vmfs3
legato-12	4304984f-d1a701...	266.42 GB	202.62 GB	vmfs3
legato-03	4309fe46-a5dedcf...	266.42 GB	66.35 GB	vmfs3
v3auto1	432b489d-628b13...	266.42 GB	261.99 GB	vmfs3
v3auto2	432b4919-d0dfaf4...	266.42 GB	262.12 GB	vmfs3
v3template1	432b4af0-aadc90...	266.42 GB	242.01 GB	vmfs3
v3template2	432b4b60-92b00b...	266.42 GB	266.12 GB	vmfs3

Details Properties...

legato-13 266.42 GB Capacity

Location: /vmfs/volumes/legato-13

132.03 GB Used

134.39 GB Free

Path Selection **Properties** **Extents**

Volume Label: legato-13

Datastore Name: legato-13

Paths **Formatting**

Total: File System: VMFS 3.11

Broken: Block Size: 1 MB

Disabled:


Recent Tasks

Name	Target	Status	Reason	Time
Add Host	DRS Cluster	Success	administrator	10/16/2005 11:14:20 P
Reconnect Host	vcy056.eng.vmwa...	Success	administrator	10/16/2005 11:12:06 P
Disconnect Host	vcy056.eng.vmwa...	Success	administrator	10/16/2005 11:11:58 P
Reconnect Host	vcy089.eng.vmwa...	Success	administrator	10/16/2005 11:11:19 P

Tasks Alarms


Simplified Service Console

- **Service console I/O more like a virtual machine**
 - Virtual I/O devices for the service console
 - All storage and network devices dedicated to the vmkernel
 - Easier install: no more dividing physical devices between virtual machines and the service console
- **Service console resource needs independent of the number of virtual machines**
 - vmx processes moved from service console to vmkernel
 - Service console not a bottleneck to scalability
 - More accurate virtual machine resource accounting



VMFS3

- **Revamped disk locking**
 - Better scaling for access by large numbers of hosts simultaneously
 - Enables large clusters for distributed resource scheduling and distributed availability services
- **Greater reliability and flexibility**
 - Distributed journaling for faster recovery
 - Logical volume manager
 - Resize LUNs, add LUNs on the fly
 - Volume availability not compromised due to spanning
- **Store more than virtual disks**
 - Exclusive repository for virtual machines and VM state
 - Virtual disks, configuration files, snapshots
 - Directories to organize files
 - Optimized for large and small files
 - Optimized for a large number of files
- **Virtual disk performance remains close to native**



1. Virtualization of Servers, Storage and Networking

VMware Infrastructure

Enterprise Virtualization

Virtual Machines



ESX
Server

ESX
Server

ESX
Server

ESX
Server

ESX
Server


ESX
Server

ESX
Server

Server Farm



Network




Storage



- Partition CPU and memory in multiple virtual machines
- Store virtual machine disks on local or shared storage. VMFS cluster file system
- Build networks within or across ESX Servers.

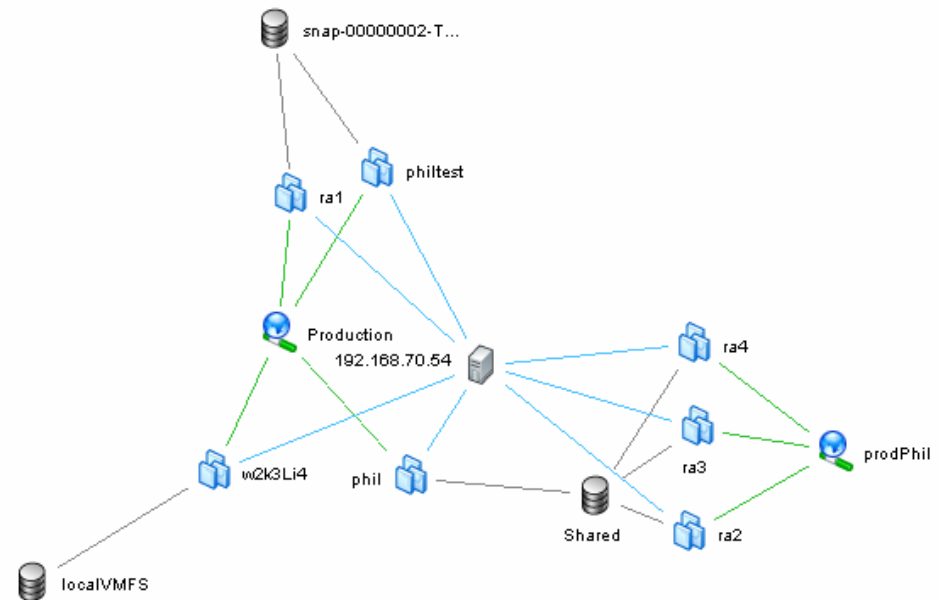
VMware Infrastructure Management



What is new in VirtualCenter 2

■ Virtual Center 2

- Common GUI
- Topology Maps
- New VMotion Capabilities



■ New Services





- Distributed Resource Scheduler
- HA
- Consolidated Backup


VirtualCenter - Key Functionality



Simplify Management - Topology Maps

Legend

-  = *Datastore*
-  = *vSwitch*
-  = *VM*
-  = *Managed host*



Map Relationships:

Virtual Machine Resources ▾

Host Options

- Host To VM
- Host To Network
- Host To Datastore

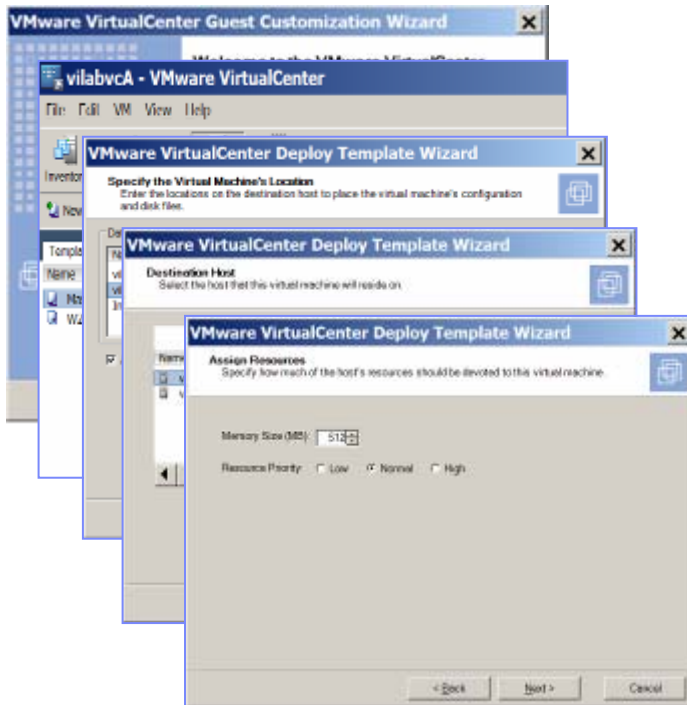
VM Options

- VM to Network
- VM to Datastore

Apply Relationships

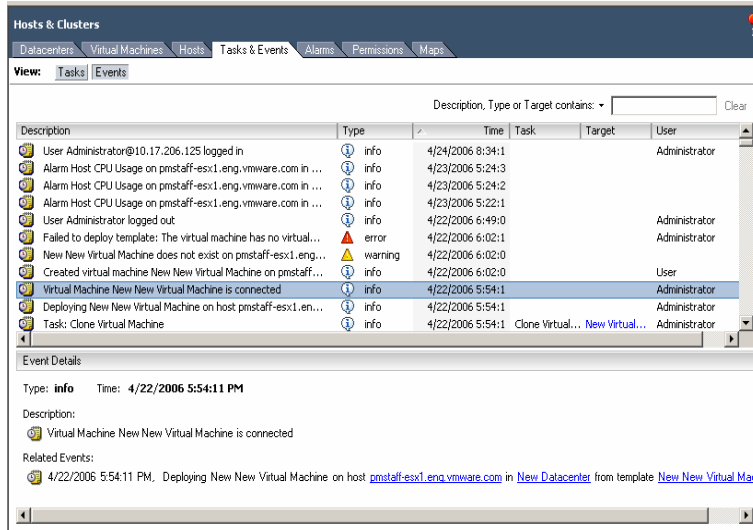
Virtual Machine Provisioning and Migration

Provision infrastructure instantly



- **Deployment Wizard**
- **New – Redesigned virtual machine templates**
 - Support easy virtual machine patching and updating by converting to VMs, powering-on, and then re-instating as templates
 - Templates are always stored in a VMFS or NAS (no longer on VC)
- **Virtual machine cloning**
- **Cold migration**
 - Drag and drop
- **Live migration with VMotion**

New – Security Enhancements

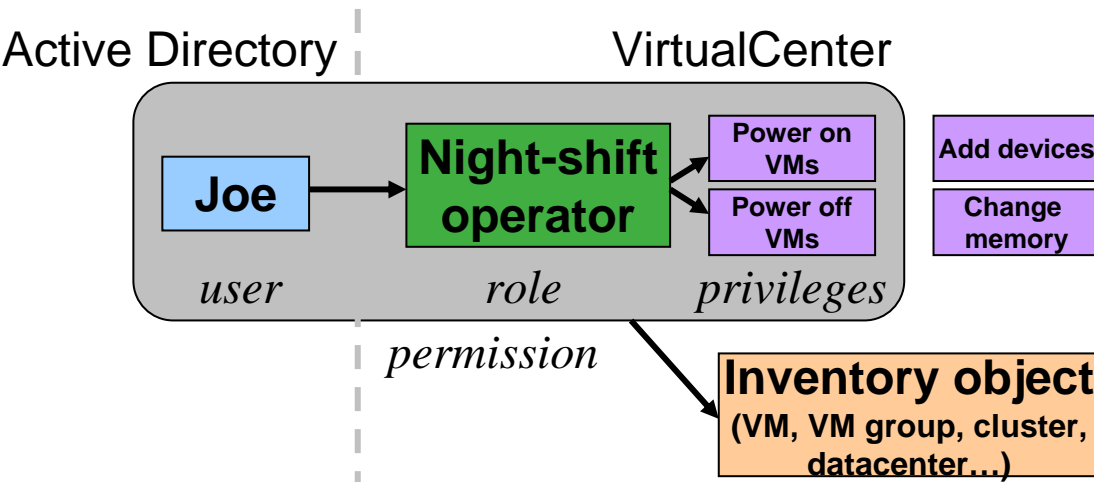


Audit trails

- Maintain a record of significant configuration changes
- Export reports for event tracking.

Custom roles and permissions

- Fine-grain control over user groups and privileges
 - User defines roles as a set of privileges
 - Down the tree propagation / privilege-inheritance is optional
 - Propagated access can be limited
- Delegate administration tasks down in the organization



Enterprise Standards - Custom Roles & Permissions

The screenshot displays the 'Add Role' dialog in the vpx interface. The dialog is divided into several sections:

- Roles List:** A list of predefined roles including 'No Access', 'Read-Only', 'Administrator', 'Virtual Machine Administrator', 'Datacenter Administrator', 'Virtual Machine Provider', 'Virtual Machine Power User', and 'Virtual Machine User'.
- Role Name:** A text input field containing 'My Custom Role'.
- Permissions Tree:** A hierarchical tree of permissions. The 'All Privileges' checkbox is checked. Under 'Host.Config', the 'Upgrade' checkbox is highlighted.

The permissions tree structure is as follows:

- All Privileges
 - Global
 - Folder
 - Datacenter
 - Host.Inventory
 - Host.Config
 - Connection
 - Maintenance
 - Upgrade
 - AutoStart
 - HyperThreading
 - Storage
 - NetService
 - Memory
 - Network
 - AdvancedConfig
 - DiskLease
 - Resources
 - Host.Local

Enterprise Standards - VirtualCenter Audit Trails

Hosts & Clusters 1

Datcenters Virtual Machines Hosts **Tasks & Events** Alarms Permissions Maps

View: Tasks **Events**

Description, Type or Target contains: Clear

Description	Type	Time	Task	Target	User
User Administrator@10.17.206.125 logged in	info	4/24/2006 8:34:1			Administrator
Alarm Host CPU Usage on pmstaff-esx1.eng.vmware.com in ...	info	4/23/2006 5:24:3			
Alarm Host CPU Usage on pmstaff-esx1.eng.vmware.com in ...	info	4/23/2006 5:24:2			
Alarm Host CPU Usage on pmstaff-esx1.eng.vmware.com in ...	info	4/23/2006 5:22:1			
User Administrator logged out	info	4/22/2006 6:49:0			Administrator
Failed to deploy template: The virtual machine has no virtual...	error	4/22/2006 6:02:1			Administrator
New New Virtual Machine does not exist on pmstaff-esx1.eng...	warning	4/22/2006 6:02:0			
Created virtual machine New New Virtual Machine on pmstaff...	info	4/22/2006 6:02:0			User
Virtual Machine New New Virtual Machine is connected	info	4/22/2006 5:54:1			Administrator
Deploying New New Virtual Machine on host pmstaff-esx1.en...	info	4/22/2006 5:54:1			Administrator
Task: Clone Virtual Machine	info	4/22/2006 5:54:1	Clone Virtual... New Virtual...		Administrator

Event Details

Type: **info** Time: **4/22/2006 5:54:11 PM**

Description:

Virtual Machine New New Virtual Machine is connected

Related Events:

4/22/2006 5:54:11 PM, Deploying New New Virtual Machine on host [pmstaff-esx1.eng.vmware.com](#) in [New Datacenter](#) from template [New New Virtual Ma](#)

Virtual Switches

reston - VMware Host Agent

File Edit View Inventory Help

Inventory Events Feedback

reston.vmware.com VMware ESX Server, e.x.p, 0

Summary Virtual Machines Performance Users/Groups System Logs Events **Configuration**

Hardware

- Storage (SCSI, SAN, and NFS)
- ▶ **Networking**
- Processors
- Memory
- Storage Adapters
- Network Adapters

Software

- Licensed Features
- DNS and Routing
- Virtual Machine Startup/Shutdown
- SNMP Agents
- Security Profile
- Service Console Resources

Networking [Add Networking...](#)

vSwitch1 [Remove...](#) [Edit...](#)

- Virtual Machine Port Group: Accounting Network (0 VMs | VLAN 103)
- Virtual Machine Port Group: Engineer Network (0 VMs | VLAN 102)
- Virtual Machine Port Group: Marketing Network (0 VMs | VLAN 101)
- Virtual Machine Port Group: Customer Support (0 VMs | VLAN 104)

Network Adapters:

- vmnic5 (disabled)
- vmnic4 100 Full
- vmnic3 100 Full

vSwitch2 [Remove...](#) [Edit...](#)

- VMotion and IP Storage Port: VMotion (192.168.10.10)

Network Adapters:

- vmnic6 1000 Full

vSwitch3 [Remove...](#) [Edit...](#)

- VMotion and IP Storage Port: EMC NAS Network (192.168.20.20)


Network Adapters:

- vmnic8 100 Full
- vmnic7 (disabled)

Tasks

Enterprise Standards - Performance Graphs

- **Redesigned to include more data, allow customization**
 - Subsumes details previously only in tools such as vmkusage and esxtop
 - Objects in the inventory and their metrics can be selected for display
- **Several levels of granularity & time-intervals**
 - Real-time statistics at a 20-second sampling rate
 - Archived statistics for the past hour, day, month, etc.
 - OR for a specific time interval specified
- **Fully exportable to Excel or HTML format**



FlexNet Licensing Changes

- **New licensing mechanisms use FlexNet license files.**
 - In order to upgrade from ESX 2 / VirtualCenter 1 to ESX 3 / VirtualCenter 2 customers will need update product binaries and product licenses
- **Encrypted content inside license files determines the type of functionality:**
 - Feature types, evaluation vs. perpetual, & quantity (per-processor) enabled
- **2 Different types of license files available for download:**

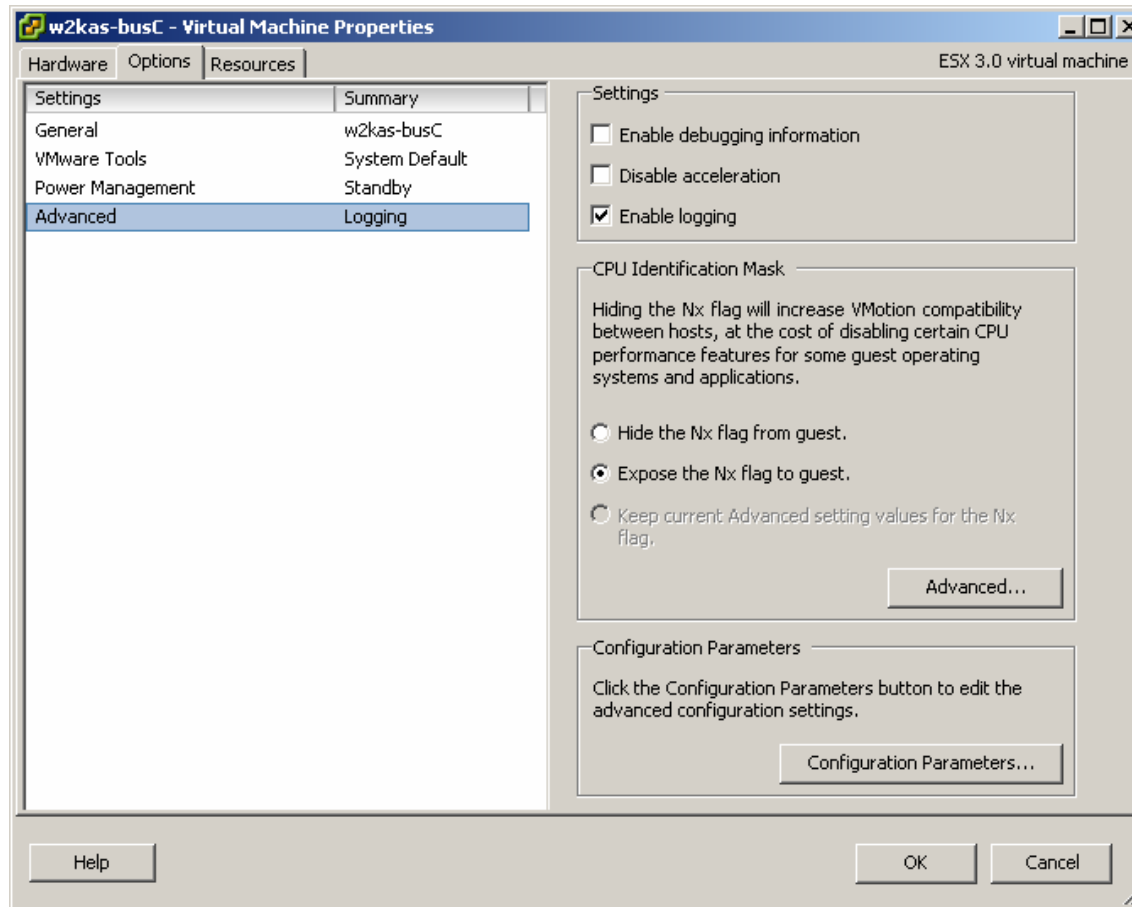
Served

- Requires license server
- Centralizes & simplifies license management in larger environments
- Available for ESX Licenses and VirtualCenter Add-on features (like VMotion, DRS, HA)

Host-Based

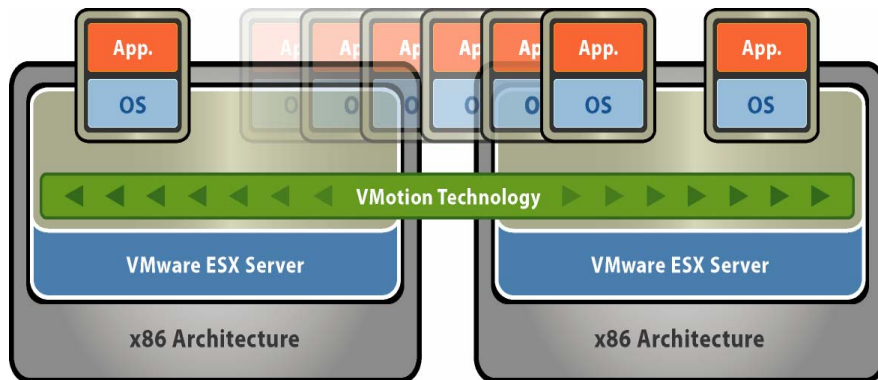
- Does not require license server
- Suitable for small deployments, without VirtualCenter
- Only available for ESX Licenses

VMotion Compatibility



Live Migration Of Virtual Machines with VMotion

70% of VMware customers have implemented VMotion




- What is it?
 - Live migration of virtual machines with VMware VMotion
- Customer Impact
 - Zero downtime
 - Continuous service availability
 - Complete transaction integrity
 - Supported on Fibre Channel and iSCSI SAN and NAS

Clusters

The Power of Many Hosts, the Simplicity of One

- **Reduce management complexity by combining stand-alone hosts into clusters for higher availability and more flexible resource controls**



Stand-alone hosts

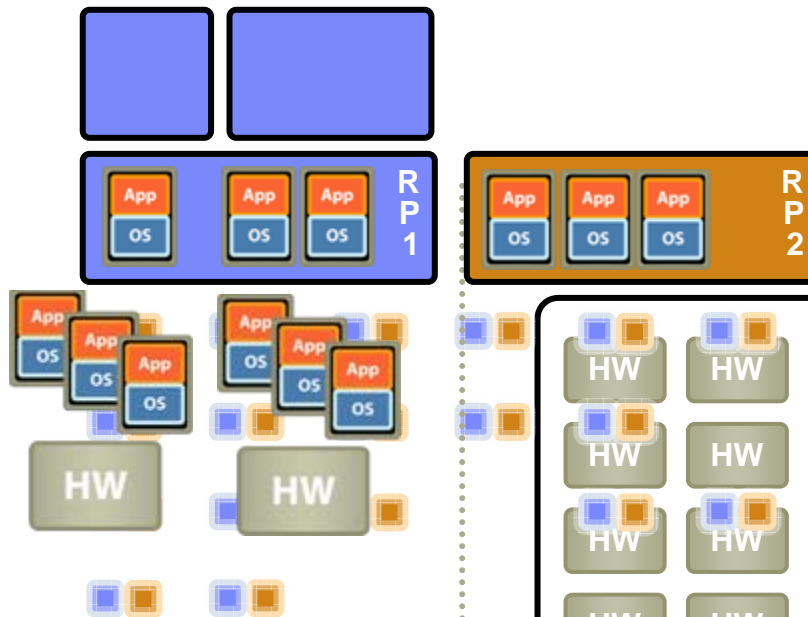
Cluster

Transforming hardware and capacity management

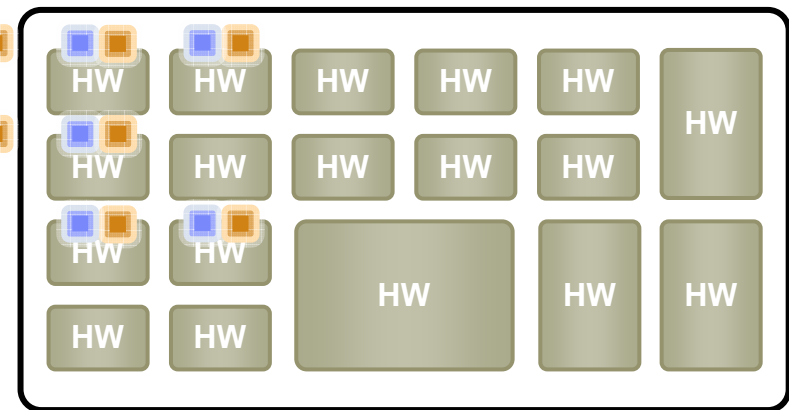
PHYSICAL



VIRTUALIZED



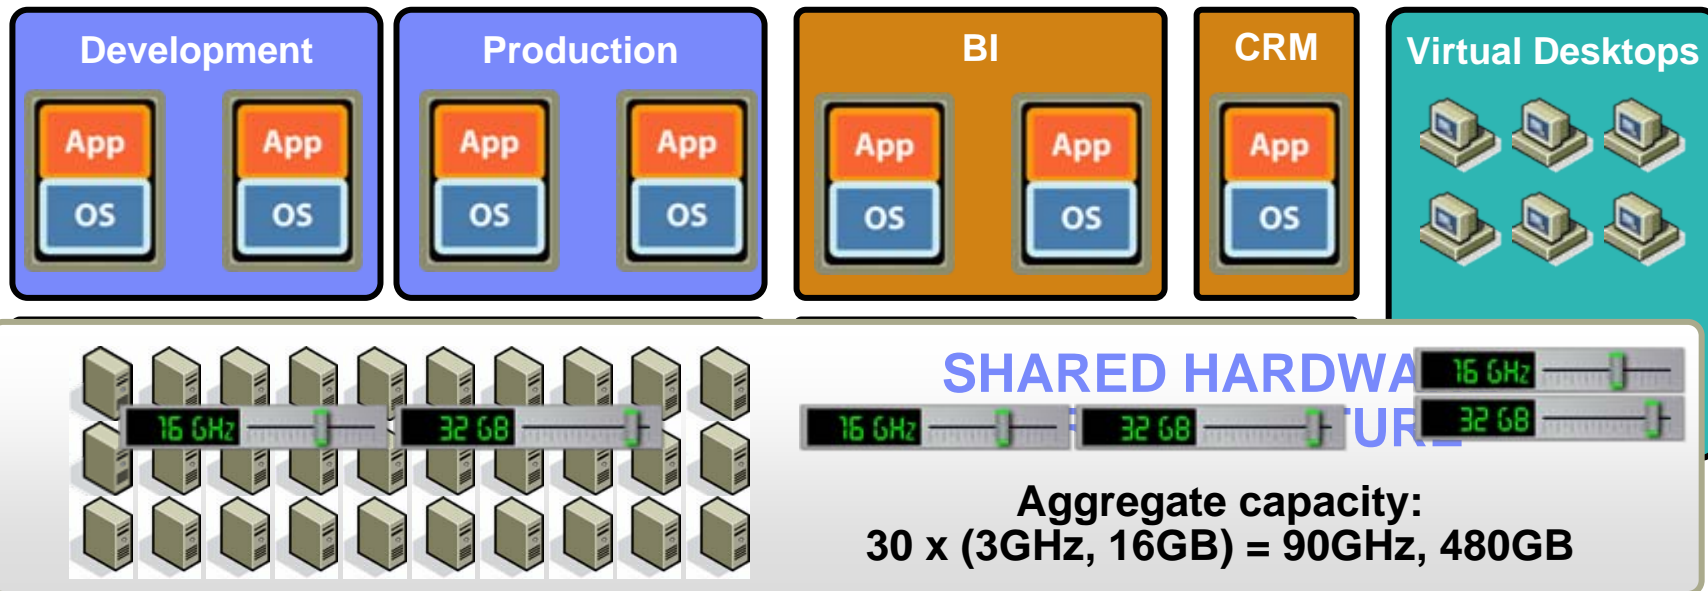
POOLED



INDUSTRY FIRSTS:

- Logical Resource Pooling (RP)
- Distributed Resource Scheduler (DRS)

Align and scale capacity to business needs




Resource Pools

Virtual Machine Containers with Assigned Resources

- **With resource pools you can...**


- Let a user create and run as many virtual machines as desired while limiting the total resources used
- Instantly add extra resources to an enterprise application
- Delegate control over assigning resources to virtual machines while maintaining complete control over hardware



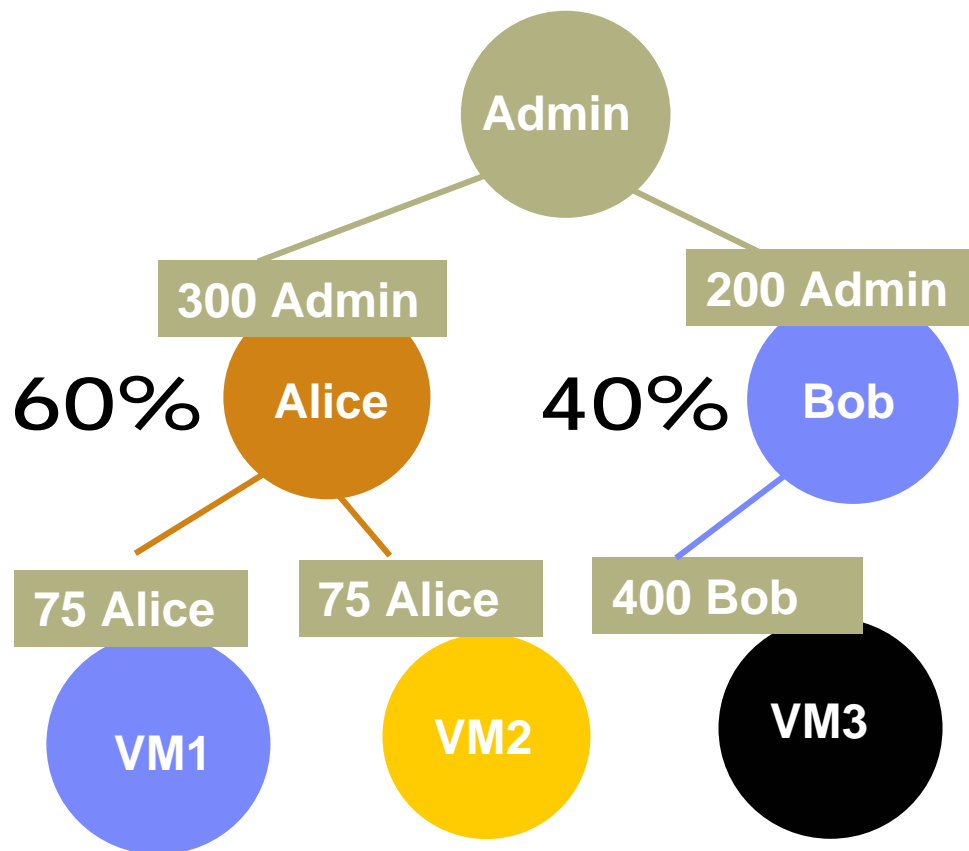
Resource Pools

Precise Resource Control


- **Virtual machines draw resources from their resource pools**
- **Resource allocations can be changed dynamically**
- **Resource pools can be nested**



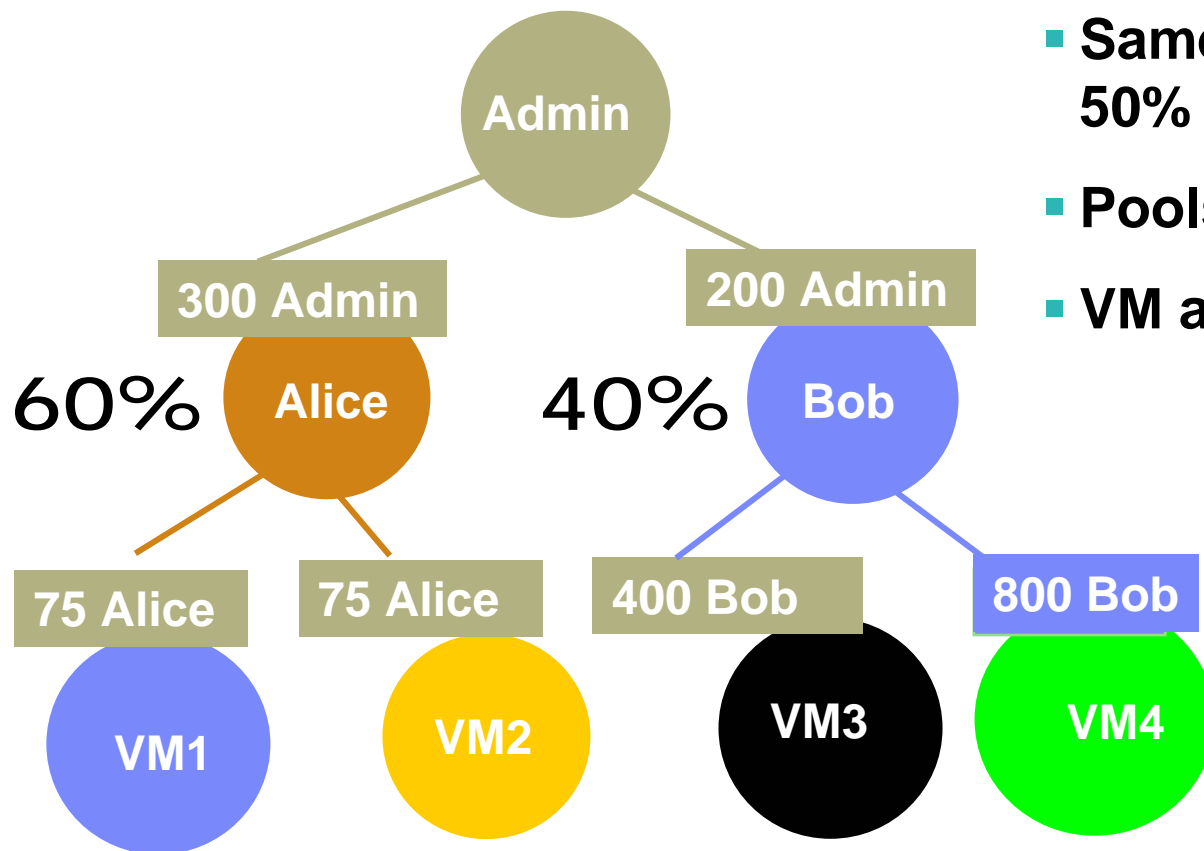
Resource Pools Example




- Admin manages users
- Policy: Alice's share 50% more than Bob's
- Users manage own virtual machines
- Not shown: min, max
- VM allocations:



Example: Bob Adds Virtual Machine



- Same policy: Alice still gets 50% more than Bob
- Pools isolate users
- VM allocations:




DRS Can Help You...

- **Manage variable loads**
 - Workloads often dynamic, time-dependent
 - Quickly shift loads in response to demand
- **Administer many virtual machines**
 - Hierarchical organization
 - Delegated administration
- **Move towards utility computing**
 - Think more about aggregate resource pools
 - Think less about individual hosts

New – Resource Optimization with VMware DRS

Dynamic and intelligent allocation of hardware resources to ensure optimal alignment between business and IT


Business Demand



Resource Pool

- What is it?
 - Dynamic balancing of computing resources across resource pools
 - Intelligent resource allocation based on pre-defined rules
- Customer Impact
 - Align IT resources with business priorities
 - Operational simplicity; dramatically increase system administrator productivity
 - Add hardware dynamically to avoid over-provisioning to peak load
 - Automate hardware maintenance

New - Capacity on Demand with VMware DRS




DRS - Key Features

- **Virtual machine placement**
 - Choose initial host when VM powers on
 - Dynamic rebalancing using VMotion
- **Configurable automation levels**
 - Manual – recommend initial host and migrations
 - Partial – automatic initial host, recommend migrations
 - Full – automatic initial host and migrations
- **Resource pools**
 - Flexible grouping, sharing, and isolation
 - Hierarchical organization and delegation

New – Ensure High availability with VMware HA

VMware HA enables cost-effective high availability for all applications



■ What is it?


- Automatic restart of virtual machines in case of server failure

■ Customer Impact

- Cost effective high availability for all applications
- No need for dedicated stand-by hardware
- None of the cost and complexity of clustering

New – Protect data with VMware Consolidated Backup

Centralized file level backup enables easy & reliable data protection



■ What is it?

- Centralized agentless backup for virtual machines
 - **Move backup out of the virtual machine**
 - **Eliminate backup traffic on the local area network**
- Pre-integrated with major 3rd-party backup products

■ Customer Impact

- Perform backup in the middle of the day



How Does Consolidated Backup Work?

- **3rd Party Backup SW schedules backup job**
 - Job represents single or group of virtual machines
- **Consolidated Backup's pre-backup script runs**
 - Create filesystem *quiesced virtual machine snapshot(s)*
 - *Make data available* on proxy
- **3rd Party Backup SW performs backup**
- **Consolidated Backup's post-backup script runs**
 - Remove mount from proxy
 - Commit backup snapshot(s)

THE END

Legal

IBM, BladeCenter, ClusterProven, DB2, ServerProven, System x, Virtualization Engine, WebSphere 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.

Linux is a trademark of Linus Torvalds. Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation. All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

(c) 2006 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.

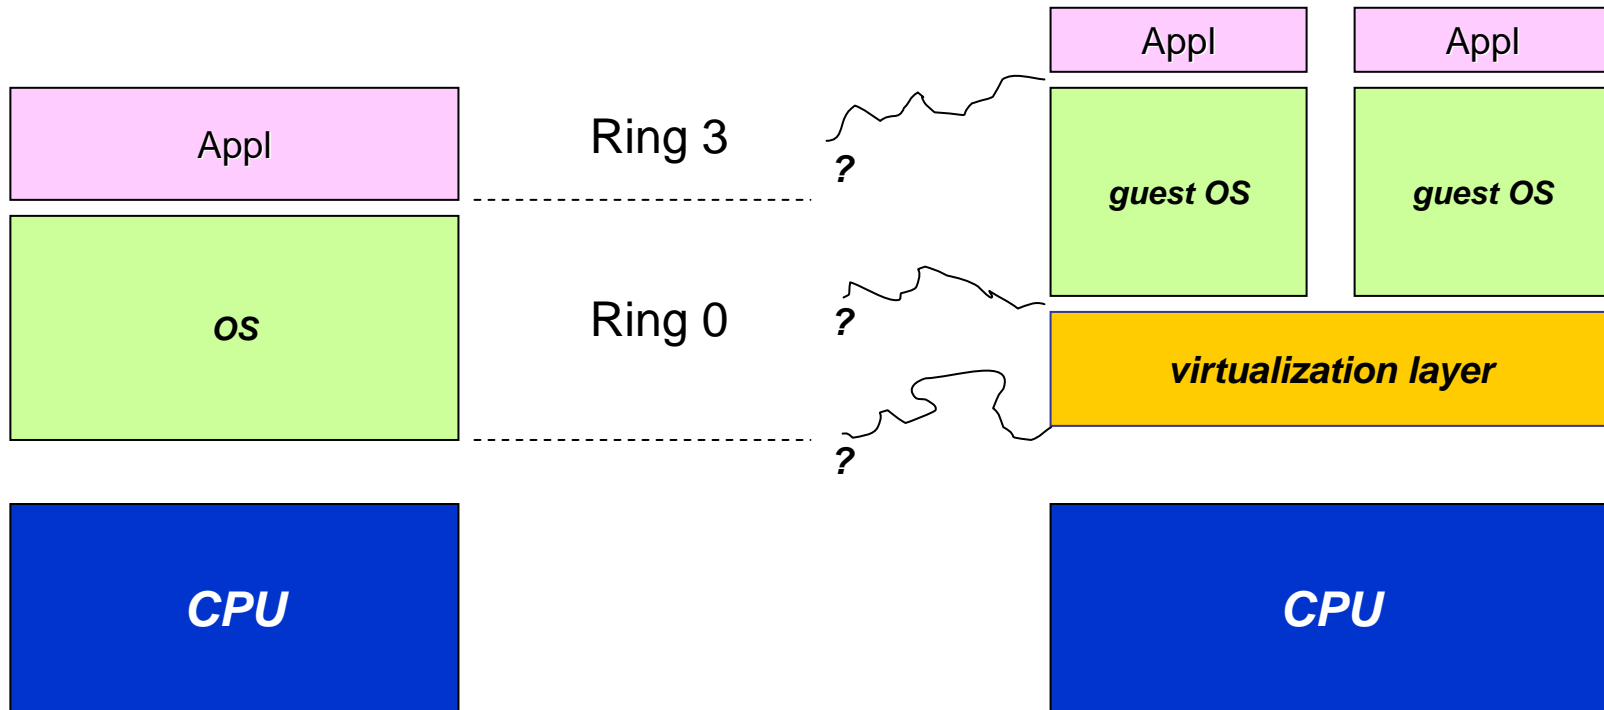
The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this 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.

Backup charts

Hardware-Assist for Virtualisation

x86 Virtualisation – The Problem




- We have been doing this for 20+ years
- OS designed to “own” the hardware
- Applications run with lower priority

- latest trend (virtualization)
- OS still designed to own the hardware
- OS does NOT own the hardware any more!
- Result = workarounds need to be found

Hardware Assist

Intel and AMD introduced x86 specific instructions to assist virtualization software

Intel calls it “VT” while AMD calls these enhancements “Pacifica”

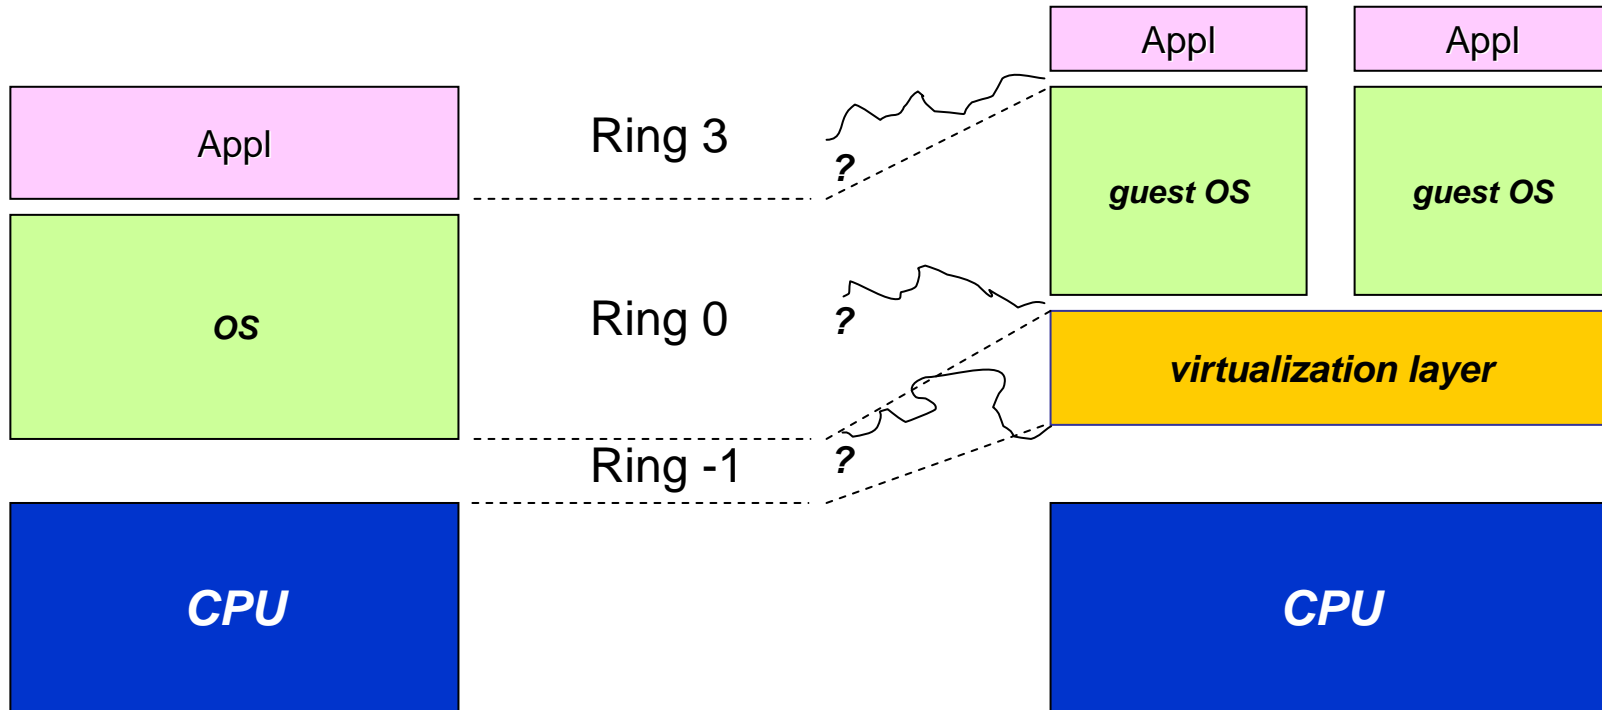


VT ships today on desktops as well as on some servers.

Why is hardware-assist important?

- Ease the development of virtualization products (Xen)
- Improve Stability

The Solution



The hypervisor runs in an extremely highly privileged ring that did not exist before



Hardware Assist

Myth

- VT/Pacifica are hypervisors

- Hardware Assist will dramatically improve performance

Reality

- In order to run virtual machines you need virtualisation software software (Xen, VMware)

- Hardware Assist's primary task is not to improve performance but to increase stability and security


...So what are we really gaining from it then ...?

- Ability to run unmodified guest OSs with Xen
- Ability to run 64-bit guest OSs

VirtualCenter 2 cont - DRS

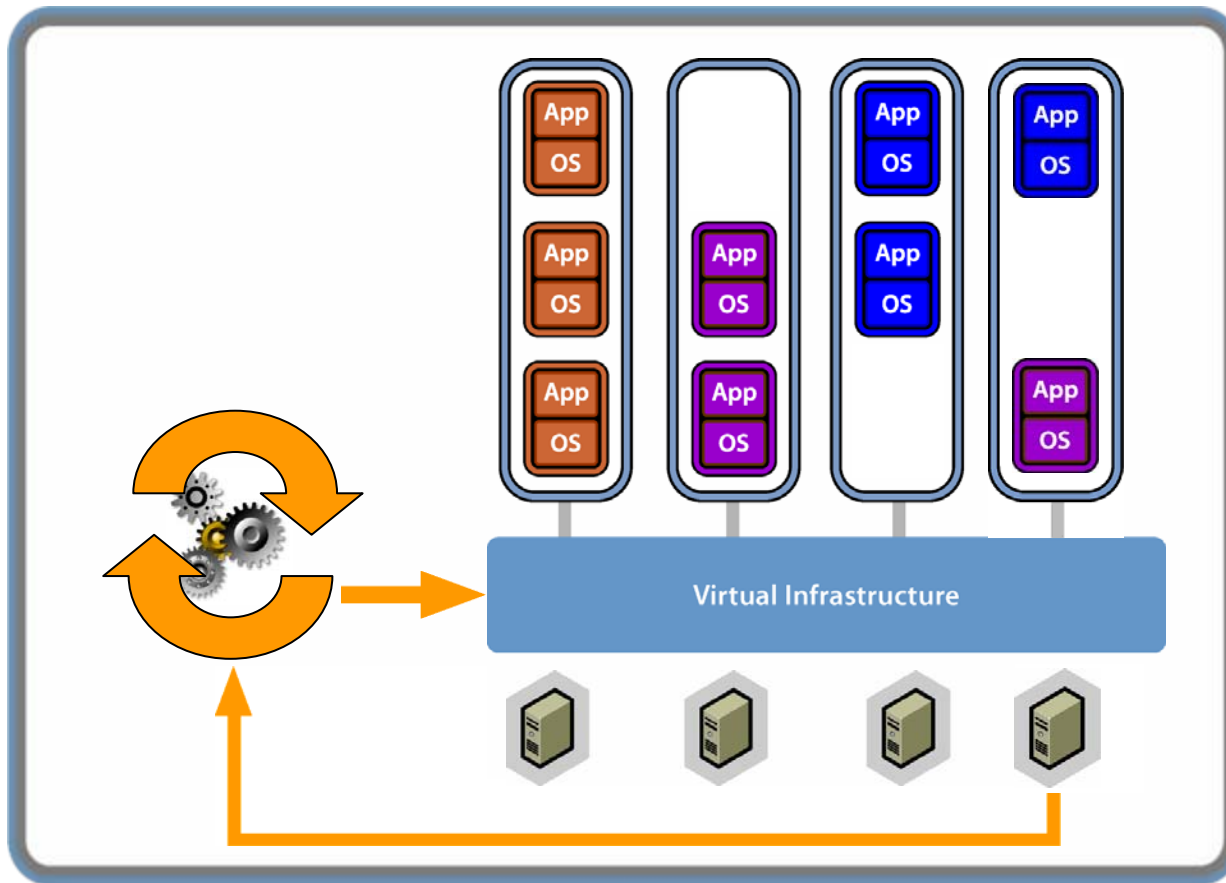
■ DRS

- (fee-based) plug-in for Virtual Center
- Automatic virtual machine placement
- Cluster-wide resource management, Resource Pools
- Policy based VMotion
- 32 hosts, LAN – not WAN



VirtualCenter 2 cont - DRS

- **Instant capacity on demand**
 - Combine with bare-metal provisioning



VirtualCenter 2 cont - HA

HA

- ▶ (fee-based) plug-in for Virtual Center
- ▶ Automatic "failover" of virtual machines between physical ESX servers
- ▶ Placement optimised by global scheduler (in conjunction with DRS)

None of the complexity of "classic" clustering, OS independent

