



# Server Virtualization: Changing the Rules...



# Technology Briefing Objectives

- The introduction of **Server Virtualization** brings large benefits in the areas of cost savings, improved productivity, reliability and flexibility.
- It provides you with new options, and is making fundamental shifts in the way that organizations architect, design, deploy and manage server resources.
- This presentation focuses on the technologies that are available from VMware today, and how these technologies are "**Changing the Rules**" ...


# Areas to be Explored Include:

- What is a Virtual Machine?
- What are its Attributes and Benefits?
- Capacity Planning
- Server Provisioning and Deployment
- Disaster Recovery and High Availability
- Technologies that are Available Now...
- Resources

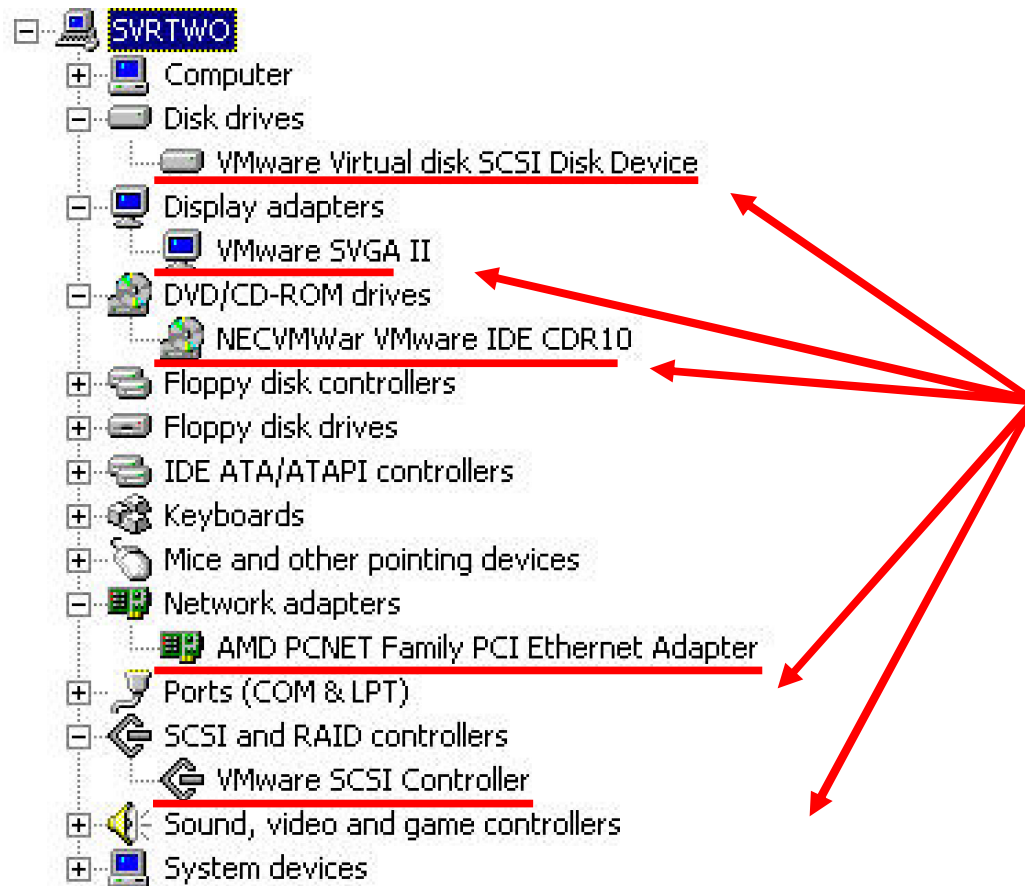
# What is a Virtual Machine?

- For the purposes of this presentation, consider a **Virtual Machine**, or **VM**, to be a software construct that appears as if it were a physical server.
- Multiple VM's can run on a **single physical server**.
- An OS doesn't know it is running in a VM, and sees generic hardware.
- A VM is isolated from other VM's. If one VM fails, the others are unaffected.

# What's in a VM?



# What's in a VM? (Windows example)



# What does a Virtual Machine consist of?

A Virtual Machine consists of 3 files:

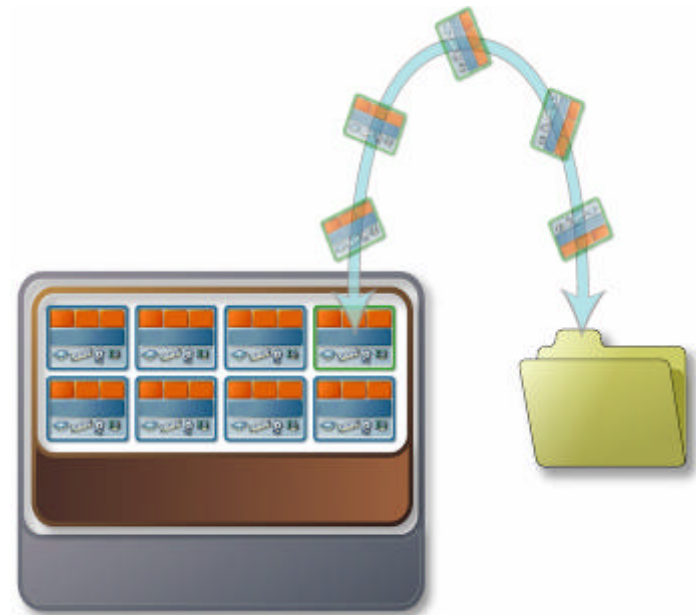
- A small configuration file, that defines the VM's hardware attributes.
- A small file that contains the “**Virtual BIOS**” settings called nvram.
- A **Virtual Disk** file, that contains all the files that are on the boot/system volume. This file is the size of the original disk or partition.

If you move these 3 files, you can move the offline VM from one physical server to another.

# System Encapsulation

Entire state of the VM is encapsulated:

- Memory, disk images, I/O device state.
- VM state can be saved to a file
  - “Checkpointing”, aka “Suspend / Resume”.
  - VM state can be transferred through time and space.
  - Time: store in a file.
  - Space: transfer over a network.





# Attributes of a Virtual Machine

- **Portability** - As the OS running within a VM does not know the real hardware that it is running on, it can be easily moved from one physical hardware platform to another.
- **Suspend Mode** - Machine state can be captured for root cause analysis.
- **Point in Time (PiT) Restoration** - Data can be saved to easily back out of changes.

# Benefits of VM Portability-Simplification

- Easy server **backup and restore**. A VM can be restored on a different physical server, with no changes, as it does not see any differences to the underlying hardware.
- **Disaster Recovery** – Replacement servers do not need to be identical hardware, and VM's can be easily consolidated during recovery.
- Easy **migration and upgrade** to more powerful server.
- Easy **replacement or consolidation** of older servers.

# Benefits of Suspend Mode

- Easy capturing application data while the system is having the problem for more effective root cause analysis.
- May facilitate certain types of system backups.
- The machine's memory state and other checkpoint information is placed in a disk-based **suspend file**.

# Benefits of Point in Time (PiT) Restoration

A VM's Virtual Disk file can be placed in “**Undoable mode**”. This creates a “**Redo File**” that captures changes. You can choose to apply or discard changes that were made after the disk was placed in this mode. This can be used for:

- Patch management - Ability to test changes, and easily commit or reverse the changes.
- Software test and application development.
- Ability to restore a VM to a prior point in time.


# Why not put all Applications on a Single Operating System Instance?

- They may conflict with each other.
- You can't test all possible interactions.
- You can't even install many apps together.
- You need to dynamically control resources.
- Some apps misbehave – i.e., leak memory or crash.

In other words, most major applications assume they are on a dedicated instance.

# What is ESX Server?

- A platform for VM's:
  - VMs can have minimum guaranteed resource allocations.
  - Allocations can be changed without rebooting.
  - SAN Path Management and protection built-in.
  - Supports Hyperthreading.
  - Supports Virtual SMP.
  - More details later...



# Attributes of VMware ESX Server

- **Dynamic Resource Allocation** - Resources are shared between all running VM's, and can be dynamically reallocated.
- **Isolation/Partitioning** – The state of one VM does not affect the state of another. If one VM fails, the others keep on functioning, unaffected.

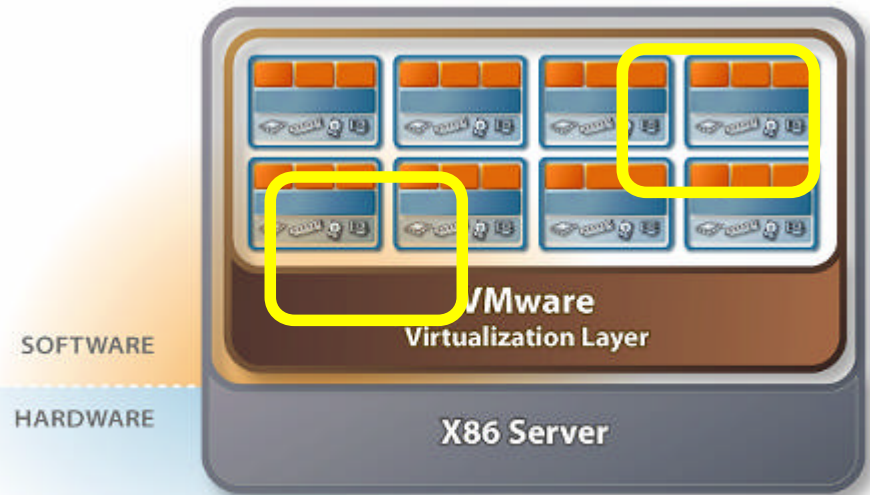
# Benefits of Dynamic Resource Allocation

- All of the VM's can have access to SAN and LAN resources, through the ESX Server's FC-HBA's and NIC's. This usage can be allocated and controlled.
- If a VM goes into a "resource runaway" condition, it will be limited by its maximum allocation, to minimize or prevent any negative effect on other VM's.
- Allows for adjustment of CPU, Memory and Disk resources to balance workload, and for changing of these allocations based on anticipated or measured workload changes.



# Partitioning

- Run multiple Operating Systems, or multiple instances of the same OS on one ESX Server.
- Fully utilize server resources, as all VM's can be allowed to access common resources.
- Each VM fully protected and resource access share can be guaranteed.



# Isolation

- CPU hardware / protection.
- Fault, performance, and security isolation.
- CPU, RAM, Disk, and network resource controls.
- Resource allocations can be changed “on the fly”.
- Guaranteed service levels.
- If one VM “crashes”, it has no negative effect on any other running VM’s.





# Capacity Planning & Workload Management Changing the Rules...




# Four Basic “Food Groups” of a Server..


- **Storage**




- **Memory**



- **Local Area Network**



- **Aggregate CPU Capacity**



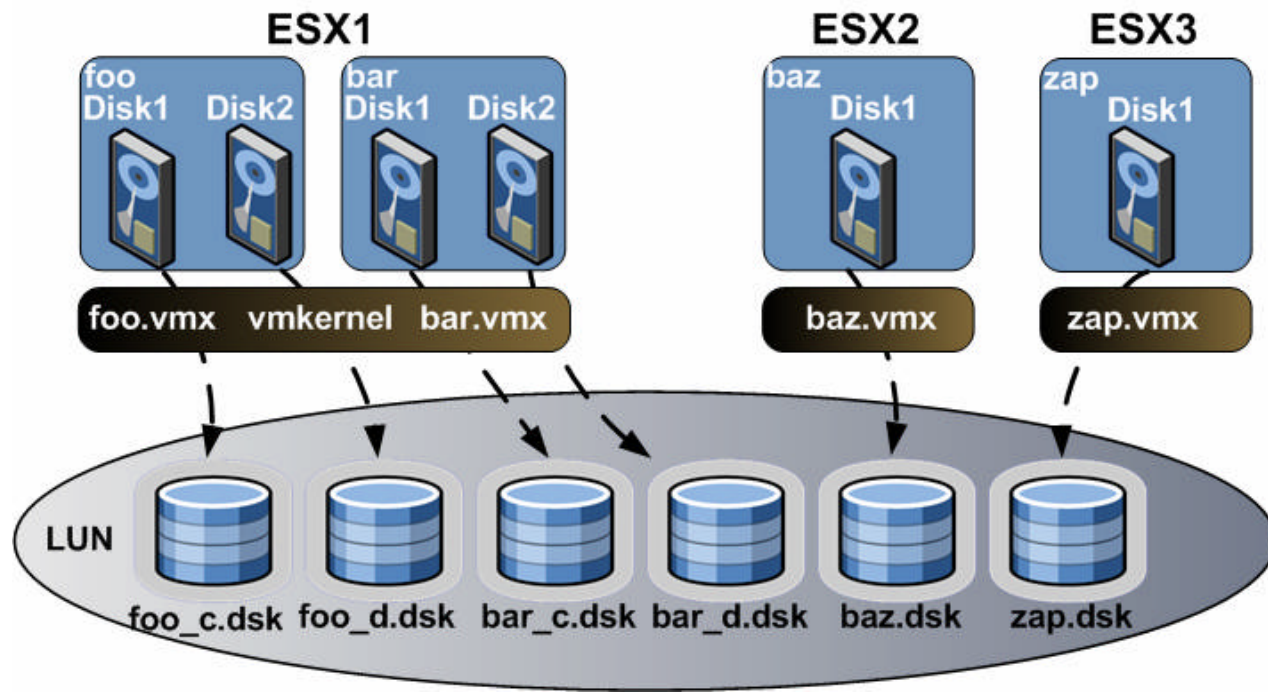
# Disk Storage

- Every physical server on a SAN requires one or two Fibre Channel HBA's, depending on the fault tolerance requirements, as well as cables and switch ports.
- As a VM sees SCSI LUN's, and doesn't know whether the LUN is on local or SAN-based storage, a pair of HBA's on an ESX Server could allow all VM's to access SAN resources. This fact is often missed when calculating server consolidation cost savings.

# SAN Cost Savings and Simplification

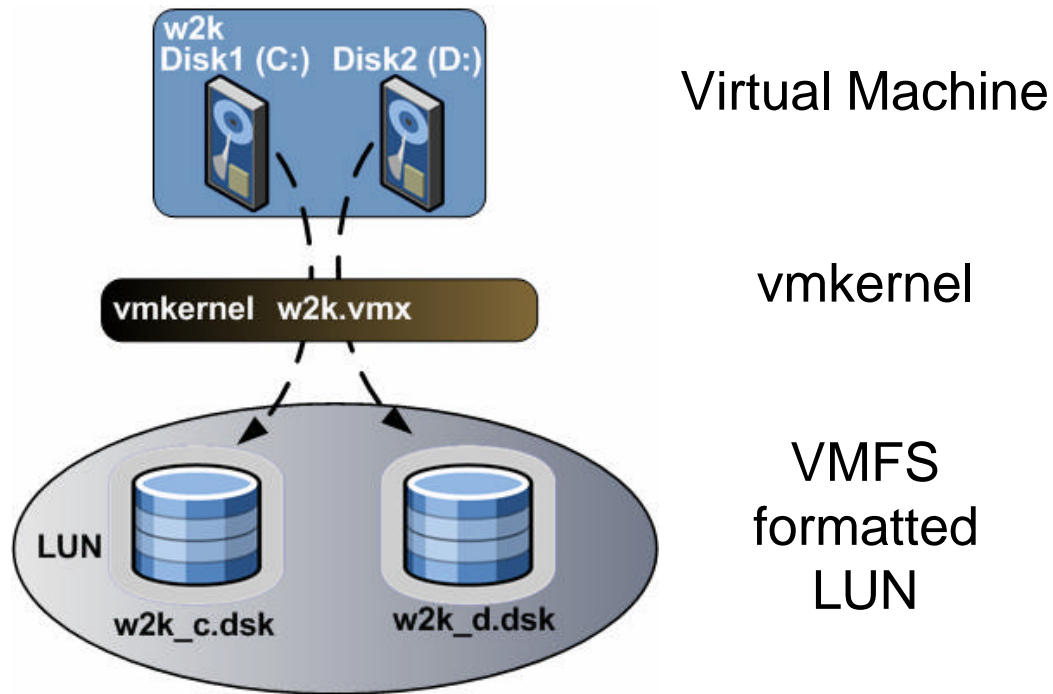
- 10 physical servers will require 20 Fibre Channel HBA's (for redundancy), 20 switch ports, cabling and path management software.
- One ESX Server will require 2 Fibre Channel HBA's (for redundancy), 2 switch ports, cabling and NO path management software, as it is built-in.
- There is also a cost associated with administering the more complex SAN and each component through its economic life.

# ESX SAN Scaling



- Multiple ESX Servers may share a single LUN
- VMFS handles concurrent access
- vmkernel manages Fibre Channel Multipathing and Failover. No additional path management software is needed.

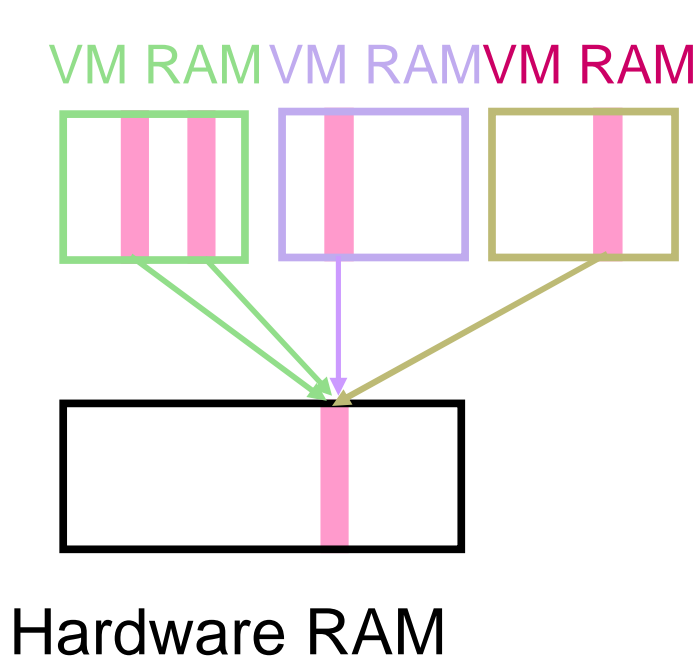
# ESX Virtual Disk to File Mapping



- Each Virtual Disk in a VM maps to a file in a VMFS or a raw LUN.



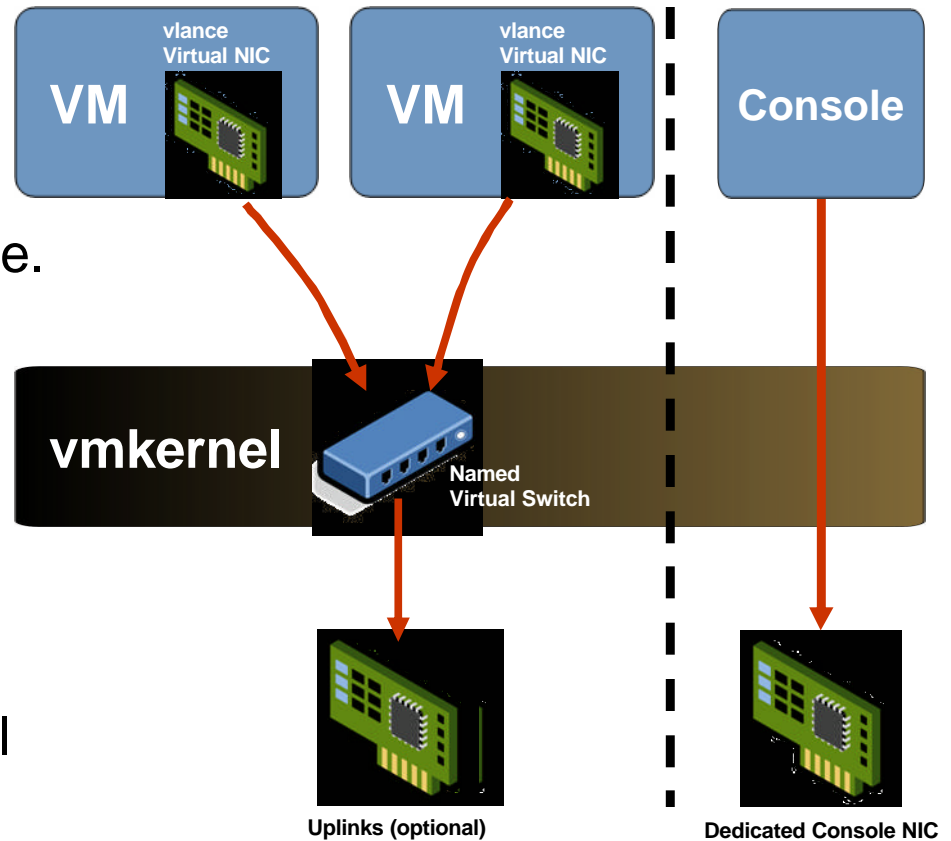
# Transparent Page Sharing Technology



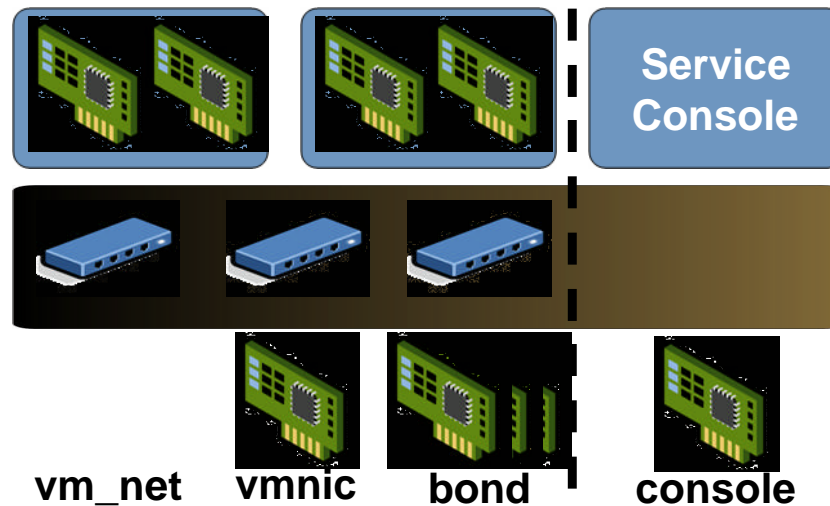
- Shares identical pages
- Periodic scans and fast match algorithm.
- Example: ~ 30% to 60% shared w/ idle Win2K VMs, due to code redundancies and “zero pages”.
- Single instance created when no longer redundant.
- No OS changes needed.

# Network Virtualization

- VM's have Virtual NICs.
- Dedicate NICs to Service Console.
- Virtual NICs plug into Virtual Switches.
- Virtual Switches are named.
- May have zero or more uplinks:
  - Zero – “host only” network.
  - One – normal external network.
  - Two or more – “bonded” external links for fault tolerance and bandwidth aggregation-802.3ad.
- Switches may have port groups.



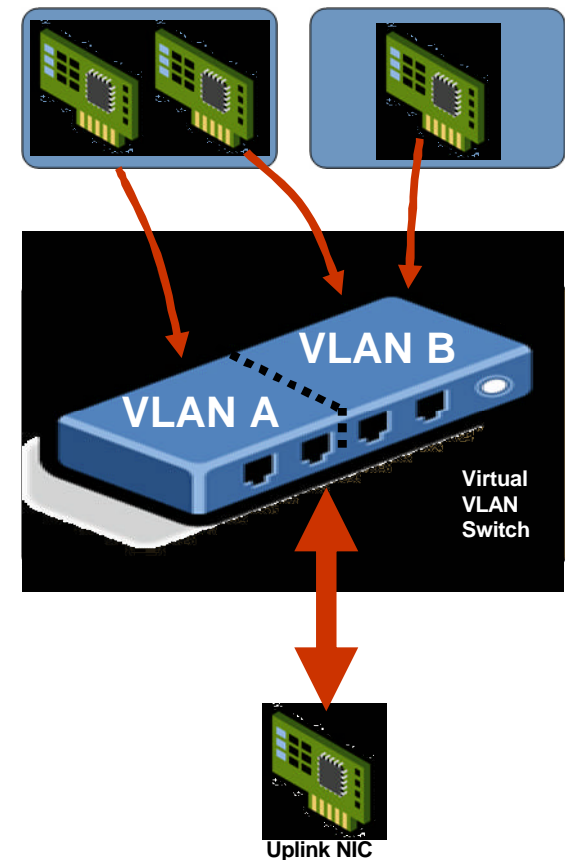
# Network Access



- Each Virtual NIC has own MAC address.
- Virtual NICs connect to Virtual Switches.
- Virtual Switches may have 0, 1, or more uplinks.
- Virtual Switches support PXE for deployment technologies such as RDM, Altiris and Symantec Ghost.
- Virtual Switch “Port Groups” may be VLANs (802.1q).

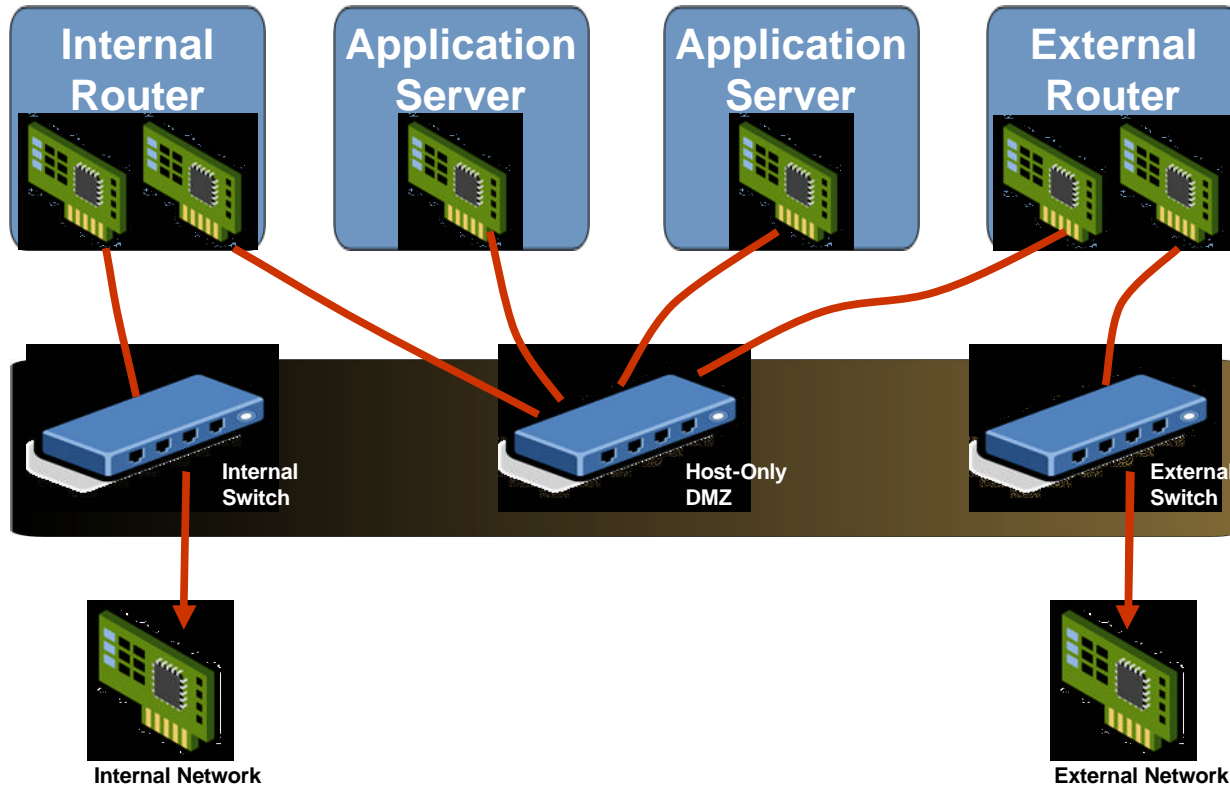
# Port Groups

- Port Groups are optional Virtual Switch attributes.
- Each Port Group has a VLAN ID.
- Uplinks plug into physical switch trunk ports.



To physical switch  
VLAN trunk port

# Firewall In A Box

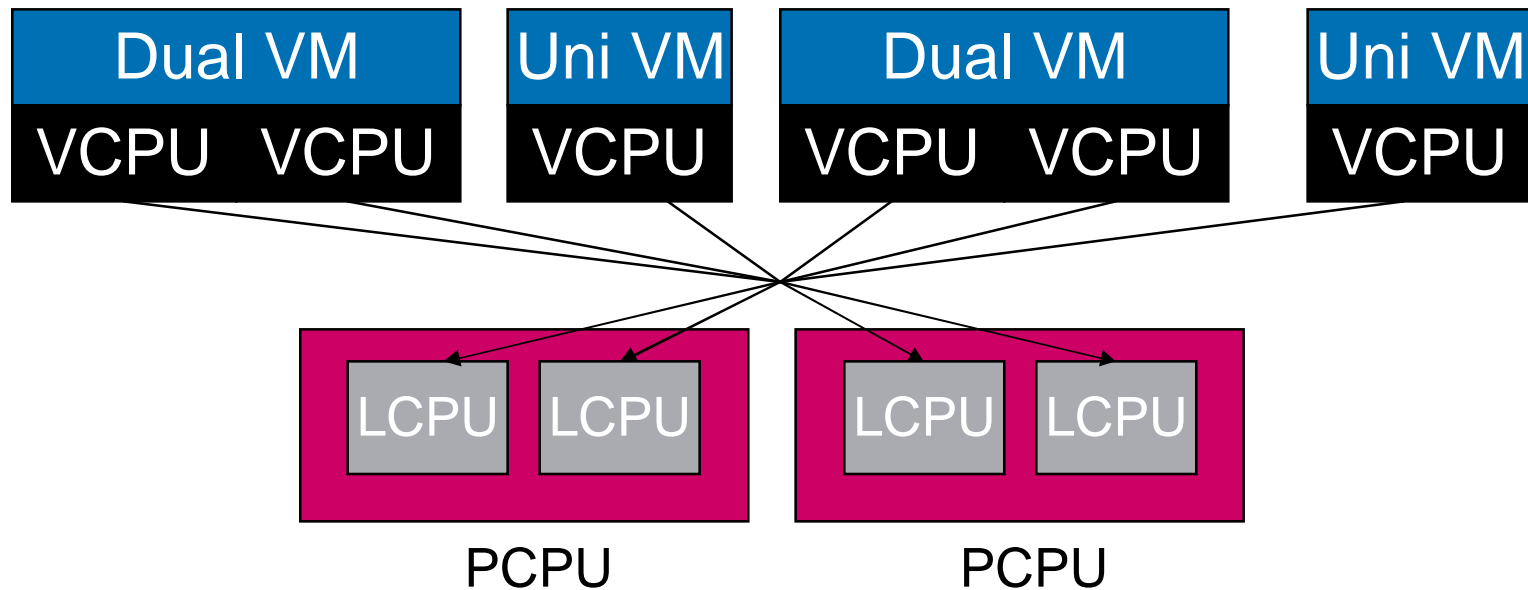


***Could be done with 1 NIC + VLANs***

# What is CPU Hyperthreading?

Hyperthreading (HT) allows one CPU to execute two threads simultaneously:

- Each Physical CPU (PCPU or package) offers two Logical CPU's (LCPU's).
- VMkernel schedules Virtual CPU's (VCPU) onto Logical CPU's.



# Hyperthreading and ESX Server

- Hyperthreading is used by default if present
  - Make sure HT is enabled in BIOS.
- VMkernel will schedule VMs intelligently
  - CPU-intensive Virtual CPU's will run in different Physical CPU's.
- VMs can be bound to Logical CPUs
  - On HT, CPUs 0 and 1 are really the two Logical CPUs in Physical CPU 0.
  - Caution! Don't bind, (by CPU affinity setting) busy VMs into the same Physical CPU.
- A VM can be isolated from Hyperthreading
  - Useful for CPU-intensive or cache-intensive workloads.

# Workload Management

- Traditionally, you needed to guess what server resources an application required at peak periods, and provision a server based on these peak requirements, even if the peak requirements were a relatively small fraction of the server's total uptime.
- Because the portability of a VM, you can migrate a VM from one ESX Server to another, either offline, which is referred to as a “**Cold Migration**”, or “on the fly”, which is referred to as a “**Hot Migration**”.




# Enhancements to Virtual Machine Portability - VMotion

- A technology called **VMotion** allows the transfer of a running Virtual VM from one ESX Server to another with imperceptible downtime.
- Workloads can be transferred to an ESX Server with more resources during peak periods.
- Workloads can be offloaded to other ESX servers during hardware maintenance for higher system availability.



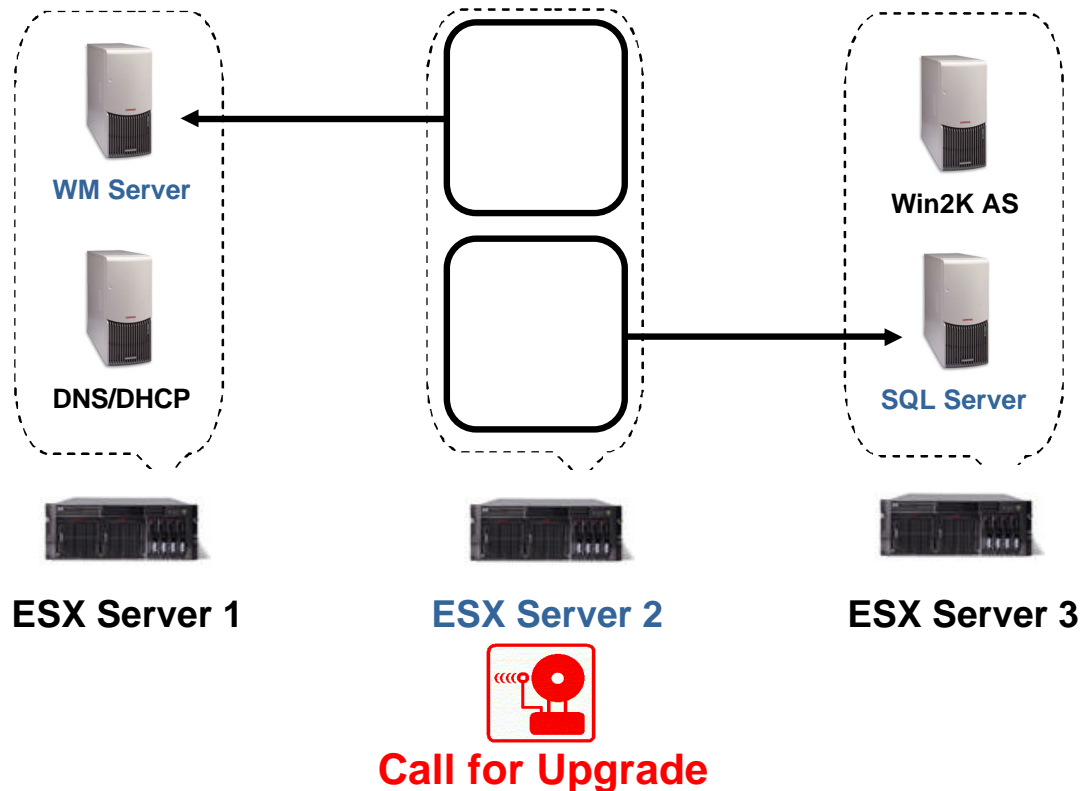
# Zero Downtime Maintenance

Upgrade and service production hardware using VMotion with zero downtime and 100% customer transparency.




# Zero Downtime Maintenance

Upgrade and service production hardware using VMotion with zero downtime and 100% customer transparency.




# Zero Downtime Maintenance

Upgrade and service production hardware using VMotion with zero downtime and 100% customer transparency.



# Zero Downtime Maintenance

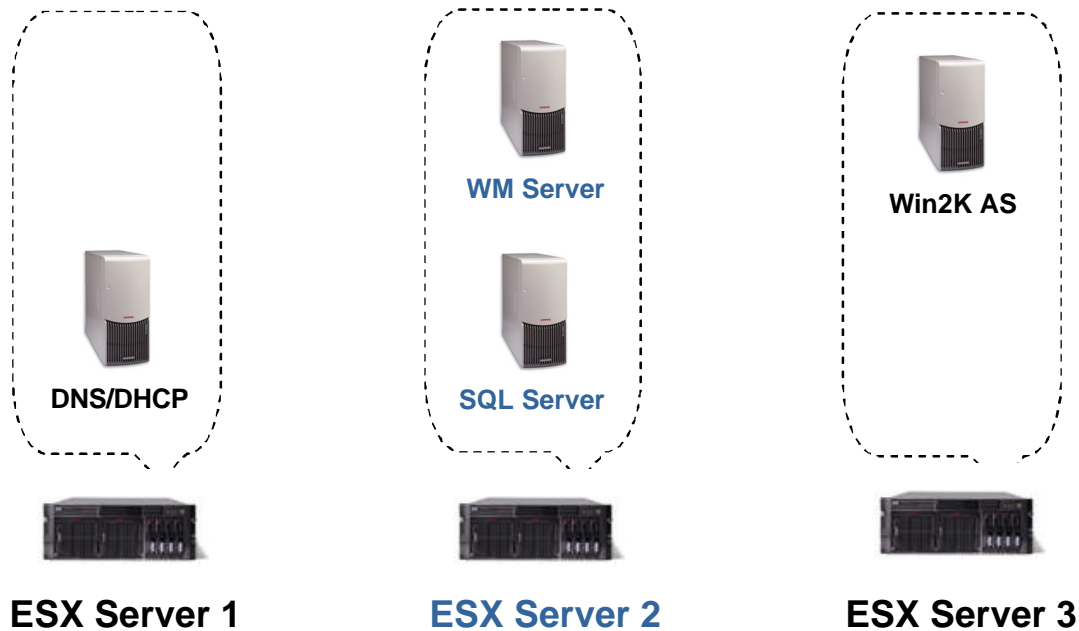
Upgrade and service production hardware using VMotion with zero downtime and 100% customer transparency.



**Upgrade Finished  
Powered On Again**

# Zero Downtime Maintenance

Upgrade and service production hardware using VMotion with zero downtime and 100% customer transparency.




# Workload Management

- A hot migration could take VMs off of one ESX server to give the remaining VMs greater access to resources.
- A hot migration could also transfer a VM to a lesser-used or unused ESX Server only during its heavy traffic periods.
- Based on performance measurement data, ESX Servers can be “**load balanced**”, and this balance can easily be changed, either based on “trend data”, or dynamic changes.



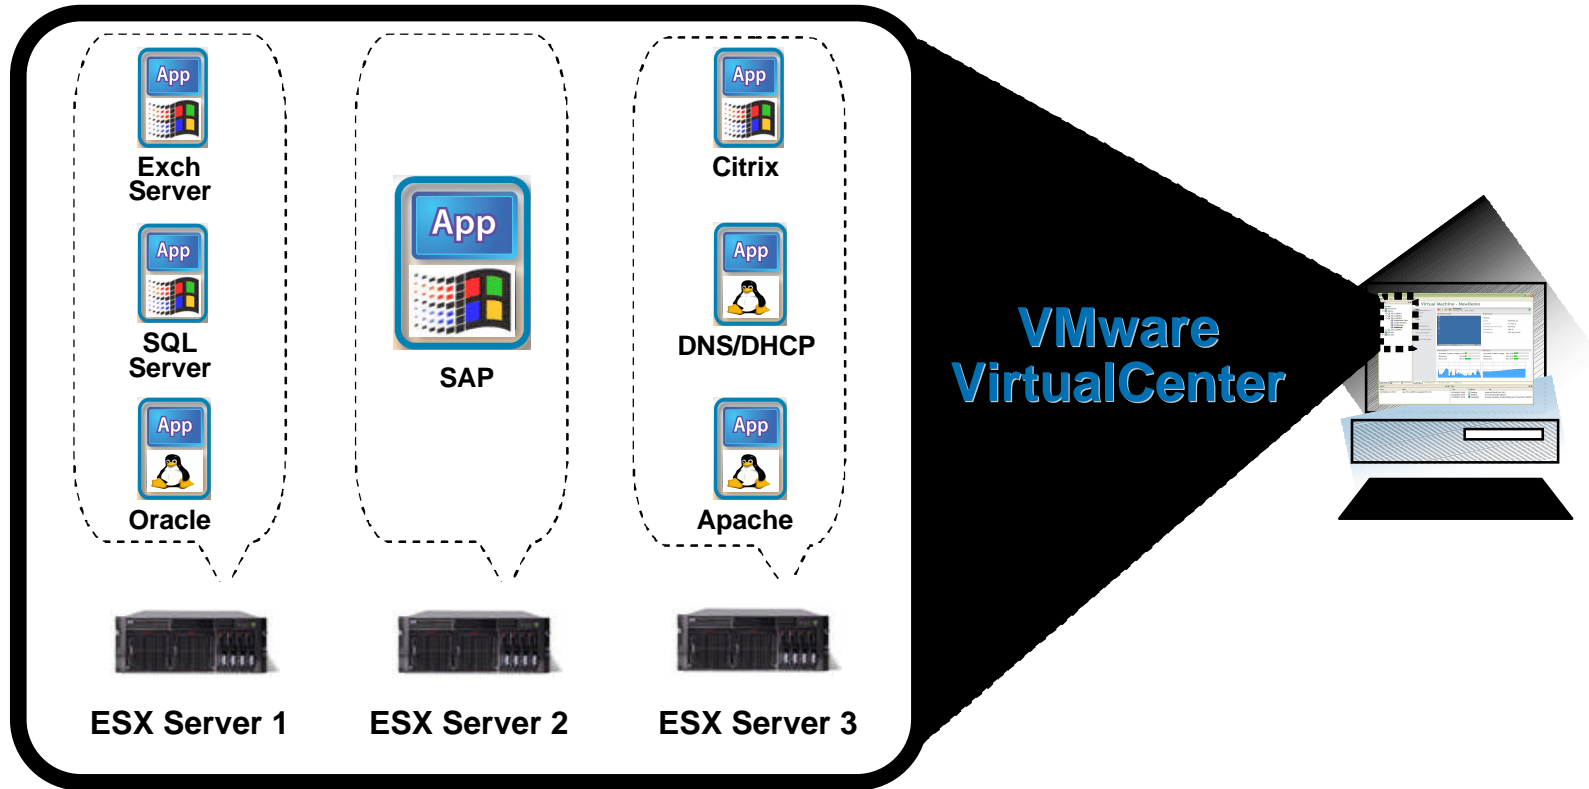
# Workload Management with VMotion

Dynamically manage workloads in response to an unexpected increase in SAP utilization




# Workload Management with VMotion

Dynamically manage workloads in response to an unexpected increase in SAP utilization



# Alternate Form Factors

Enable a Virtual Scale-Up solution from multiple server blades




- Achieve 60-80% average utilization.
- Share headroom across multiple server blades to handle spikes.
- Applications are now mobile across server blades.
- Transparently shift systems across blades to match complementary workloads.


# Increased System Utilization

Run Four Chassis Worth of Workloads on a Single Chassis

Blade Servers without VMware VirtualCenter



Blade Servers with VMware VirtualCenter



# Workload Management

- The allocation of CPU, Memory and Disk resources can also be changed dynamically, through guaranteed minimums, maximums, and a “**Proportional Share**” mechanism. Network traffic can be throttled through **Traffic Shaping**.
- This means that predictable service levels can be guaranteed, and resource allocations can be adjusted as needed, or even changed on a scheduled basis.

# Capacity Planning –Other factors

- Anticipated Workload Increases.
- Planning for Peak vs. Average loads.
- Time of day, day of week, day of month, day of year.
- Hardware Failover / Clustering.
- Inter-relationship between one VM's workload and another.
- Remember that VMotion, and dynamic resource allocation changes can help manage this.

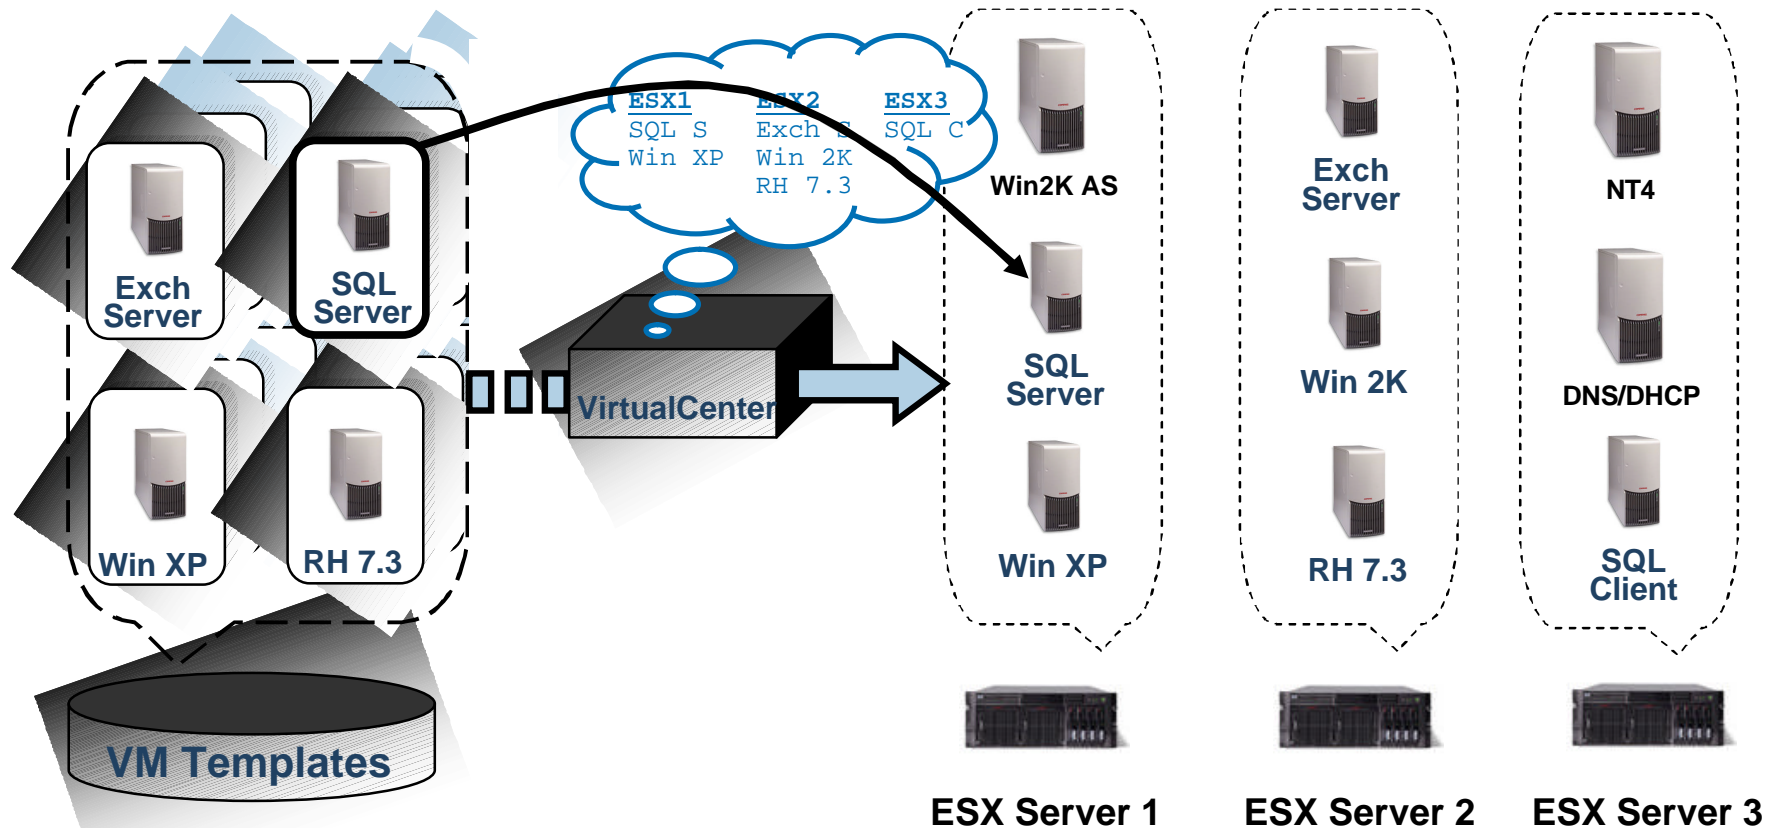


# Server Provisioning / Disaster Recovery / High Availability Changing the Rules...



# VMware Automated Server Provisioning

Automate the deployment of new servers into a server farm using a single repository of VM templates





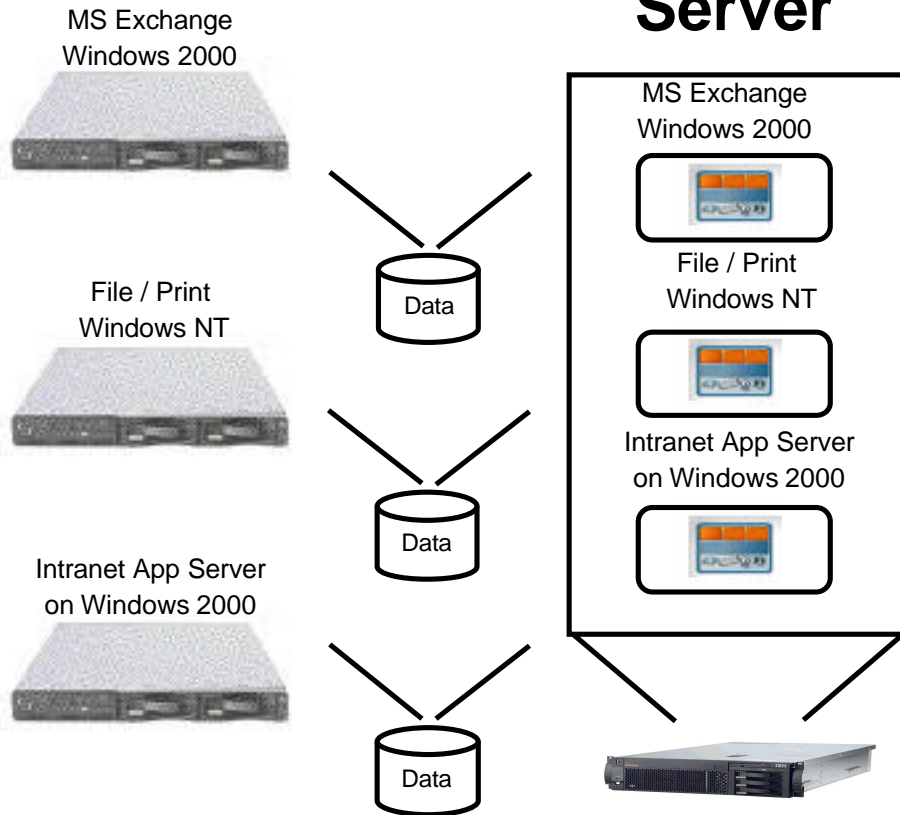
# Disaster Recovery and High Availability

- VMs can be run on different hardware.
- Several VMs can be consolidated on a single physical server.
- Virtual Disks are stored on VMFS, allowing for multiple ESX Servers to access the same volume.
- VMFS can be replicated, and standby site can be brought up quickly.

# Failover Clusters – 3 Configurations

- **Two nodes within the same ESX Server** Protects against operator error, application and OS crashes.
- **Two nodes on two different ESX Servers** Protects against operator error, application and OS crashes, hardware failures.
- **One active node on a physical server and its backup on an ESX Server** - Protects against operator error, application and OS crashes, hardware failures. Low-cost N+1 redundancy.

# ESX Server based Failover Solution



## Benefits:


- Low cost: organizations for whom 1-1 failover is unaffordable.
- Flexible: allows failover of heterogeneous-OS servers.
- Scalable: additional failover machines can be added as needed.

# Physical to Virtual Server Consolidation

## Mail System Example

Traditional Approach

Apache on  
Linux



Sendmail  
Gateway  
on Linux




Exchange + IIS  
on  
Windows 2000



## Mail System Example


Consolidated Environment



## Customer Benefits


- No need to standardize on one OS.
- Removes concerns about application interaction or dependencies.
- VMware resource governing guarantees that critical apps will get the resources they need.
- Reduced system setup and management costs.
- Manage “Legacy” OS and Apps with current physical hardware.

# Backup Options with VM's



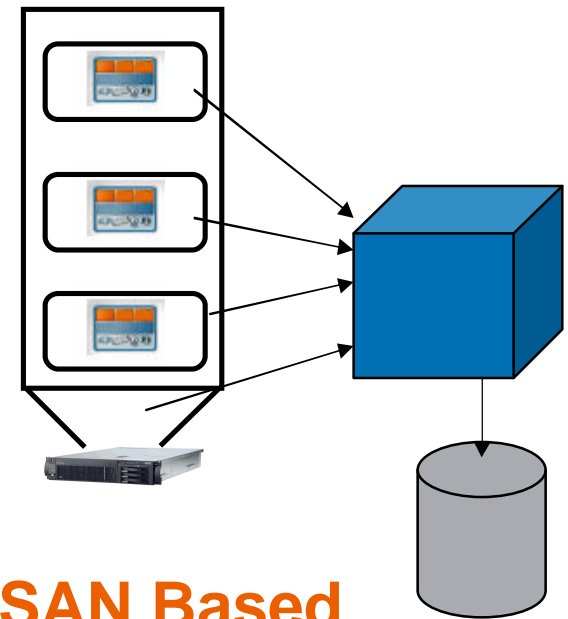
## Backup Each VM

- Agent based inside each VM.
- Like physical servers.



## Whole Server Backup

- Use backup software with ESX Server Console OS.



## SAN Based

- Use SAN snapshot.
- SAN handles checkpointing and tape archiving automatically.

# Disaster Recovery Restoral Simplification

- In any of the backup scenarios, VMware makes recovery simpler.
- List of items required to bring to the DR site is smaller.
- ESX Server licenses and install.
- .DSK files (VM backups)
- Way to restore VM configurations.

VS

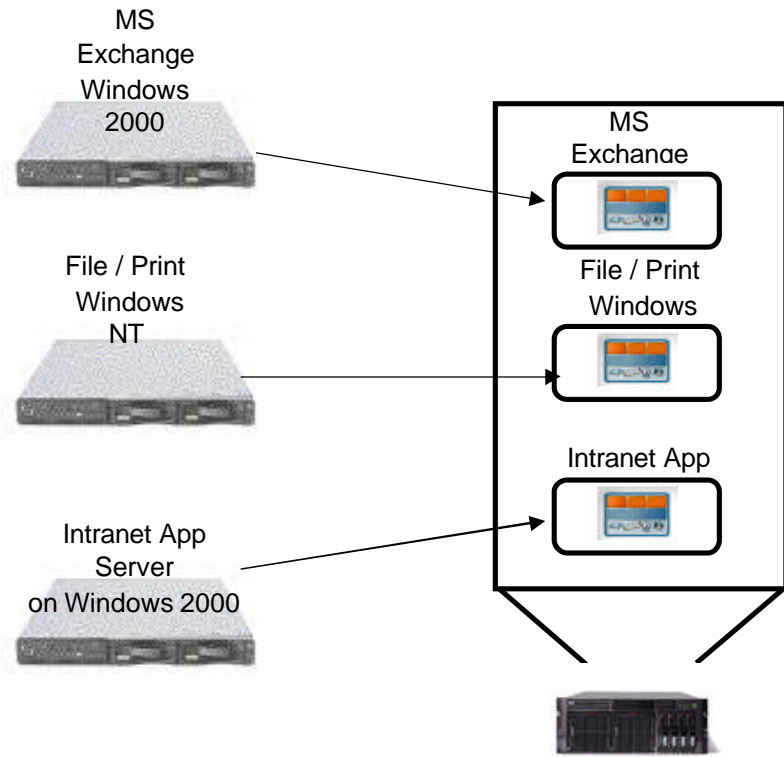
- Bringing full set of tapes and way to cold boot.
- Rebuild server on slot-for-slot identical hardware models.

# Replication for Disaster Recovery

Replication technology such as IBM LVM, HP DRM, EMC SRDF

Continuous Replication for applications such as:

- Exchange / Notes
- SQL Server / Oracle
- IIS
- Custom Application



1-to-1 or Many-to-1 Continuous Replication

# Some Tasks that Change...

<b>Key Task</b>	<b>Traditional Approach</b>	<b>With Server Virtualization</b>
<b>Provision a new server</b>	<ul style="list-style-type: none"><li>• 3 - 10 days hardware procurement</li><li>• 1 - 4 hours provisioning new server</li></ul>	<ul style="list-style-type: none"><li>• 5 - 10 minutes provisioning new VM</li></ul>
<b>Moving an application to a new server; Repurposing a server</b>	<ul style="list-style-type: none"><li>• 4 - 6 hours for migration</li><li>• Service interrupted for duration of maintenance window</li><li>• Requires days/weeks of change management preparation</li></ul>	<ul style="list-style-type: none"><li>• 2 - 5 minutes using VMotion (no service interruption)</li><li>• 10 - 30 minutes without VMotion</li></ul>
<b>Hardware maintenance</b>	<ul style="list-style-type: none"><li>• Requires 1 - 3 hour maintenance window</li><li>• Requires days/weeks of change management preparation</li></ul>	<ul style="list-style-type: none"><li>• Zero downtime hardware upgrades with VMotion</li></ul>





Virtualization  
Technologies that are  
available now...



# Technologies that are Available Now...

- VMware Workstation
- GSX Server
- ESX Server
- VirtualCenter
- VMware P2V Assistant

# VMware Virtualization Technologies...

## Workstation

- Productivity Tool
- Performance 70-90% of Native

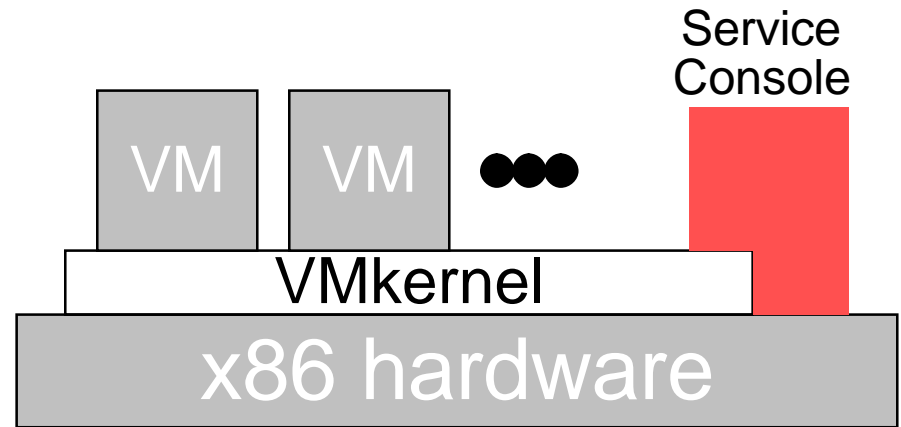
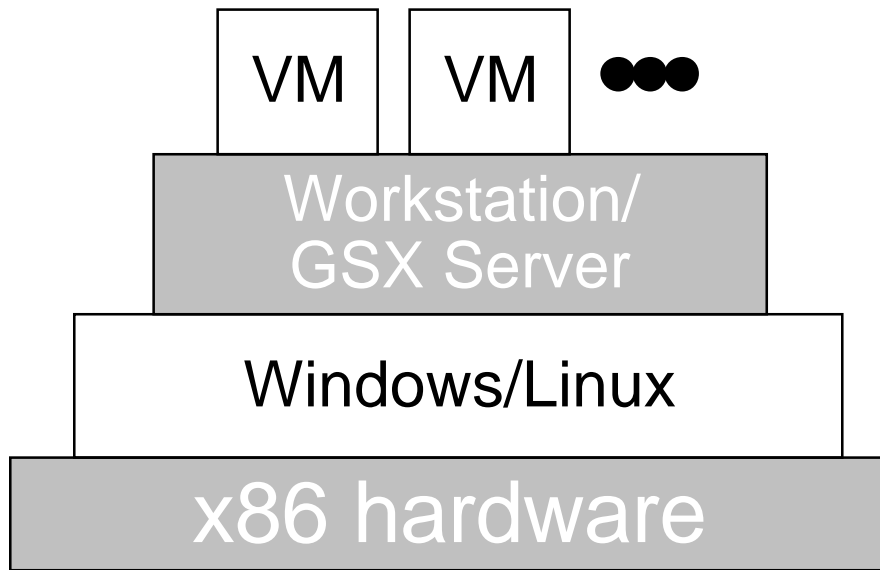
## GSX Server

- For Departmental Services
- Remote Management
- Performance 70-90% of Native

## ESX Server

- For Data Center
- High Performance and Scalability
- Advanced Resource Management
- Remote Management
- Performance 83-98% of Native

# Virtualization Architectures



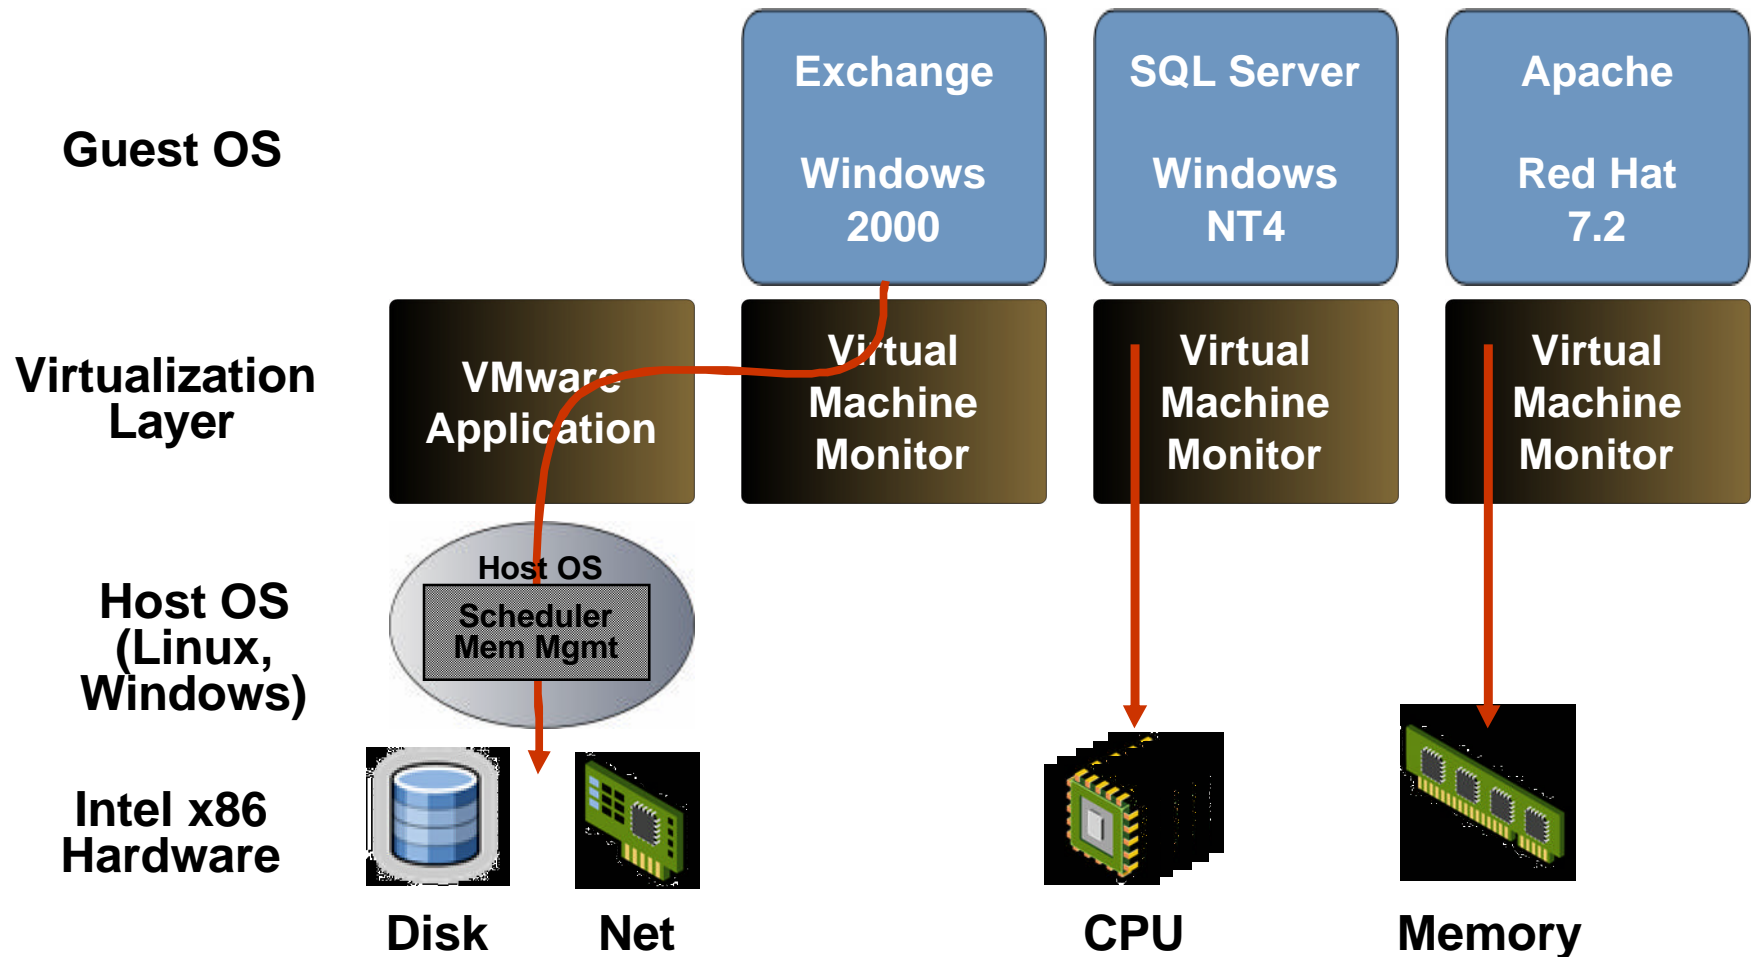
## Hosted (Workstation, GSX Server)

- Maximum device compatibility
- Installs like an application
- Lower price point

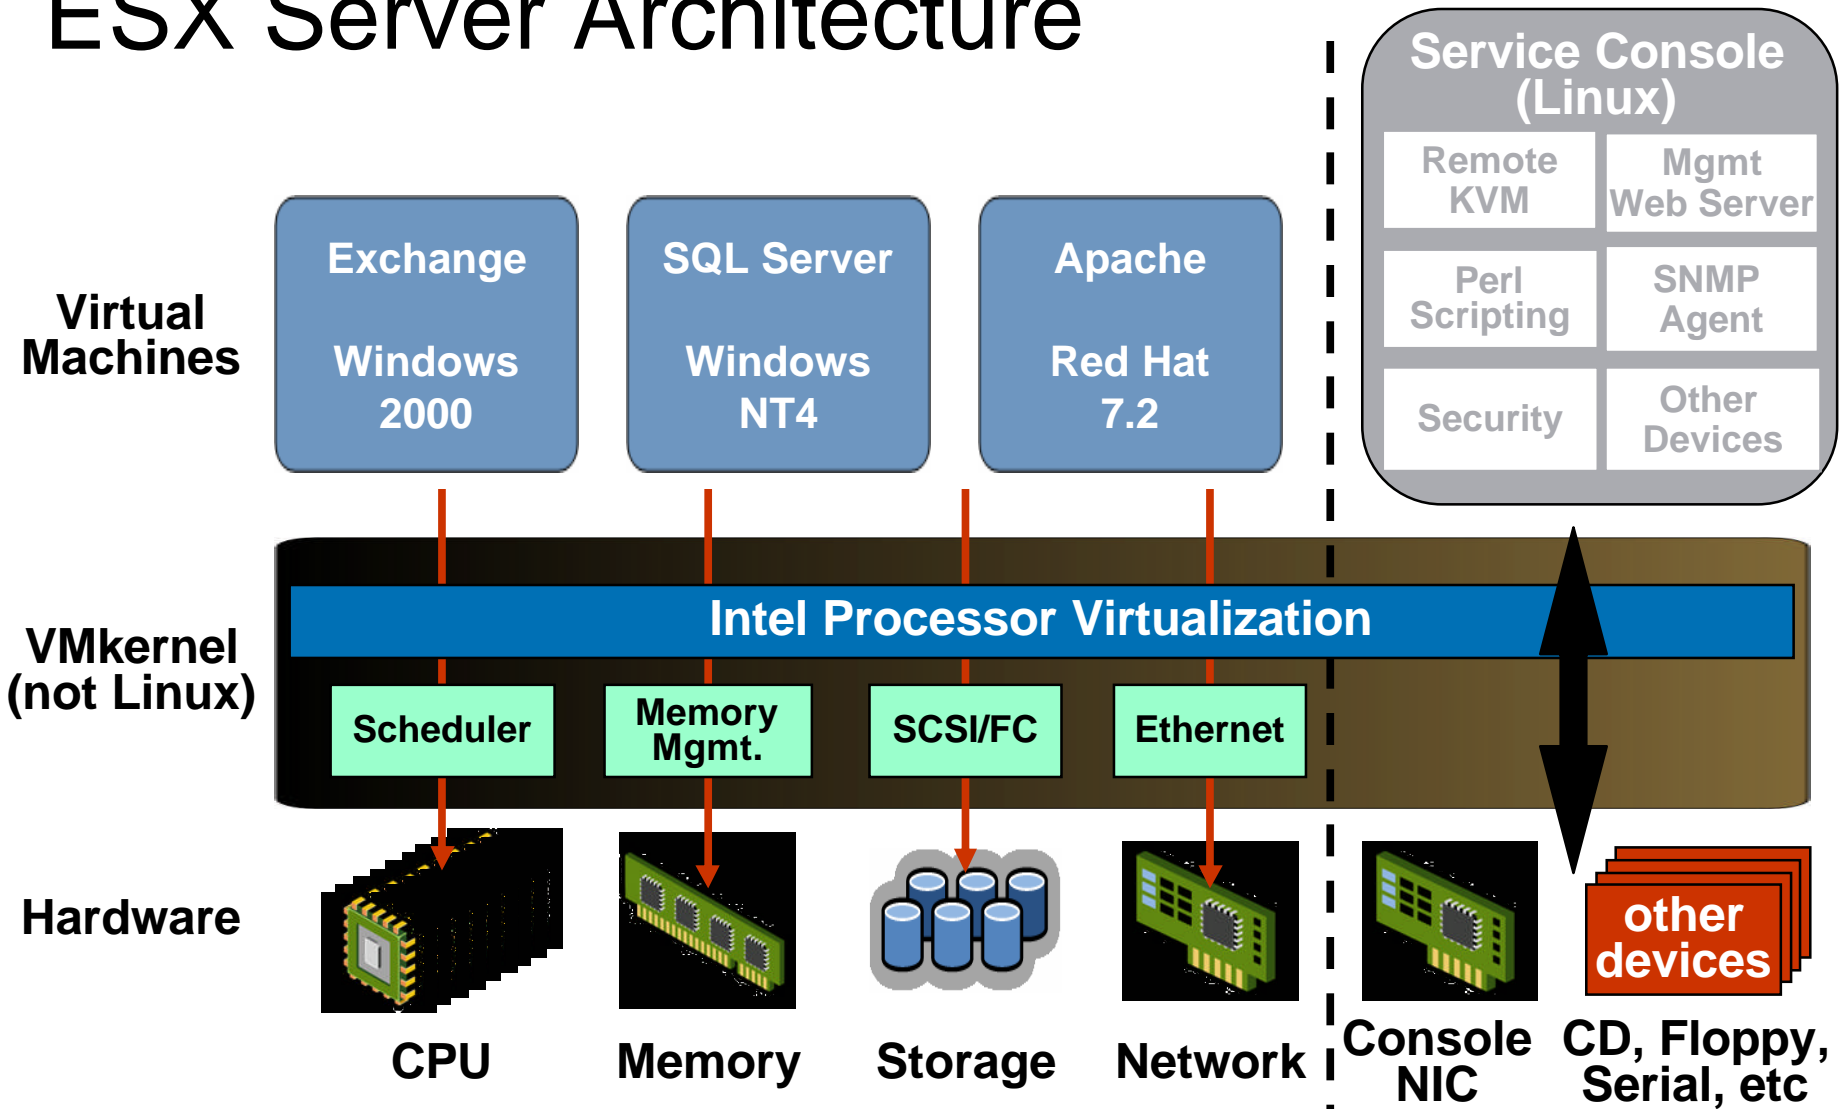
## “Bare-Metal” (ESX Server)

- Maximum performance
- Lowest overhead
- Dynamic resource controls
- Virtual SMP

# Workstation and GSX Server Architecture



# ESX Server Architecture



# Comparing GSX and ESX Server...

## VMware GSX Server



- Hosted on Windows or Linux
- Maximum hardware compatibility
- Up to 3.6GB RAM per Guest OS
- Up to 64 VMs per GSX Server
- Lower price point

## VMware ESX Server



- Runs natively (“Bare-Metal”)
- Maximum scalability and performance
- Up to 3.6GB RAM per Guest OS
- Up to 80 VMs per ESX Server
- Virtual SMP support
- Advanced memory management
- Dynamic resource allocation

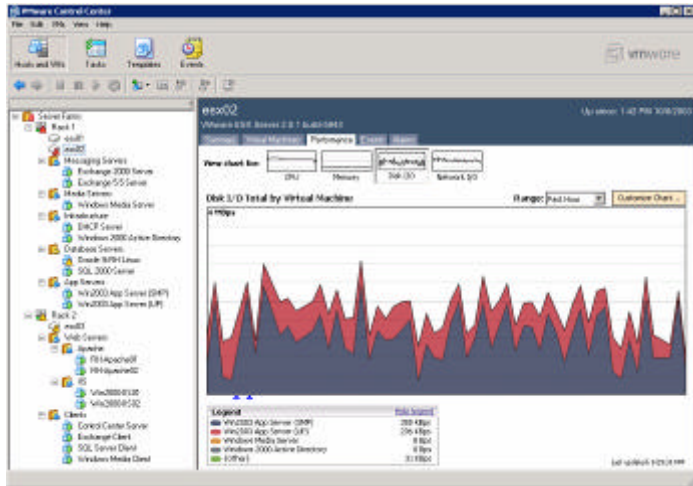
# VirtualCenter

- **VirtualCenter**: a management front end to ESX Servers and GSX Servers, and their VMs
  - A Windows application with a database back-end
  - Uses:
    - Centrally manage all VM images
    - Deploy new VMs using standardized templates
    - Monitor VMs' state and performance
    - Manage access to VMs with one point of control
    - Coordinates VMotion transfers
    - Provides notification of system events
    - Allows for scheduling of tasks
    - API's for 3<sup>rd</sup> party software developers



# VirtualCenter Overview


An enterprise solution for deploying and managing virtual machines across the data center.



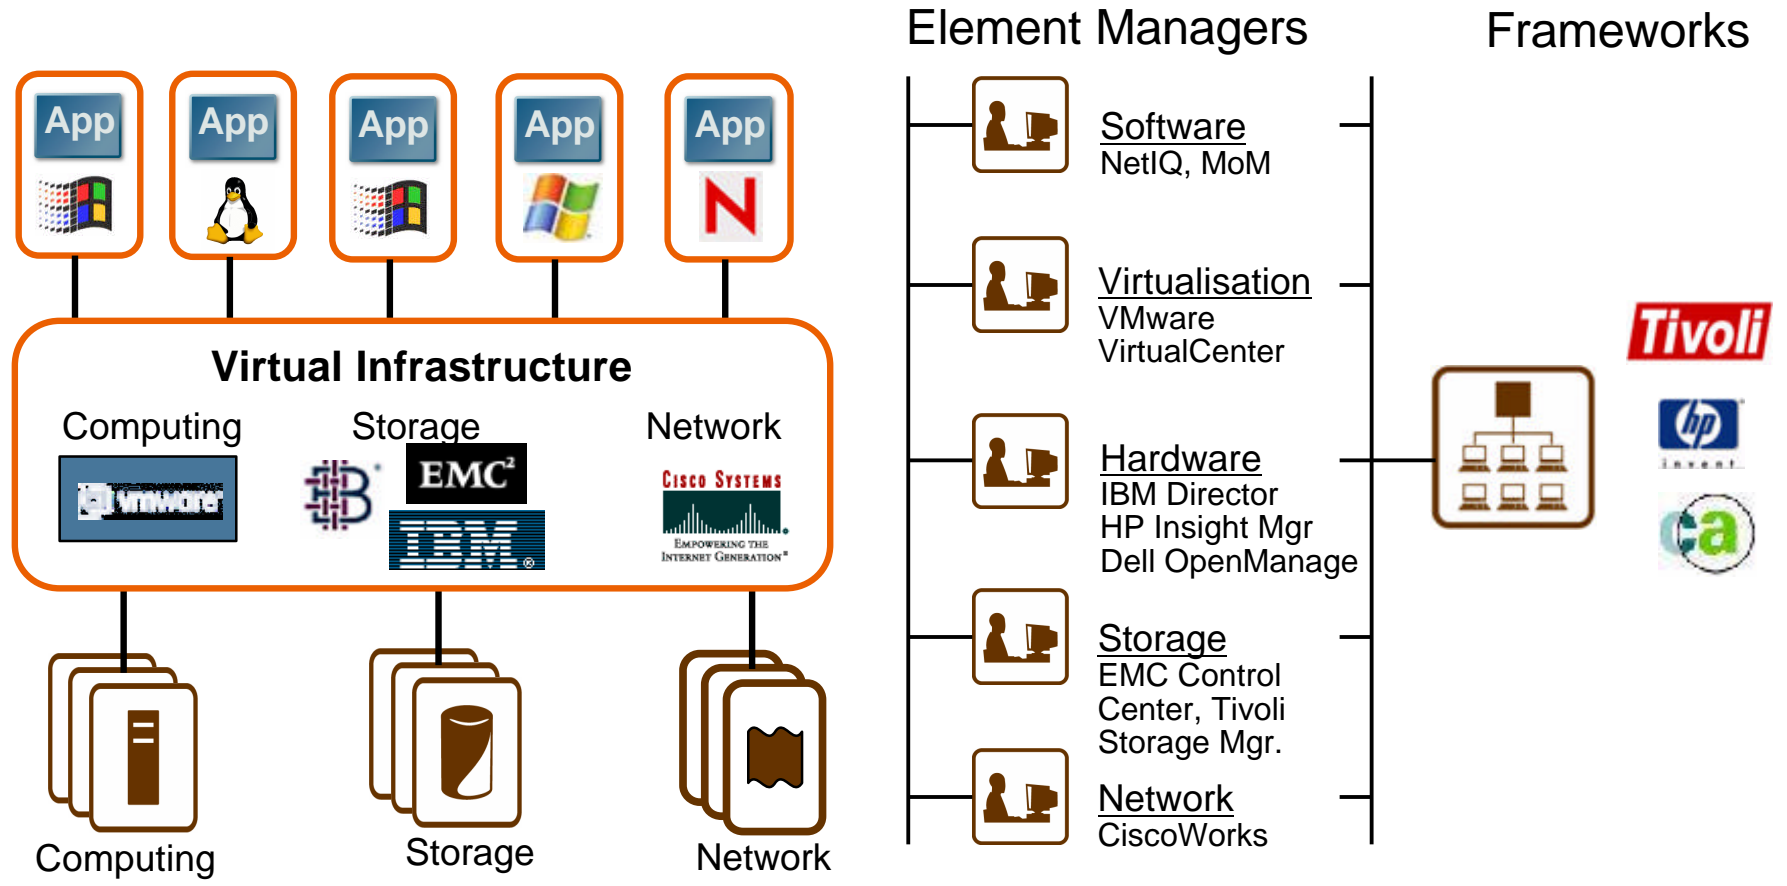
- **Manage** Intel hardware as a single pool of computing resources.
- **Streamline** server provisioning and management.
- **Monitor** system availability and performance.
- **Move** workloads dynamically across distributed servers without service interruption.
- **Eliminate** scheduled downtime.
- **Automated** notifications and email alerting
- SDK to **integrate** with existing management tools

# Centralized Management Console

Centrally manage a heterogeneous computing environment from a single graphical user interface.

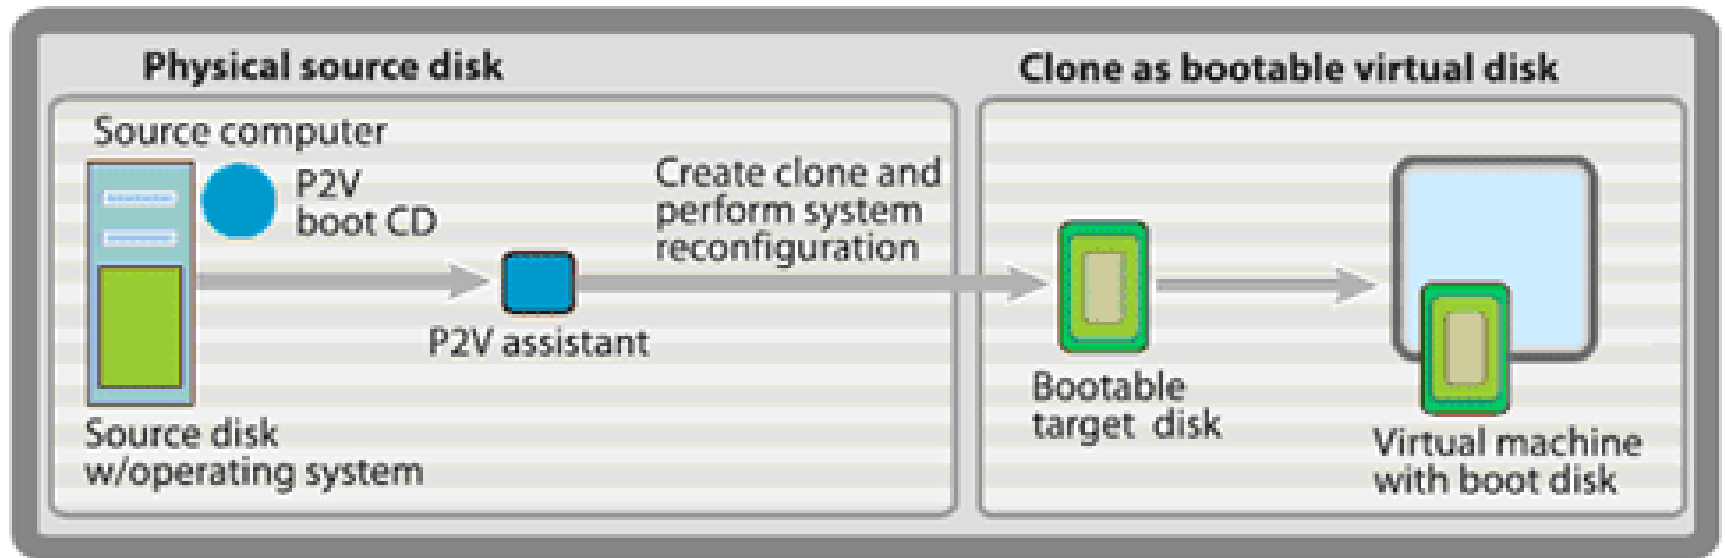


# Virtualization Management Architecture



# VMware P2V Assistant

- Takes a snapshot of an existing physical system and transforms the physical server image into a VMware virtual machine. Simplifies conversion to a Server Virtualization environment



# Resources...

- **Webinars**

<http://vmwareevents.webex.com>

- **Training**

<http://www.vmware.com/services/education.html>

- **Consulting Services**

<http://www.vmware.com/services/consulting.html>

- **Products**

<http://www.vmware.com/products>



Questions ?





# Backup Slides... Server Virtualization- Changing the Rules...



# Hosted vs. Bare-Metal Virtualization

## Hosted

- Granularity of resource controls limited by design of underlying operating system.
- Security of any virtual machine is subject all all vulnerabilities of the underlying host operating system.
- Host operating systems are typically 10-40 million lines of complex error prone code.
- I/O of virtual machines must go through two protocol stacks that are not aware of each other.

## “Bare-Metal”

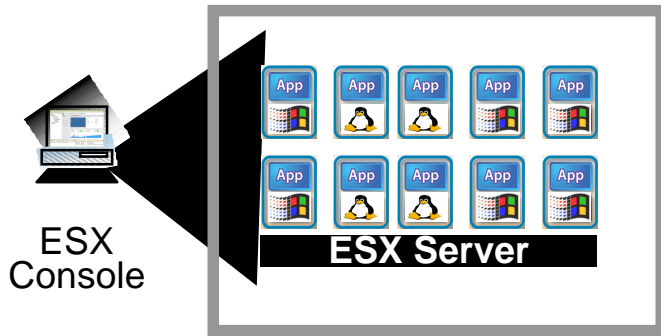
- Enables fine-grained resource management of all major system resources.
- Micro-kernel virtualization layer is highly secure. There are no O/S services vulnerable to abuse.
- Virtualization layer of only 100K's of lines of code which enhances availability.
- Highly efficient direct I/O pass-through mimics mainframe architecture for both Network and Disk.




# VMware on SMP and Blade Server Platforms

*Same key benefits, differences in scale and flexibility*

## Large SMP


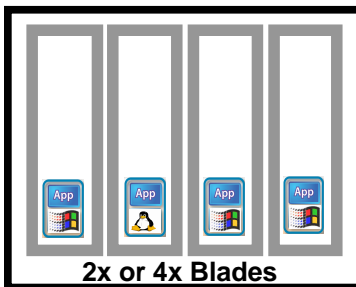


## Blade Server



ESX Server on 4-16x SMP System	VirtualCenter w/VMotion on Blade Server
<p>Large single system image</p> <p>Automatic load-balancing of large number of VMs in same system</p> <ul style="list-style-type: none"> <li>• Simplified resource sharing and management</li> </ul> <p>Lower entry cost for smaller systems</p> <ul style="list-style-type: none"> <li>• SCSI disk can be used</li> </ul>	<p>Greater resiliency and availability</p> <p>Multiple 2X/4X blade and ESX instances</p> <p>“Standby” blade within chassis for availability</p> <p>Modular scalability</p> <p>Increased scaling of memory and I/O capacity</p> <p>Blade form factor advantages</p> <p>Increased density, standardization, power efficiency</p> <p>Workload balancing with mixed blades</p> <p>e.g. temporarily migrate VM to higher performance blade to complete task more quickly</p>

# VMware Multiplies Blade Server Benefits



BENEFIT	BLADES w/o VMware	BLADES & VIRTUALCENTER
SERVER DENSITY (# OF SERVERS IN 42U)	<b>84 (2X)</b>	<b>168 - 508 (4-12X)</b>
UTILIZATION	<b>10-60%</b>	<b>50-80%</b>
SERVICEABILITY (BLADE MAINTENANCE DOWNTIME)	<b>1hr end-user downtime</b>	<b>ZERO end-user downtime</b>
AVAILABILITY	<b>N+N failover</b>	<b>N+1 failover</b>
PROVISIONING TIME	<b>&lt;1 Hour</b>	<b>Minutes</b>
CROSS-BLADE HOT MIGRATION	<b>No</b>	<b>Yes</b>

# When to have an Application in a Separate Virtual Machine

- The more critical or sensitive the application, the greater the reason it should be isolated on a separate VM.
- You should not group VMs handling critical or sensitive data with applications that handle routine or unrestricted information.
- By isolating critical applications, restoration priority becomes more straightforward in the event of a disaster.

# When to have an Application in a Separate Virtual Machine

- The more prone the application is to be security attacked, the greater the reason it should be isolated on a separate VM.
- The more complex the application, and/or the greater possibility of undesirable interactions with other applications, the greater the reason it should be isolated on a separate VM.

# When to have an Application in a Separate Virtual Machine

- The more vulnerabilities that, historically, have been identified with an application, the greater the reason it should be isolated on a separate VM.
- Security applications, such as firewalls and Network-based IDS should be isolated on a separate VM.
- The greater the risk for data corruption, or the more inherently unstable the application, the greater the reason an application should be isolated on a separate VM.

# When to have an Application in a Separate Virtual Machine

- Applications that have differing performance characteristics or resource utilization, should be isolated on a separate VM. This makes it easier to characterize and optimize resource requirements.

# Leveraging Transparent Page Sharing

- Similarity of consolidated operating systems.
- Similarity of consolidated applications.
- Current RAM utilization.
- **Design tradeoffs** -The more of the same instances of the same application, the greater the potential for duplicate pages to be reclaimed by transparent page sharing. However, there is also a greater likelihood of the same limiting resource to become saturated. I.E. multiple VM's that have CPU intensive workloads.

# How do you Access a VM?

- Remote Console software gives secure remote KVM access to the VM, as well as power and CD/DVD/Floppy drive access control.
- The CD/DVD drive access can be a physical drive, or an .iso file that can be mounted to be presented as if it were a CD/DVD. This can be on an ESX Server's local disk, SAN based disk, or network share. You can create a repository of OS and application installation disks, that can be easily accessed, shared and changed.



# Server Provisioning and Deployment

- As a VM is portable, and tools exist to remove “Globally Unique” information, a library of VM based templates can be easily made that have patches, hot-fixes, organizational customization, and applications pre-installed.
- These templates can be easily deployed, dramatically reducing the time it takes to provision new servers, in some cases, shortening the time from weeks to minutes.
- These templates can also provide a foundation for restoration of backup tapes.