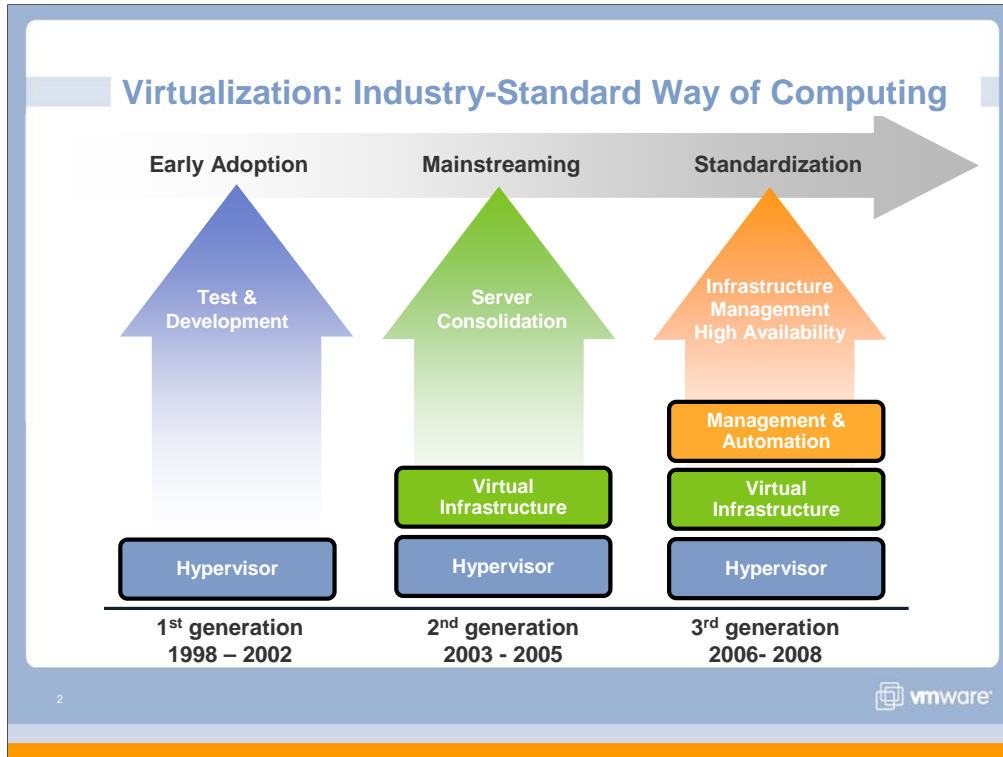





## **What's New with VMware Virtual Infrastructure**

November 2007









- Over the last few years, virtualization has gone from a technology being tried out in test/dev to a production server consolidation technology, and now is gathering steam as the industry standard way of computing.
- Early adopters of our technology used the hypervisor for basic partitioning. As we matured and provided means to aggregate multiple virtualized nodes and centralized management – customers rolled us out into mainstream production environments.
- As customers and our technology matured, virtualization began to go far beyond its original use for server consolidation and live migration of virtual machines. Ensuring availability and uptime helped customers achieve better service levels. Using virtualization for business continuity and DR helped customers achieve better RTOs and RPOs at a fraction of the cost. With the end to end management and automation capabilities available from VMware, it became very easy for customers worldwide to make VMware virtualization the default in the datacenter.
- 43% of customers surveyed last year said that their default policy for all or most new machines was a virtual machine.



- Qualcomm's experience illustrates well the typical customer adoption phases
- Qualcomm started a POC in the first half of 2003. Today, Qualcomm has implemented a VMware first policy, i.e. they have standardized on VMware for x86 workloads. Today 60% of Qualcomm's x86 environment is virtualized (1900 total servers/1150 are virtualized). The number of physical servers has grown from 950 to 1900 over the past 2.5 years, and because of the much simplified provisioning with virtualization, they have been able to maintain the same number of server admin's today. They provision 68 new VM's/month. This would be impossible in the physical world without dramatic staffing increases. Which means that the number of physical servers a single sys admin can manage has more than doubled. This translates into substantial operational savings for the company.
- Provisioning has gone from 6-8 weeks to between an hour and 2 days depending on the group's workload. They're is a joke within IT at Qualcomm that: "Projects kickoff in the morning and the servers are ready in the afternoon." They've built a web based, front ends and business users simply request a server via the web and it gets routed to the server admin's who determine where to place the VM etc.
- In aggregate, they've saved \$4.5M over 3 yrs with VMware. This calculation doesn't include the additional cost for storage in the virtual world (all VM's are SAN connected now), but it also doesn't include the cost savings from power, cooling etc.

## Leading Organizations are Standardizing on VMware Virtual Infrastructure

Company	Virtual Machines	% of x86 Infrastructure
	3,000+	30%
	1,700+	75%
	1,000+	75%
	1,100+	50%
	1,600+	50%+
	2000	50%

4



- ▶ The fast adoption of virtual infrastructure has given rise to a new type of enterprise: the virtualized enterprise.
- ▶ Some leading companies have already deployed thousands of virtual machines based on VMware technology, and have virtualized the overwhelming majority of their x86 infrastructure.
- ▶ Other companies with large number of VMs

## Drivers for Standardization are Emerging

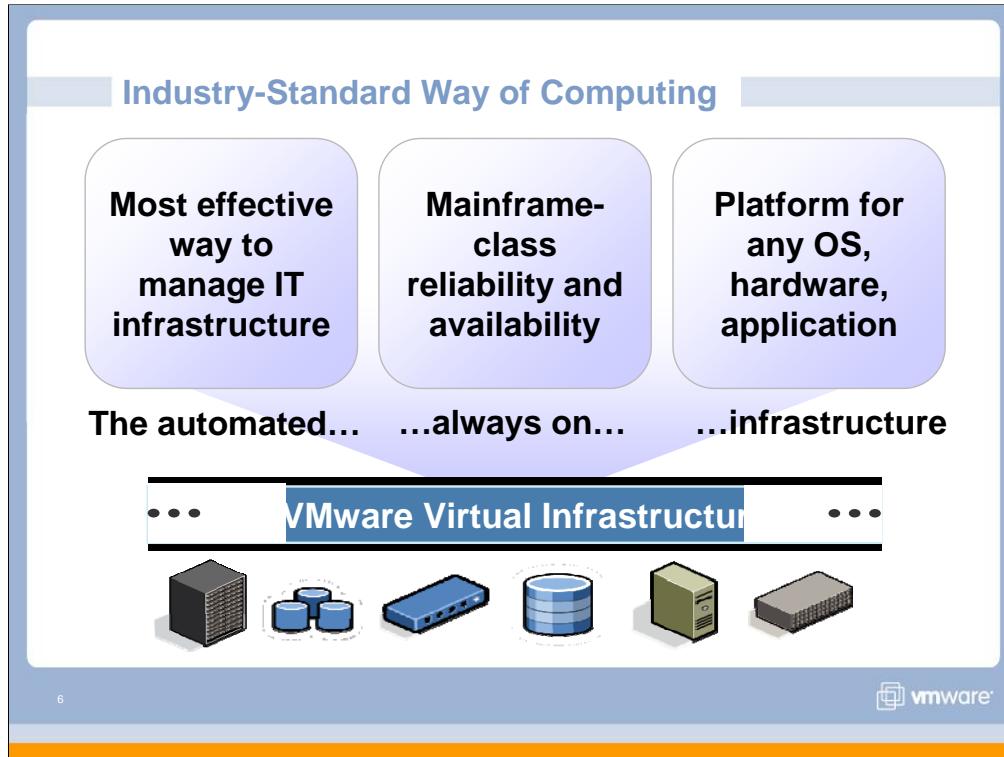
“While we initially chose VMware virtual infrastructure to address development hardware problems by reducing hardware costs and decreasing server deployment time, we soon discovered additional benefits to adopting the technology, including **server portability, protection, and availability.**”

Keith Leahy  
VP, Technology, Merrill Lynch

5



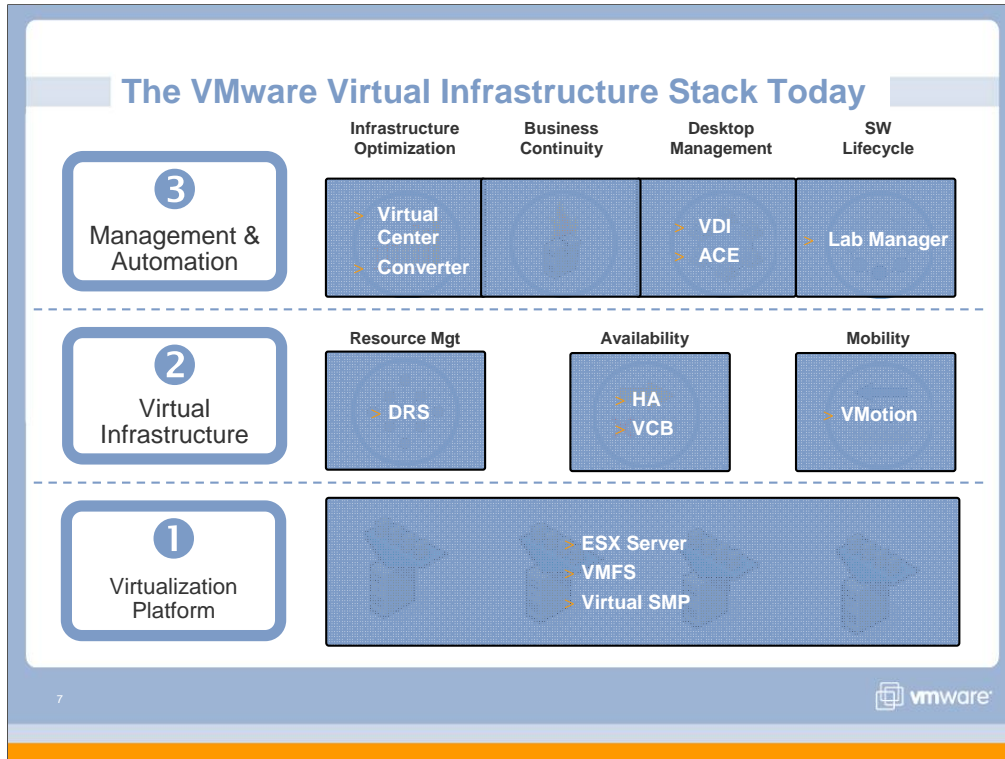
- Why this rapid adoption?
- Customers that initially started deploying VMware for server consolidation found enormous benefits not just from the cost savings- the reduction in provisioning time, the improved overall manageability of infrastructure, the abstraction away from physically disparate hardware and higher levels of availability all make for virtual environments being more flexible, adaptable and secure than physical environments.



➤ Why are customers standardizing on VMware Infrastructure?

➤ The three main pillars driving standardization on VMware Infrastructure are:

- VMware virtualization is a fundamentally more efficient of managing industry standard datacenters
- VI brings unprecedented ( mainframe class) levels of reliability and availability to x86 datacenters
- Because of its OS agnosticism and near native performance on a wide range of hardware, it is relevant for every application.



- Today our existing product line is represented in 3 layers –
- The base layer /virtualization platform is ESX Server – the highest performing, tried and tested, production –proven hypervisor. Over 20000 customers deploy ESX Server ( over 60% in production environments) for a wide variety of workloads.
- The virtual infrastructure capabilities have to do with farms or groups of virtualized servers. These capabilities leverage the abstraction away from physical hardware: to aggregate physical resources and present them uniformly to virtual machines for dynamic load balancing, to provide higher levels of availability than physical servers and allow mobility of virtual machines between different physical hardware with no disruption or downtime.
- The top layer represents end to end management and automation capabilities from VMware that automate specific IT processes from end to end.

**Industry-Standard Way of Computing**


**Most effective way to manage IT infrastructure**


**Mainframe-class reliability and availability**

**Platform for any OS, hardware, application**

The automated... ...always on... ...infrastructure

VMware Virtual Infrastructure



8 

- ▶ Let's talk first about how innovations in VI3.5 make it even more efficient to manage IT infrastructure.

## Customers Have Been Looking for a Better IT Management Model

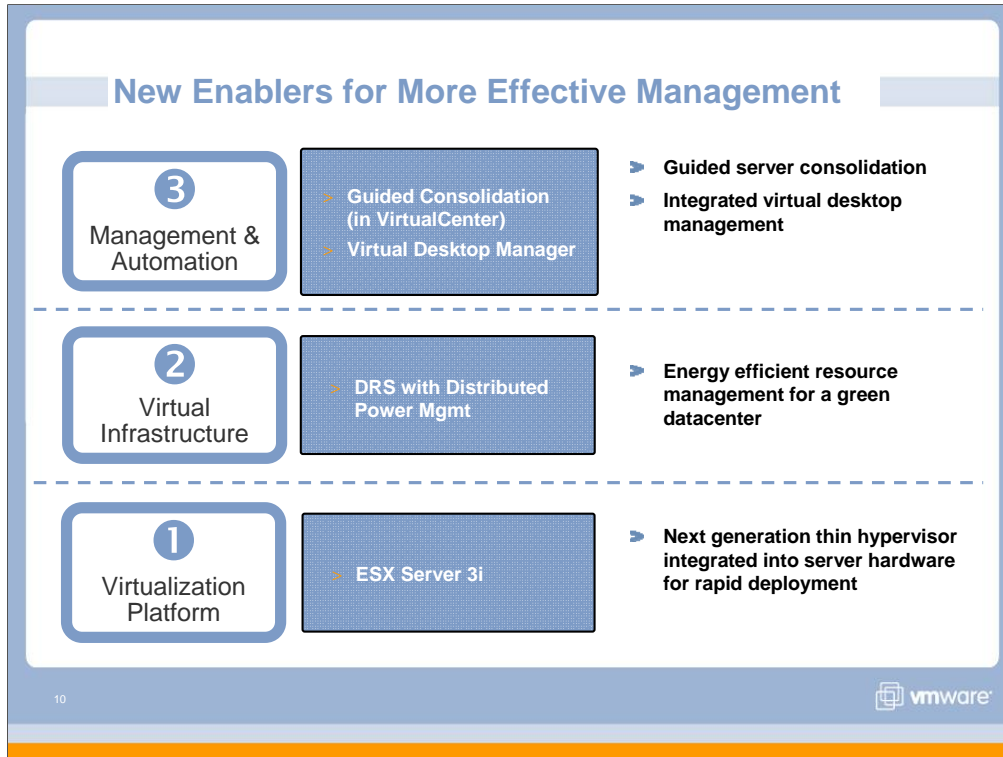
*“ Why can't my data center be as easy to use as the USB port on my Macintosh? It's plug and play, mix and match, multi-function, and it just works—every time, and on the fly. That's what I want, I want a USB Data Center! Whoever figures that out, will really be on to something big...”*

CIO  
Fortune 200 Manufacturer

9



- What is the customer requirement today?
- Customers don't want to deal with complexities of managing disparate moving parts. They don't want to spend time building complex bridges between things that don't talk to each other. They want an easy to use plug and play datacenter.



➤ So VMware introduces new capabilities across the three layers of the suite that make the VI3.5 suite an even better way to manage IT infrastructure.

➤ **Platform**

- ESX Server 3i – This is the next generation architecture hypervisor that comes integrated with the hardware simplifying install and deploy processes to a matter of minutes

➤ **Virtual infrastructure capabilities**

- DPM (experimental) automates energy efficiency in datacenters – shut servers off when resources are not needed by virtual machines, bring them back on when resource requirements increase

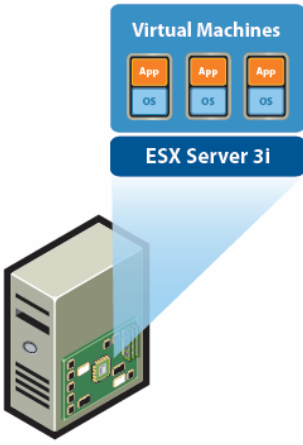
➤ **Management and automation**

- Guided Consolidation – this is focused on new virtualization users . In smaller/simpler windows environments, achieve benefits of consolidation quickly with this new interface.
- Virtual Desktop Manager is a new connection broker being introduced to provide a complete end to end solution for VDI deployments.

➤ **GENERAL** info on availability of these new capabilities:

- ESX Server 3i and ESX Server 3.5 are functionally equivalent – expected to be GA by end of year.
- DPM is a component of DRS- requires VC 2.5
- Guided consolidation is a component of VC 2.5
- VC2.5 is expected to be GA by end of year\VDM is part of the VDI solution sold separately from the VI editions. It is expected to be available by end of the year as well.


## ESX Server 3i



- Compact, 32MB footprint
- Only architecture with no reliance on a general purpose OS
- Integration in hardware eliminates installation
- Intuitive wizard driven start up experience dramatically reduces deployment time
- Standards-based management of the underlying hardware

- Server boot to running virtual machines in minutes
- Simplified management
- Increased security and reliability

11



- **ESX Server 3i is the next generation, thin hypervisor integrated in server hardware**
  - Hypervisor: on its own, ESX Server 3i offers basic partitioning of server resources. However, it also acts as the foundation for virtual infrastructure software, enabling VMotion, DRS, etc, the keys to the dynamic, automated data center
  - Thin architecture: Small footprint (32 mb) for security, reliability and simplified management
  - Server integration: Default feature makes deployment easy and fast
- **Two target markets:**
  - Existing customers: Get all the performance and sophistication of ESX Server and fully supports VI3 features and benefits in a very small footprint architecture
  - New customers: Easiest way to get started with virtualization
- **Additional functionality on top of the hypervisor just requires the right licenses, not any changes to the code itself. No reinstalls and no VMFS changes to go from running a standalone instance of the hypervisor to a full VI3 Enterprise deployment**
- Management is simplified because:
  - No Linux command line skills required
  - No user accounts or passwords to create and maintain
  - No OS security hardening, antivirus, or backup effort required
- Security and reliability is increased because:
  - Fewer interfaces minimize attack profile
  - Locked-down, BIOS-like interface prevents users from running arbitrary code
  - No dependence on failure-prone hard-drive enables disk-less servers
  - More than 50% fewer patches\*
- **ON STANDALONE AVAILABILITY:**
- A standalone , hard disk installable version will be available after the initial ESX Server 3i starts shipping embedded in servers. The date for this is still TBD.
- This standalone version may have a smaller compatibility list as compared to ESX Server 3

## From Server Boot to Running VMs in Minutes




1. Power on server and boot into hypervisor
2. Configure Admin Password
3. (optional) Modify network configuration
4. Connect VI Client to IP Address
  - Or manage with VirtualCenter

12



- We've also simplified the user experience for this new form factor – it becomes as easy to set up as a home router appliance.
- For e.g. It picks up IP addresses automatically using DHCP – or it can be manually set
- The whole setup time takes a few minutes, customers are ready to run VMs within minutes of booting into ESX Server 3i

## Enabling the 'Plug-and-Play' Datacenter



➤ **Plug:** Power on a new server with ESX Server 3i. The new server joins a DRS cluster.

➤ **Play:** All VMs in the cluster are automatically rebalanced taking in consideration the newly available resources

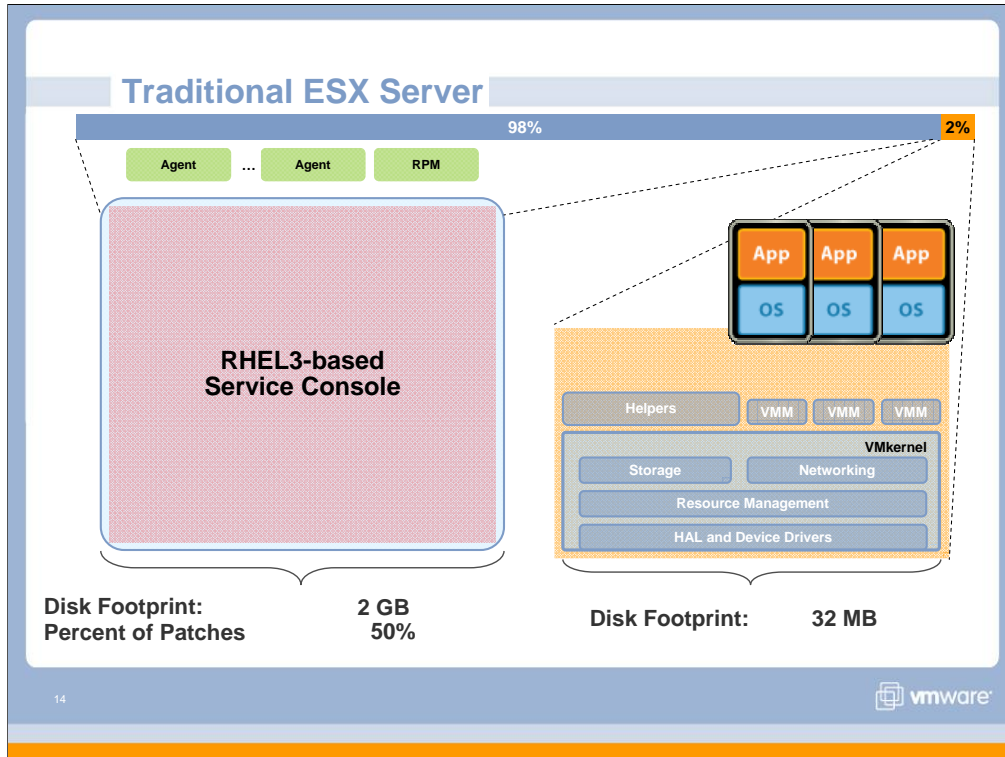
➤ On-demand capacity

➤ Easy scalability

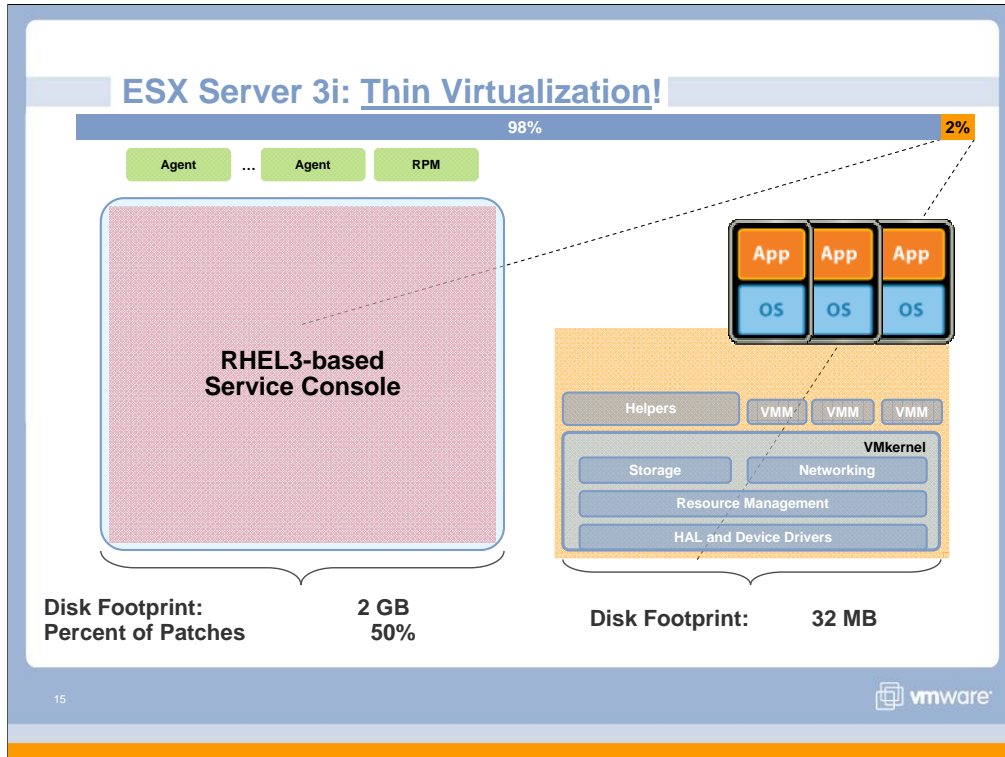
13



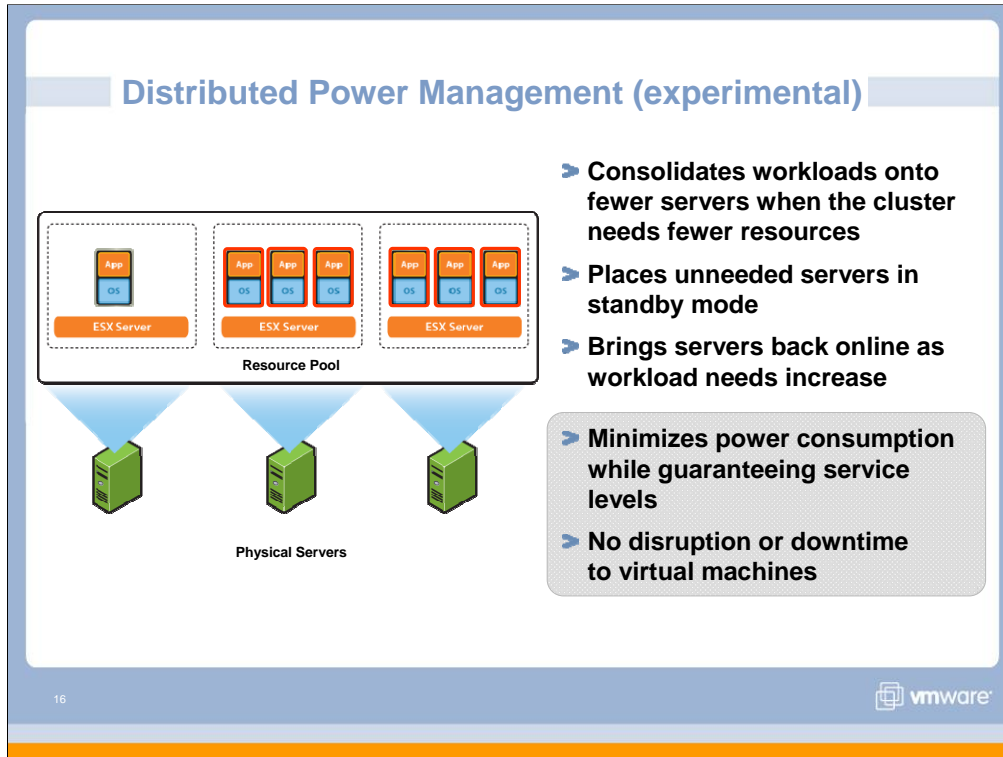
- How does this enable a plug and play datacenter?
- Plug in and run ESX Server 3i. DRS takes over and automatically manages this new compute capacity, migrating Virtual Machines to optimally load balance. Adding capacity and scaling your datacenter becomes a simple, non-disruptive process.
- ESX Server 3i nodes can be viewed as standardized compute nodes:
  - No installation required
  - Small footprint
  - Reliability and manageability of hardware



- Traditional ESX Server has a RHEL service console attached to it that is used primarily for booting up from a variety of different hardware, to run scripts and hardware monitoring agents etc.
- While ESX Server is not dependent on this service console ( the Console can be non-responsive but have no effect on VMs) – there is still extra effort administrators have to spend in securing this console, patching it and maintaining it



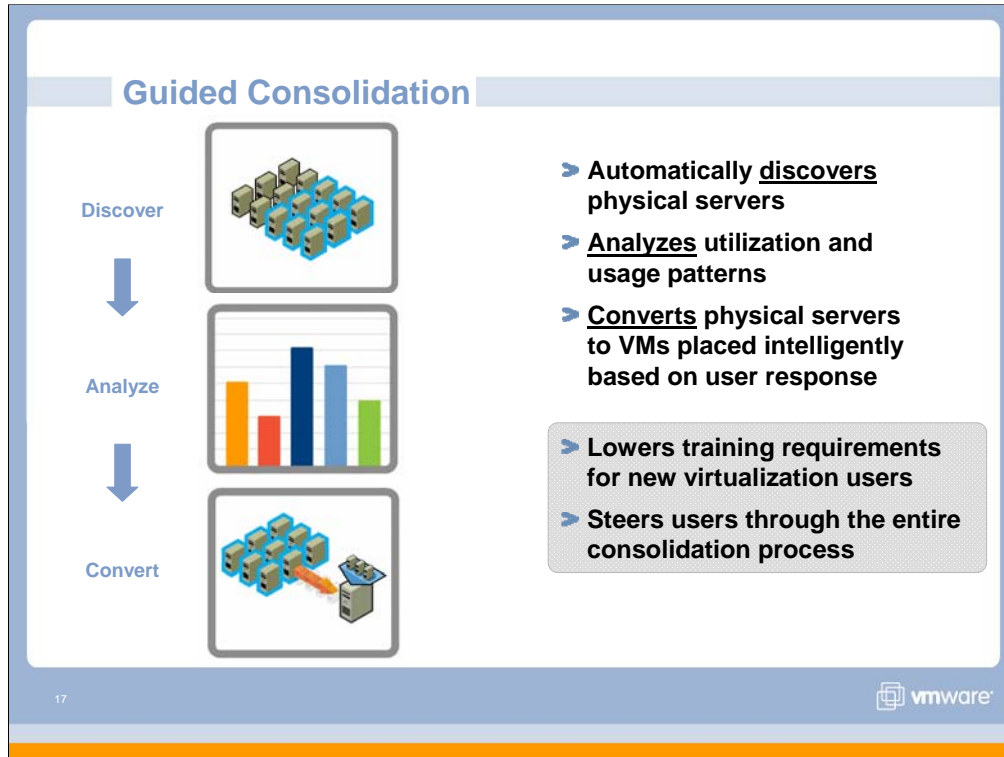
- ESX Server 3i now comes pre installed in the hardware eliminating the need for a service console. Scripts can be run in a remote CLI. Hardware monitoring agents can talk to the underlying hardware through CIM.
- The resulting compact 32MB footprint hypervisor is more easily maintained, updated and secure.



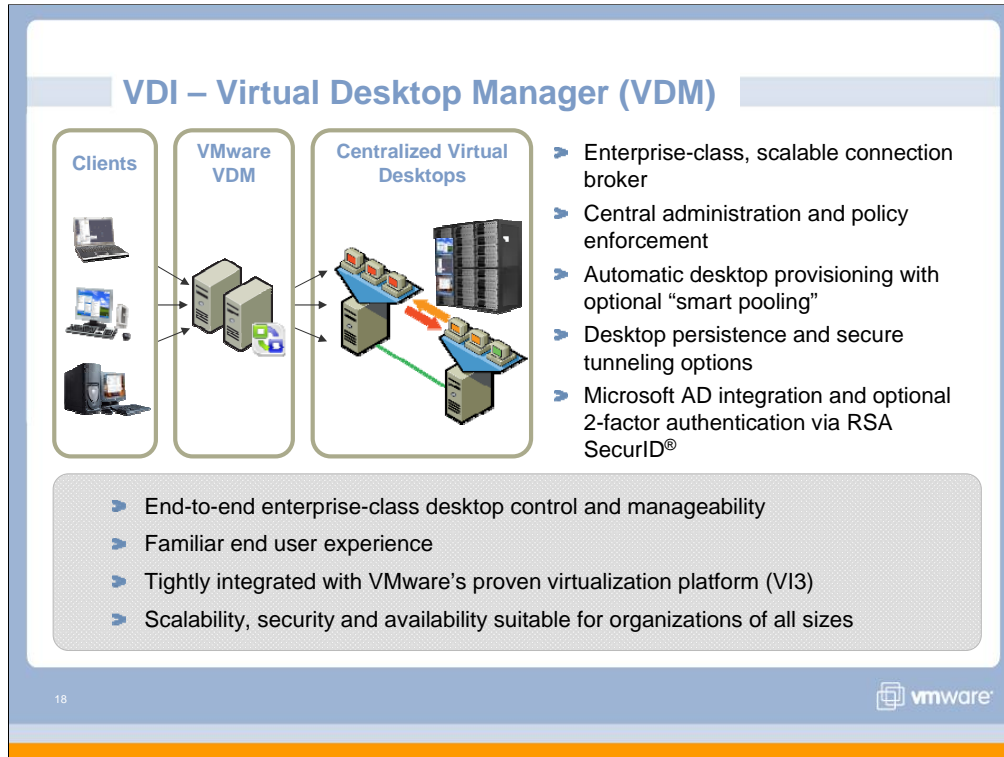
- Most servers consume 50% of their peak power requirement even when idle. Distributed Power Management helps you really manage your power bill without compromising on resource availability to virtual machines
- Put host in stand-by mode if:  $\text{total demand} + \text{reserve} \leq \text{total capacity} - \text{host capacity}$
- MORE DETAIL FOR INTERESTED CUSTOMERS:
- Users can define:
  - Reserve capacity to always be available
  - Time for which load history can be monitored before the power off decision is made.
  - Time for which load history can be monitored before the power on decision is made.
  - Power on will also be triggered when there aren't enough resources available to power-on a VM or when more spare capacity needed for HA.
- Stand-by mode means the host is powered off (S5). There are sub-modes to standby mode:
  - When a recommendation to go into standby is accepted, the host immediately enters standby/entering mode. In this mode the machine is still powered on and may still have some VMs running on it. They must be migrated off, and no new VMs can be started/migrated on the host.
  - When the host is evacuated, it can be powered off (standby/off)
  - At some later point, if a recommendation to leave standby mode is accepted, the machine is powered up and begins to boot (standby/leaving). No VMs can be migrated on it until it's fully up.
  - Automation options:
- Automation level
  - a) Automatic – DPM makes power on/off recommendations and executes them
  - b) Manual – DPM makes power on/off recommendations and admin chooses whether or not to execute them
  - c) Off – DPM is disabled
- Automation level defined on per-cluster & per-host basis
- Advanced options for tuning load parameters
- Manual power off/on of hosts

Requirements and considerations:

- Running Virtual Center with DRS
- VMotion compatible cluster
  - Shared storage
  - Compatible processor architectures
  - VMkernel/VMotion network
- Hosts with Wake on LAN
- NICs with Wake on LAN
- Configure these as the VMotion NICs



- > For first time virtualization users, a new feature in VirtualCenter 2.5 guides users through the process of server consolidation. Recommended for smaller, simpler environments – this feature steers users through discovering physical servers, collecting performance data from these servers and converting these servers to virtual machines placed intelligently on the most appropriate hosts.
- > Allows new users to quickly realize benefits from server consolidation, reduces training requirements for first time “virtualizes”
- > Current constraints in Guided Consolidation – FOR CUSTOMERS WHO ARE INTERESTED:
- > Guided Consolidation currently requires default username and password (needed for discovery) to only contain ASCII strings. As such the discovery feature may not work on non-English OSes. This is essentially a bug in collector which currently does not support Unicode. (we hope to fix this by GA)
- > Users cannot limit discovery to a certain IP range or filter search results based on wildcards or other search criteria
- > Guided consolidation does not discover across multiple AD sources as the collector engine can only discover and report on domains managed by a single AD source
- > Guided Consolidation relies on AD and/or LANMAN to discover machines on the network and gather data on the same. Systems that have LANMAN service disabled (typical in larger enterprise environments) will not show up in discovery in the absence of AD. In such a scenario, VC needs to be installed on a system that is part of an AD domain and the user needs to provide credentials that has read access on the AD domain to begin discovery and **data gather**.
  - **Needless to say in the absence of both LANMAN and AD, guided consolidation will not be able to discover any physical systems**



### ➤ Control and manageability in an end-to-end product:

- Central administration of desktops from any location including policy enforcement options such as session time-outs
- Enterprise-class connection brokering, managing connections between remote clients and their centralized virtual desktops
- Automatic provisioning of desktops using optional “smart pooling” capabilities
- Desktop persistence options: maintains desktop state on log-out or reverts to “known” state
- Secure tunneling via a proxy server so that all connections are completely encrypted
- Two-factor authentication via RSA SecurID®
- Full Microsoft Active Directory (AD) -based integration
- Policy enforcement:

### ➤ Familiar end-user experience:

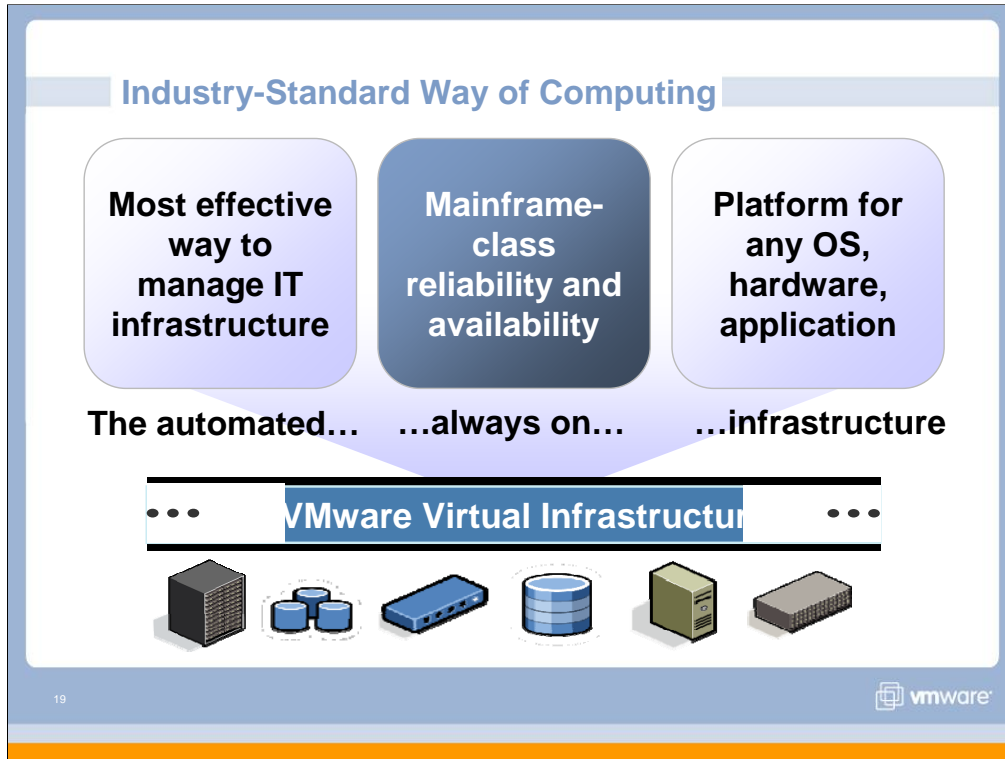
- Give users similar desktop experience, just running over a remote connection. No changes to the apps. No retraining.
- Individuals can customize their desktop. Desktops are isolated, not shared.
- Usability improvements to RDP for local printing, PDA, smart cards, etc.
- Secure Single Sign On to desktops on same domain as client

### ➤ Enterprise-class deployment, scalability & high availability:

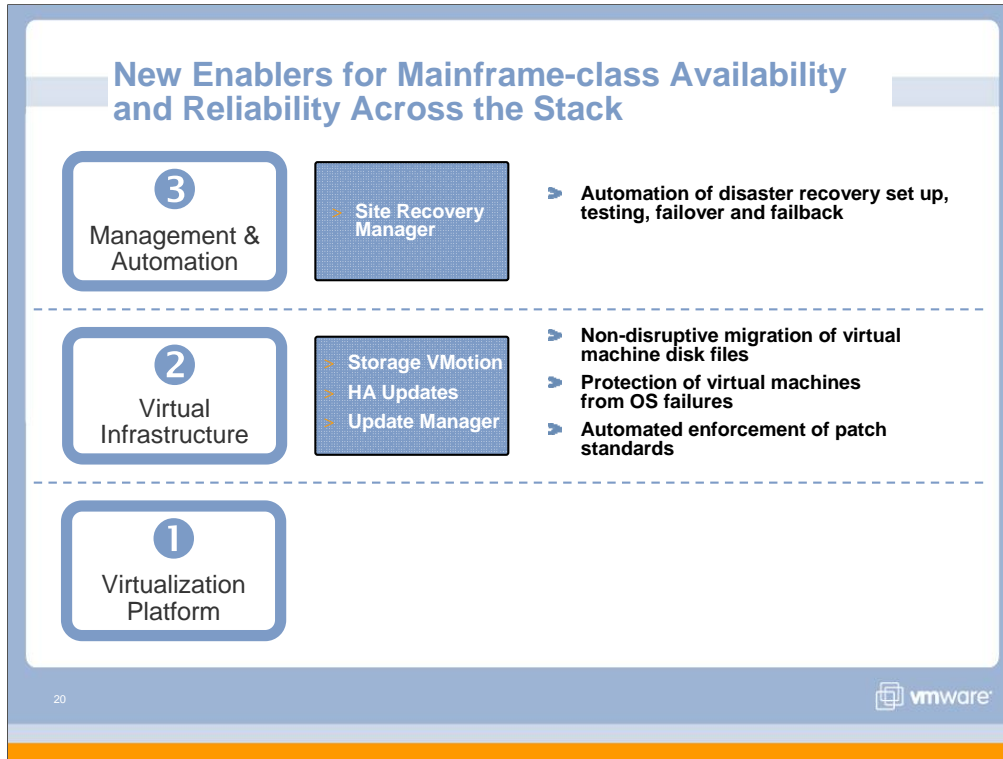
- Brings powerful VI3 and backend data center server capabilities to the desktop
- Deploy desktops in minutes instead of weeks through automated desktop provisioning
- Simplify backups and DR with centralized desktops that can leverage shared storage
- No single point of failure: fully synchronized connection mgmt servers and support for industry standard server load balancing
- Tighter integration with VI3 – the leading virtualization platform:
- Familiar VI3 look and feel for consistent administration across desktops and servers in the data center

### ➤ Third party extensibility

- Built by the leader in virtualization software (VMware)



- Next we talk about how new products/features from VMware improve reliability and availability in the datacenter



- Storage VMotion extends the concept of VMotion to storage arrays and allows virtual machine disks to be moved from one datastore to another with no disruption or downtime.
- HA enhancements now allow HA to protect against OS failures, allow HA to scale to 32-nodes and improve its robustness with more isolation addresses per hosts to prevent false restarts.
- Update Manager enforces greater security and reliability in the datacenter by automating enforcement of patch standards for ESX server hosts and virtual machines. ESX Server hosts are patched non-disruptively and automatically. Secure patching of offline or suspended virtual machines allows greater levels of compliance than physical environments.
- Site Recovery Manager is a new solution from VMware that is enabled by VC 2.5.
- Details on availability:
  - Storage VMotion is available as part of VI3 Enterprise. ( Customers with active Sns or SIs for VI3 Enterprise get this as a free upgrade with VC2.5).
  - Enhanced HA – is just updated HA in VC2.5
  - Update Manager is available as part of all three VI3 editions –starting with VC2.5. Customers with active Sns or SIs will get Update Manager when VC2.5 is generally available
  - Site Recovery Manager is not part of VI3; it is an addon product for VI3 environments. It will require VC 2.5. It is likely to be available early next year.

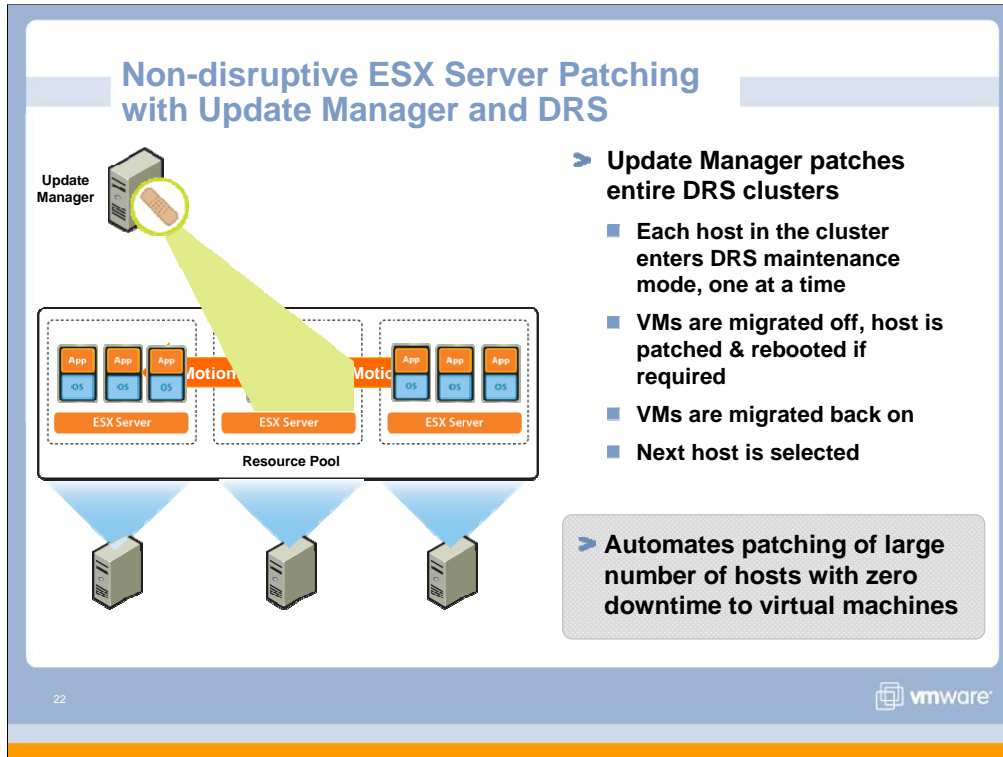
## VMware Update Manager

- **Automates patch management for ESX Server hosts and select Microsoft and RHEL virtual machines**
  - Scans and remedies online as well as offline virtual machines\* and online ESX Server hosts
  - Snapshots virtual machines prior to patching and allows rollback to snapshot
- **Eliminates manual tracking of patch levels of ESX Server hosts and virtual machines**
- **Automates enforcement of patch standards**
- **Reduces risk through snapshots and offline virtual machine patching**

21 \* Note: RHEL guests can only be scanned, not remediated

vmware

- What is VMware Update Manager?
  - It is an automated patch management solution for VMware ESX hosts as well as Microsoft and Linux virtual machines
- Two main benefits compared to traditional patching solutions :
  - 1. **patching of offline /suspended machines is done securely. Non compliant machines are patched in a quarantined state so that the rest of the network is not exposed to them**
  - 2. **Can patch and update ESX server and ESX Server 3i**
- MORE DETAIL FOR INTERESTED CUSTOMERS:
- VMware Update Manager is used to enforce compliance to patch standards in four steps:
- 1. **Getting information on the latest patches:** V
  - VMware Update Manager automatically gathers the latest patch data from VMware as well as application vendors such as Microsoft, Adobe and Mozilla via the Internet.
  - To enable secure off-network patching, VMware Update Manager has a disconnected patch mode. A separate download service assists Update Manager in gathering patch data used for off-network patches and updates.
- 2. **Set baselines :** The information collected by Update Manager is used to set baselines . Baselines contain one or more service packs, patches and /or updates
  - The baselines data that Update Manager gathers gives IT administrators granular control in defining patch levels.
  - These baselines updates can be static baselines defined manually or dynamic baselines that are set automatically depending on the significance of the patch data from the system vendor.
- 3. **Compare physical hosts and virtual machines against the baselines:**
  - VMware Update Manager scans the state of the physical VMware ESX Server hosts as well as select Microsoft and Linux guest operating systems and compares it with baselines set by the administrator.
  - Scans can be initiated on entire data centers, clusters, resource pools, templates, folders or individual hosts and virtual machines. They can be run immediately or scheduled as necessary. After a scan is complete, non-compliant machines are flagged for patch updates.
- 4. **Remediate the selected set of physical or virtual machines.**
- Virtual machine patching:
  - VMware Update Manager supports either manual or scheduled patching of the non compliant machines. If a reboot is required on a manual patch or update, the administrator has the option to reboot immediately or delay system restart by up to 60 minutes.
  - To reduce the risk of virtual machine patching failures, the VMware Update Manager automates snapshotting the virtual machine state prior to applying a patch. Snapshots are stored for a user-defined period so administrators can rollback a patched virtual machine to a known working state if there are any problems
  - VMware Update Manager also patches offline or suspended virtual machines. When remediating offline or suspended virtual machines, VMware Update Manager disables their NICs during the patching so the network is not exposed to non-compliant virtual machines.




- DRS enabled patching brings new levels of automation and resource management to the datacenter
- With DRS enabled patching, system admins can be almost completely hands-off in applying patches to ESX Server hosts. DRS manages the entire process for entire clusters of physical hosts with NO disruption or downtime to any of the virtual machines

## Update Manager Benefits

### Guest Patching

Administrative time – 3064 hrs, \$153,200 saved annually  
 Calculated for 100 virtual machines, assuming 75 patches per machine



### Offline Machine Patching

- Reduces exposure from non-compliant offline/suspended virtual machines
- Systems have NICs disabled during patching to reduce risk

23



- In a typical 100 VM environment, you save on 1 FTE dedicated to patching and maintaining VMs. This value is delivered in all three VI 3 Editions at no extra price!

## Manage All Types of Downtime

	Prevent planned outages	Quick recovery from unplanned outages
<b>Component</b>	<b>NIC Teaming Multipathing</b>	
<b>Server</b>	<b>DRS with Maintenance Mode VMotion</b>	<b>HA</b>
<b>Storage</b>	<b>Storage VMotion</b>	<b>VCB</b>
<b>Data</b>	<b>N/A</b>	<b>VCB</b>
<b>Site</b>	<b>Site Recovery Manager</b>	

24



- A variety of solutions available with VMware virtual infrastructure support management of all types of downtime

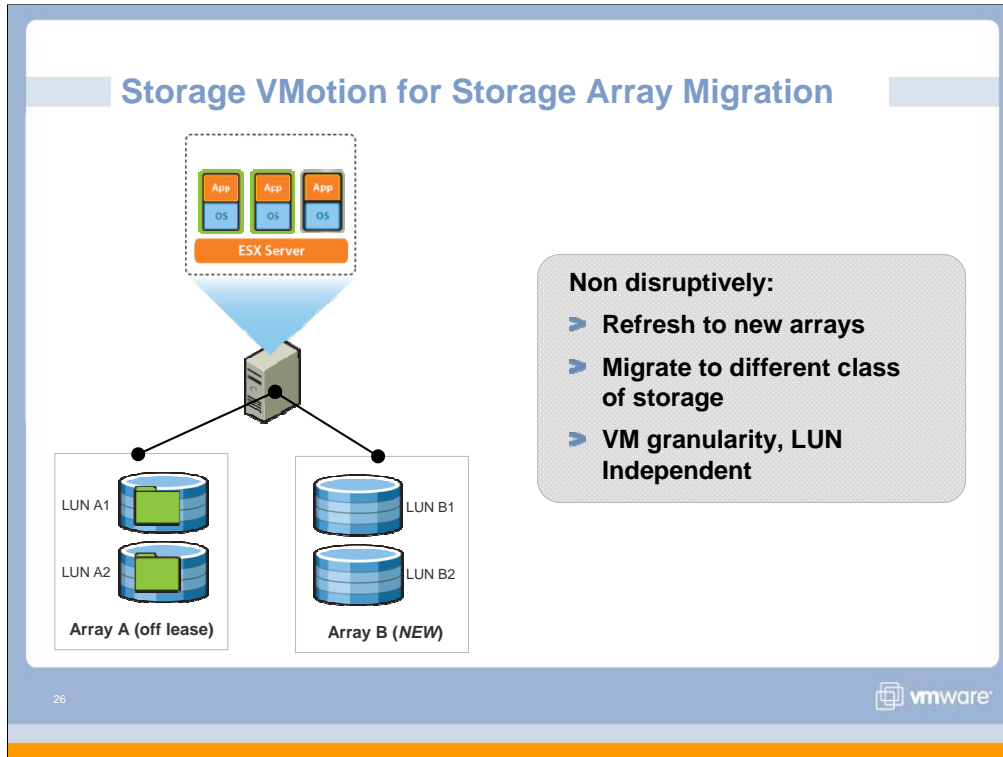
## Storage VMotion

- > **Storage independent migration of virtual machine disks**
  - Zero downtime to virtual machines
  - LUN independent
  - Supported for Fibre Channel SANs
- > **Minimizes planned downtime due to storage**
- > **Complete planned downtime management solution across servers and storage with VMotion and Storage VMotion**

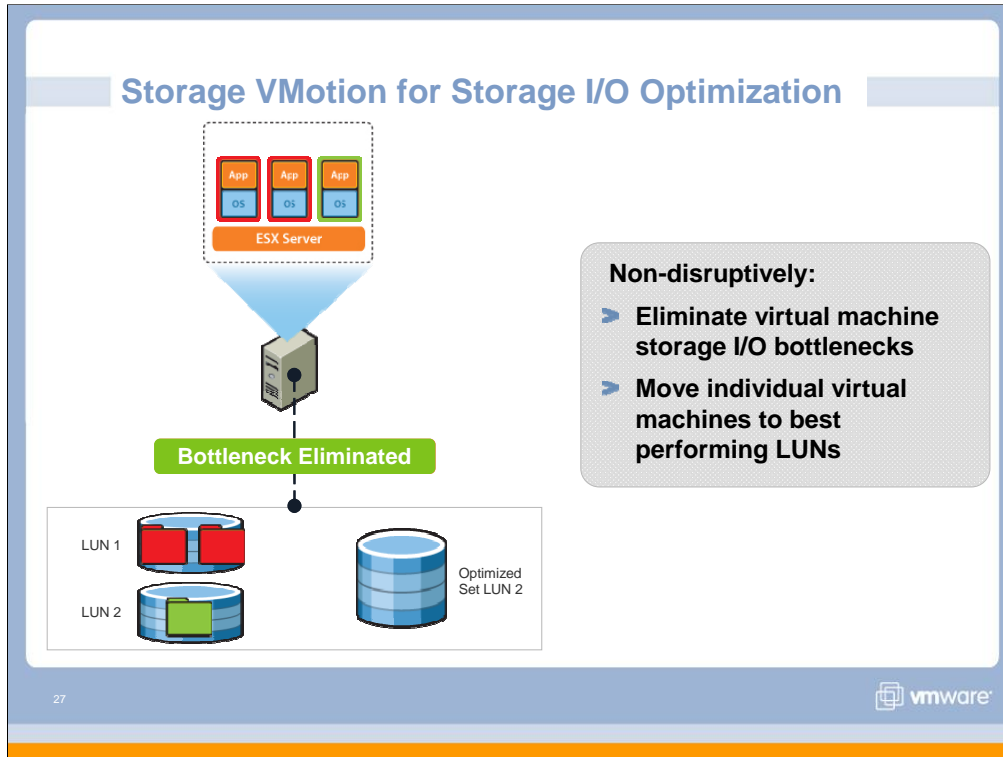
25

vmware

- > Storage VMotion completes the non-disruptive management of planned downtime
- > Migrating single VMs and their disks from one array to another is accomplished today by moving entire LUNs with data movers/SAN tools – almost always involving downtime
- > Now performance optimization of VMs with the right type of storage becomes a trivial problem to solve
- > HOW DOES IT WORK:
  - VMware Storage VMotion allows virtual machines' disks to be relocated to different data store locations completely transparently while the virtual machine is running, with zero downtime. Storage VMotion takes advantage of core technologies that VMware has developed, such as disk snapshots, REDO logs, parent/child disk relations, and snapshot consolidation. Before moving a virtual machines disk file, Storage VMotion moves the "home directory" of the virtual machine to the new location. The "home directory" contains meta information about the virtual machine i.e. configuration, swap, log files. It next "self VMotions" to the new VM home location. The disk movements follow after the home directory migration. First, Storage VMotion creates a "child disk" for each virtual machine disk to be migrated. Once the migration operation has started, all disk writes are directed to this "child disk". Second, the "parent", or original, virtual disk is then copied from the old storage device to the new storage device. In step three, the child disk that is capturing the write operations are then "re-parented" to the newly copied parent disks. And, in the final step the child disk is consolidated into the new parent disk and the ESX host is now re-directed to the new parent disk location. The switchover process of home directory and disk migration, creation of child disks and parent disks, re-parenting and consolidation of child disks happens within sub-2 seconds, fast enough to be unnoticeable to the application user.



- The traditional process of adding new storage disks and arrays is cumbersome, time-consuming and disruptive, with storage data migrations consuming IT resources and requiring significant service interruption. (Often these migration activities require extensive coordination between multiple IT groups)
- Storage VMotion helps customers embrace new storage choices, take advantage of flexible leasing models, cost effectively embrace new disk file formats and retire older, harder to manage arrays. Storage VMotion gives IT organizations the ability to non-disruptively move underlying virtual machine disk files from existing storage to any new storage of choice
- This operation can be performed within, and across storage arrays.
- It also helps IT organizations to conduct storage upgrades and migrations based on usage and priority policies (assumes that their arrays support storage tiering).



#### Use Case: Dynamically optimize storage I/O performance

**Customer Problem Statement:** Application performance is key metric to SLA success. Often, managing storage LUN allocations to support dynamic virtual machine environments can be a time consuming process, requiring extensive co-ordination between application owners, virtual server owners and storage administrators, also requiring off-lining of storage and causing application downtime. Customers sometimes compensate by over allocating high I/O LUNS to support an application which may never need it.

- The two VM's on the left have their data stores (red folders) on a LUN that is seeing increasing I/O activity leading to a hot spot that will soon cause the application performance on those VM's to suffer or get sluggish. Essentially this LUN is reaching its limit to supporting high IO and if left un-attended cause application performance issues.
- In this example – VMware Storage VMotion will migrate the affected VM's data store's, to the "optimized for High I/O LUN set" on the right, which the storage administrator has created on the existing array.
- Normally moving data store for a running application would involve off-lining the application. Storage VMotion will do this on-line (non-disruptively) in sub-2 seconds (we need not mention this), imperceptible to application users. This can be done even across arrays as shown in previous examples (**Important:** assumes that LUN sets have already been created – Storage VMotion does not create LUN's).

## HA Enhancements


- **32-node clusters**
- **Proactive cluster configuration checks**
- **Individual VM Failure monitoring (experimental) :**
  - **Monitors virtual machines for guest OS failures**
  - **Automatically restarts VM after specified interval**

- **Simple, cost effective availability for any workload**
- **Minimizes unplanned downtime due to hardware and OS failures**

28 vmware

- What does HA do today – its protects against hardware failures by restarting VMs automatically on alternate ESX hosts
- With ESX 3.5/VC 2.5, HA scalability is improved to 32-node clusters
- Cluster configurations are now proactively monitored for abnormal configurations such as : common DNS misconfigurations, right amount of network redundancy in the service console.
- VM failure monitoring extends HA protection to OS failures as well. Initially this feature is experimentally supported.
- Guest OS agnostic – this feature extends HA protection to OS failures across the board
- Heartbeat monitoring through VMware tools
- Users can define parameters around how long before a VM is declared failed, how much time to restart, how many attempts to restart.

## Consolidated Backup (VCB) Enhancements



- Support for VMs running on iSCSI, NAS or local storage
- Run VCB within a VM
- Use VMware Converter to restore VCB images
- Full GUI Integration with leading partners

- Simplified data protection setup and configuration
- 1 step restore procedure for virtual machines


29



- VCB now supports storage that is attached to the VM by iSCSI, NAS, or is locally attached (released in 3.0.2 but not widely known). This is a cost savings element that makes VCB more attractive in the SMB space.
- Ability to run VCB in a VM and eliminates the need for a dedicated backup proxy server running on Windows. This is a cost savings element that makes VCB more attractive in the SMB space and also enables the use of VI features such as VMotion and DRS to maximize the use of existing resources
- VMware Converter to restore VCB images – (Released with 3.0.1 but not widely known) Restore VCB images of VMs to running virtual machines. This is a simple 1 step graphical way to take VMs restored from tape and put them back into VI
- Full GUI Integration with Leading Partners – use the NetBackup 6.5 example (just released in August) that will enable VCB to run w/o scripting and be configured completely through the NetBackup GUI

## Introducing VMware Site Recovery Manager

Site Recovery Manager leverages VMware Infrastructure to transform disaster recovery



- **Simplifies and automates disaster recovery workflows:**
    - Setup, testing, failover, failback
  - **Provides central management of recovery plans from VirtualCenter**
  - **Turns manual recovery processes into automated recovery plans**
  - **Simplifies integration with 3<sup>rd</sup>-party storage replication**
- **Makes disaster recovery rapid, reliable, manageable, affordable**

30



### ➤ Provides a central point of management for disaster recovery plans

- Plug-in to VMware VirtualCenter
- Manage recovery plans for virtual machines managed by VirtualCenter

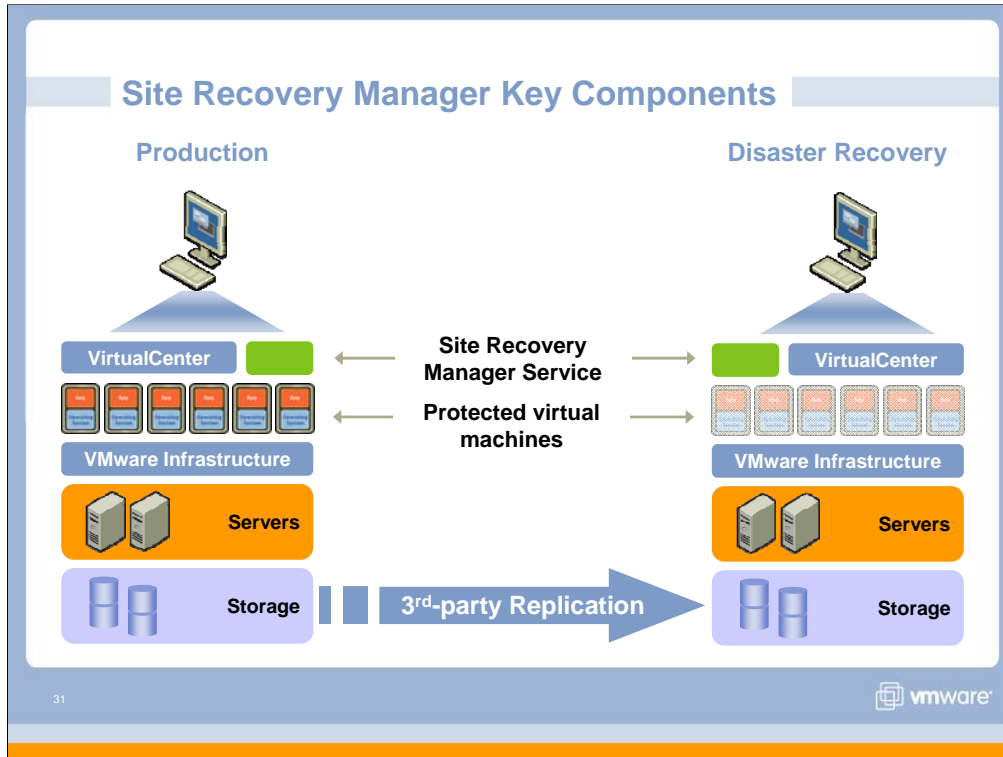
### ➤ Enables pre-programming of disaster recovery plans

- Map resources between production and recovery sites
- Program steps of recovery process

### ➤ Automates key disaster recovery workflows

- **Simplified setup of recovery plans.** Create, configure, and manage recovery plans from VMware VirtualCenter, making it possible to manage disaster recovery plans as an integral part of virtual infrastructure. Turn the complex, manual runbooks that are commonly required in traditional disaster recovery solutions into automated recovery workflows maintained and executed automatically at a push of a button from within Site Recovery Manager.
- **Automated failover.** Once recovery plans have been created, automate execution of those recovery plans in order to make recovery dramatically faster by eliminating the slow and unreliable manual processes otherwise required.
- **Easier testing of recovery plans.** Perform non-disruptive, automated tests of the same recovery plans that will be used in a failover scenario.

### ➤ Expected to be available in H1 2008. will be a separately licensed component



- This slide illustrates the key components of a site recovery manager deployment
- The requirement is for protected VMs to be running off of storage that is being replicated to a secondary site. Virtual Center is required to manage both sites
- Site Recovery Manager manages the mapping of components ( VMs, resource pools, networks etc) between the two sites and provides workflow automation for setup, failover, failback and test of this DR environment.

## Summary of Benefits



### Site Recovery Manager Leverages VMware Infrastructure to Make Disaster Recovery

- > **Rapid**
  - Automate disaster recovery setup, failover, failback, and testing
  - Eliminate complexities of traditional recovery
- > **Reliable**
  - Ensure proper execution of recovery plan
  - Enable easier, more frequent tests
- > **Manageable**
  - Centrally manage recovery plans
  - Make plans dynamic to match environment
- > **Affordable**
  - Utilize recovery site infrastructure
  - Reduce management costs

32 These features are representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.



## Industry-Standard Way of Computing

**Most effective way to manage IT infrastructure**

**Mainframe-class reliability and availability**

**Platform for any OS, hardware, application**

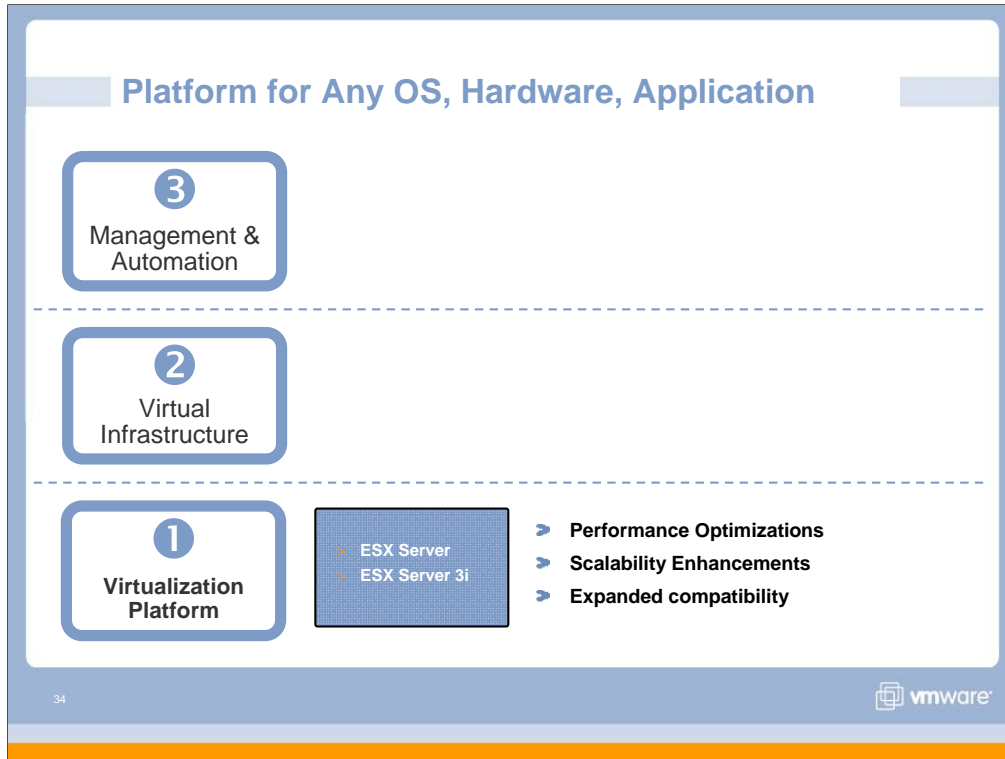
The automated...

...always on...

...infrastructure

... VMware Virtual Infrastructure ...





- Lastly we talk about enhancements in ESX server 3.5 and ESX server 3i 3.5 that optimize virtual machine performance, prove the scalability of the platform and expand its compatibility

## Virtualization Platform Enhancements

**Virtual Machines**

**ESX Server**

**CPU Memory**

**Network**


**Storage**

Performance	Scale	Compatibility
<ul style="list-style-type: none"> <li>➤ Paravirtualization</li> </ul>	<ul style="list-style-type: none"> <li>➤ 64GB virtual RAM</li> </ul>	<ul style="list-style-type: none"> <li>➤ Ubuntu</li> <li>➤ Windows Vista</li> </ul>
<ul style="list-style-type: none"> <li>➤ Hardware page table assist</li> <li>➤ Large memory pages</li> </ul>	<ul style="list-style-type: none"> <li>➤ 128 GB of physical RAM</li> </ul>	
<ul style="list-style-type: none"> <li>➤ TCP Segment Offload</li> <li>➤ Jumbo Frames</li> </ul>		<ul style="list-style-type: none"> <li>➤ 10 GigE</li> <li>➤ Infiniband</li> </ul>
<ul style="list-style-type: none"> <li>➤ NPIV Support</li> </ul>		<ul style="list-style-type: none"> <li>➤ SATA devices</li> <li>➤ SAS arrays</li> </ul>

35

- First we cover Scale & Compatibility – then a deep dive into performance.
- With ESX server 3.5 ; VMs as big as 64GB of memory can run, ESX server can run on hardware shipping with up to 128 GB of physical RAM ( experimental support for 256GB memory).
- On the compatibility side, a growing ecosystem around VMware have partnered with us to create expanded storage and networking choices for ESX Server by contributing through the community source program.
- High speed networking interconnects (10 Gig E and Infiniband ) are now supported with ESX server 3.5. (Note: iSCSI over 10 GigE is currently experimental.)
- SATA local storage is also supported.

## Virtualization Platform Enhancements Paravirtualization



- Paravirtualization is making the guest OS virtualization aware and more performant running on a hypervisor
- VMware led successful industry collaboration effort to create a standard interface - paravirt ops - that allows standard Linux kernel (2.6.21) to be VMware virtualization aware
- Ubuntu 7.04 now uses this interface

- Enables Linux to run at higher performance by making it virtualization-aware
- Performance benefits seen in large databases such as Oracle

## Virtualization Platform Enhancements – Memory Management

The diagram shows a stack of Virtual Machines on top of an ESX Server. Below the ESX Server are three hardware components: CPU Memory (highlighted with a red box), Network, and Storage.

This diagram shows two processes, each with its own Virtual Memory (blue blocks) and Guest Physical Memory (grey blocks). Arrows indicate that the Guest Physical Memory is mapped to Host Memory (green blocks) by the ESX Server.

ESX Server translates guest VM memory to physical memory

- Uses **CPU hardware page tables** to reduce memory translation overhead
- Large memory pages for more efficient memory access

- Benefits memory intensive workloads such as databases, Citrix, Terminal Services

37 vmware

- ESX Server translates the guest OS physical memory to actual hardware physical memory. This translation when done in the CPU hardware is faster than when done in software
- Larger memory pages for the guest OS and the hypervisor enable more efficient memory access
- The net effect of this is that memory intensive workloads such as databases, Citrix etc see benefits in performance

## Virtualization Platform Enhancements Network

The diagram illustrates the ESX Server architecture. At the top, five Virtual Machines are shown. Below them is the ESX Server layer. Underneath the ESX Server are four hardware components: CPU Memory, Network (highlighted with a red border), and Storage. The Network component is represented by a blue network switch icon.

- TCP Segmentation Offload\* reduces CPU overhead by streamlining TCP processing
- Jumbo frames\* allow reduced TCP/IP processing by enabling larger frame sizes
- Supported for ESX Server hosts and select guests
- Support for 10 GigE for network intensive workloads\*\*


- Benefits network intensive workloads (e.g. backups over LAN, webservers, Citrix, Win Terminal Services, iSCSI) in virtual machines

\* Supported for ESX Server hosts and select guests  
\*\*10 GigE not supported for the ESX iSCSI software initiator

38

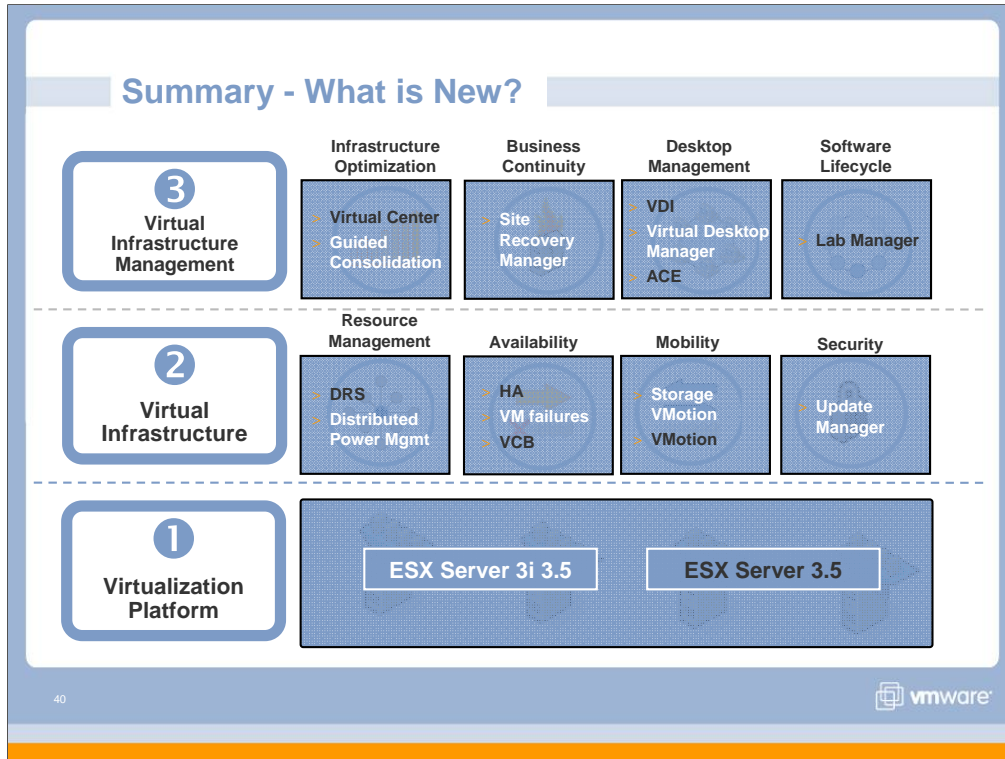
- **Networking performance** ESX Server 3.5 introduces these enhancements that reduce the overhead associated with processing network i/o.
  - o **TCP Segment Offload (TSO)** dramatically improves performance by reducing the CPU overhead involved with sending large amounts of TCP traffic. It is supported both in the guest operating system as well as in the ESX Server kernel TCP/IP stack. Thus it improves performance for TCP data coming from a virtual machine and for traffic, such as VMotion, that is sent out of the server. TSO is enabled for most virtual machines automatically. With NICs that support segmentation offload there is some additional benefit.
- To enable TSO in a virtual machine, you must select either the Enhanced vmxnet or e1000 virtual device. iSCSI with TSO is supported. TSO is only enabled with select guests. Guest list will be available with GA documentation.
- **Jumbo Frames** allow ESX Server 3.5 to send larger frames out onto the physical network. The network must support jumbo frames (end-to-end) for jumbo frames to be effective. Jumbo frames up to 9kB (9000 Bytes) are supported. NICs must support jumbo frames- there will be a separate HCL for this list. iSCSI with Jumbo Frames is not supported. This feature is available only for select guest operating systems and needs to be enabled via the command line.

## Virtualization Platform Enhancements Performance





- NPIV allows assignment of virtual HBAs to individual virtual machines
- Enables QoS through SAN tools


- Enables performance optimization through more granular traffic monitoring



- > In Summary – a wide range of products and features from VMware enable further rapid adoption across the board. Extending out the virtualization capabilities of VMware virtual infrastructure helps make it the right choice for virtualization in datacenters worldwide.

## Other Enhancements with VI 3.5

<b>3</b> <b>Virtual Infrastructure Management</b>	<ul style="list-style-type: none"> <li>➢ VirtualCenter scales to 200 hosts and 2000 VMs</li> <li>➢ Multimedia application support for VDI deployments</li> <li>➢ Automation for VMware Tools batch and scheduled install</li> <li>➢ Network CDP Support</li> </ul>			
VirtualCenter				
<b>2</b> <b>Virtual Infrastructure</b>	<b>Resource Management</b> 	<b>Availability</b> <ul style="list-style-type: none"> <li>➢ Increased HA isolation addresses</li> <li>➢ Failover customizations</li> </ul>	<b>Mobility</b> <ul style="list-style-type: none"> <li>➢ VMotion with local swap files</li> </ul>	<b>Security</b> 
VirtualCenter				
<b>1</b> <b>Virtualization Platform</b>	<ul style="list-style-type: none"> <li>➢ Experimental IOAT v1 support</li> <li>➢ IPv6 support for guests</li> </ul>			

41


➤ Other enhancements of interest – possibly for specific audiences.