

Virtual Infrastructure 3: Beta to Production

Speaker Name: Doug Baer
Senior Consultant, IT Partners



VMWORLD 2006

Agenda

- **Welcome/Introductions**
- **Upgrade Overview**
 - Plan ahead
 - Use a systematic, phased approach
 - A typical upgrade
- **Upgrade Detail**
 - Diving into individual phases
- **New Features**
 - Resource Pools
 - VMware DRS
 - VMware HA
 - VMware Consolidated Backup
- **Q&A**

Welcome / Introduction

■ Doug Baer, Senior Consultant, IT Partners

- Based in Phoenix, Arizona
- Partnerships
 - VMware VIP Enterprise Premier, VAC
 - Microsoft (Gold)
 - Hewlett-Packard



Upgrade Overview

■ Plan ahead

- Discovery: know what you are getting into
- Be prepared to lab activities to verify steps

■ Systematic, phased approach

- Break into manageable pieces
- Have a plan – this is not just a simple software install

■ Typical upgrade tasks

- Discovery and pre-upgrade tasks
- Upgrade VirtualCenter Server and Clients
- Upgrade ESX Hosts
- Upgrade Virtual Machines
- Upgrade VMware Tools
- Clean up and implement new features

Upgrade Detail – Outline

■ Pre-Upgrade

- Licensing
- Discovery/Validation
- Upgrade method?

■ Phase I: VirtualCenter

- VirtualCenter Server
- VirtualCenter Database
- Deploy Virtual Infrastructure Clients

■ Phase II: ESX Hosts and Datastores

- VMFS-2 to VMFS-3

■ Phase III: Virtual Machines

- New virtual hardware for additional functionality

■ Phase IV: VMware Tools

■ Post-Upgrade: Clean up and implement new features

Upgrade order is
very important:
each phase
constitutes an
irreversible process

Pre-Upgrade

- **Backup:** Backup everything prior to any upgrade
- **Licensing:** Acquire all licenses in advance
- **Order:** Complete one phase prior to starting the next phase

- **What makes this upgrade different?**
 - Completely new VMFS partition type, VMFS-3
 - To ESX 3 hosts, VMFS-2 partitions are *read-only*
 - To ESX 2 hosts, VMFS-3 partitions are *unreadable*
- **Run `preupgrade.pl` Script on ESX Server 2 host**
 - Script on ESX Server 3 Installation CD-ROM
 - `mount /mnt cdrom`
`perl /mnt/cdrom/scripts/preupgrade.pl`

- **Notes from the field**
 - Licensing!
 - Commit all VMDK REDO files prior to upgrade
 - Clusters must be migrated from virtual to RDM disks

Risk Mitigation



■ Have a Backup Plan!

- VirtualCenter + VirtualCenter database
- ESX Server host
 - Service Console
 - VMs: all .vmx files
 - Local image files: .iso, .flp, exported .vmdk files
 - Modified configs, `/etc/passwd`, `/etc/groups`, custom scripts
 - Local VMFS
 - All virtual disk and virtual machine template files
- SAN-based VMFS
 - Use SAN snapshots or backup/restore software
- VMs
 - Clone prior to virtual hardware upgrade

Risk Mitigation (continued)

■ Understand and Test Back-Out Procedures!

- VirtualCenter
 - Reinstall original VirtualCenter Server
 - Restore full backup of VirtualCenter database
- ESX Server host
 - Reinstall original ESX host software / reimage
 - Restore Service Console and local VMFS files
- Shared VMFS
 - Reformat VMFS-2 from ESX 2.x host and files
- VMs
 - Revert to clone taken prior to upgrade

Upgrade Types

- **In-place:** take down, install, bring up, clean up
 - **Benefits:** Speed, simplicity
 - **Drawbacks**
 - All virtual machines must be powered off simultaneously
 - Upgrade installation requires 2GB root partition w/850MB free
 - **Process**
 - Shut down all virtual machines
 - Boot host from ESX Server 3 Install media
 - Upgrade to ESX 3 and VMFS-3
 - Reboot host, upgrade and power up virtual machines
 - **Candidates**
 - Standalone ESX hosts
 - No SAN connectivity, or HVL aligned (more on this later)
 - Possible to take all virtual machines down at once
 - Downtime is easier to obtain than additional hardware and storage

Upgrade Types

■ **Migration:** managed, gradual transition

> **Benefits**

- Minimized downtime for critical systems
- Not an “all or nothing” proposition

> **Drawbacks**

- Requires additional resources (ESX host, shared storage)
- Longer upgrade process = coexistence considerations

> **Process**

- Install ESX 3 Server on new host and new VMFS-3 partition
- Move virtual machines from ESX 2.x host(s) to new host
- Once ESX 2 host is empty, rebuild as ESX 3 host; Repeat

> **Candidates**

- Environments with critical virtual machines (sensitive to downtime)
- Not possible to shut down all virtual machines simultaneously
- New hardware and storage is cheaper than downtime

Upgrade Types



■ esxMigrator

> Benefits

- Disks migrated while virtual machines are online
- Cutovers can be scheduled for low-traffic times
- Single virtual machine reboot accomplishes all upgrade tasks for that VM

> Drawbacks

- Requires additional ESX host and software purchase

> Process

- Install esxMigrator on Windows machine (can be virtual)
- Attach to ESX 2 and ESX 3 hosts from esxMigrator
- Select per-virtual machine migration options and start sync
- Once ESX 2 host is empty, rebuild as ESX 3 host; Repeat

> Candidates

- Software is cheaper than hardware, storage, or downtime
- ESX 3 host hardware different from ESX 2 (no VMotion)
- Not possible to shut down all virtual machines simultaneously

Discovery/Validation

- **VirtualCenter Server on a shared server?**
 - Not on a web server – VirtualCenter uses tcp/80 and tcp/443
 - Not on a GSX/VM Server host – VirtualCenter uses tcp/902

- **VirtualCenter Server in a virtual machine?**
 - Do not use same virtual machine as VirtualCenter database
 - VMware License Server should be on a physical server

- **Collect additional information**
 - ESX 2 version and VMFS-2 partition checks (`pre-upgrade.p1`)
 - Location of VMDK files for all registered VMs on an ESX host
 - Map running virtual machines to current ESX 2.x hosts
 - Backup `/home` to preserve `.vmx` files (virtual machine configurations)

Upgradeable Versions

- Minimum versions are VirtualCenter 1.2 and ESX Server 2.1.1
- VirtualCenter versions prior to 1.2 must be upgraded to 1.2 or higher in order to preserve the VirtualCenter database
- Upgrade from VMFS-1 is not supported
- Upgrade from VMFS-2 with >8 MB block size is not supported
- Use migration method instead of in-place upgrades to handle unsupported ESX Server versions

<i>VirtualCenter Version</i>	<i>Upgrade Support</i>
Beta release (any)	NO
1.0	Upgrade to 1.2 first
1.1	Upgrade to 1.2 first
1.2	YES
1.3	YES
1.3.1	YES
1.4	YES

<i>ESX Server Version</i>	<i>Upgrade Support</i>
Beta release (any)	NO
1.x	NO
2.0.x to 2.1.0	NO
2.1.1, 2.1.2, 2.1.3	YES
2.2.0	YES
2.3.x	NO
2.5.0	NO
2.5.1 or higher, 3.0	YES

Discovery/Validation

■ VirtualCenter Server OS

- Windows 2003, any 32-bit release
- Windows 2000 SP4 + Update Rollup 1 + MDAC 2.6
- *Windows XP Professional*

■ VirtualCenter Database

- Microsoft SQL 2000 SP4
- Oracle 9iR2, 10gR1 (versions 10.1.0.3 and higher only), and 10gR2
- *MSDE*

Upgrade Detail – Phase I

■ Preparation

- License file(s)
- Install media
- Database credentials

■ Expected Downtime

- Virtual Machines: None
- ESX hosts: None
- VC Server: approx. 30 mins, depending on database size

■ Complete

- VirtualCenter 2.x

■ Remaining

- ESX 3.x
- VMFS-3
- Virtual Machines
- VMware Tools

■ Upgrade VirtualCenter

➤ VirtualCenter Server

- Windows software install, removes VC Server 1.x

➤ VirtualCenter Database

- Save Performance or Events & Tasks logs?
 - Sacrifice historical data to speed upgrade?
 - Significantly faster if not selected
- Large Databases
 - May time out waiting for **vpzd** service to start
 - Does not mean the upgrade has failed!

➤ Virtual Infrastructure Clients

- May be deployed in advance of server upgrade

Templates

■ Options for Handling Templates

> Pre-upgrade

- Deploy each to a virtual machine; clone/upgrade as needed

> Post-upgrade

- Import old templates – may be deployed to ESX 2 hosts or ESX 3 hosts as legacy virtual machines only
- Convert “legacy” templates for use by ESX 3 hosts
 - Template datastore must be upgraded to VMFS-3

> The approach depends on

- Extent of template use
- Duration of coexistence
- Directives (all new virtual machines provisioned to ESX 3 hosts?)
- Current datastore for templates (VC server vs. shared storage)

Upgrade Detail – Phase II

- Preparation
 - Phase I complete
 - Licenses
 - Install media
- Expected Downtime
 - Virtual Machines: depends on upgrade method
 - ESX host: 40-60 minutes per host
 - VC Server: None
- Complete
 - VirtualCenter 2.x
 - ESX 3.x
 - VMFS-3
- Remaining
 - Virtual Machines
 - VMware Tools

■ Upgrade ESX Hosts and Datastores

- **ESX 3 hosts only run VMs from VMFS-3 partitions**
- In-place Upgrade
 - Downtime incurred from ESX 2.x host down until ESX and VMFS upgraded
- Traditional Migration (ESX 3.0/VC 2.0)
 - Power off VM
 - Cold migrate to ESX 3.x host and VMFS-3 datastore
- Enhanced Migration (ESX 3.0.1/VC 2.0.1)
 - Live migrate (VMotion) running VM from ESX 2 host to ESX 3 host **with datastore relocation**
 - Requires VMotion-compatible hardware
 - **One-time, one-way operation**
- Vizioncore esxMigrator

Upgrade Detail – Phase II

■ VMFS-3 Upgrade

- VMFS-2 to VMFS-3 upgrade is non-destructive
- ~1.2 GB free space required for conversion
- Virtual machine swap files now stored with VMDK and VMX


■ ESX Upgrade

- Upgrade options
 - Boot from the ESX Server 3 installation media
 - Select the Upgrade option
 - Perform a clean installation
 - Download the tarball and upgrade similar to patching ESX 2.x hosts

- **DO NOT move to Phase III until all hosts have been upgraded to ESX 3.x and all VMDKs have been moved to VMFS-3 Datastores**

Upgrade Detail – Phase II

- Virtual machines become inaccessible to ESX 2.x hosts as partitions are upgraded to VMFS-3
- Keep this in mind when designing the migration path



Upgrading VMFS

■ Pre-Upgrade Preparation

- Back up files in VMFS datastore
- Resolve unsupported configurations (clustered VMs, raw disks)

■ Upgrade VMFS

- Put host in maintenance mode (requires all VMs be powered off)
- Upgrade VMFS datastore
- Take host out of maintenance mode

The screenshot shows the VMware ESX Server 3.0.0 interface. The 'Configuration' tab is selected, displaying a list of storage components:

Storage Component	Path	Capacity
san1 (1) (Readonly)	vmhba1:0:2:1	
volume1 (1) (Readonly)	vmhba0:0:0:6	
SharedVMs (Readonly)	vmhba1:0:25:1	99.99 GB
SharedTemplates (Readonly)	vmhba1:0:26:1	35.99 GB
SAN (Readonly)	vmhba1:0:31:1	29.99 GB

Two error and warning dialogs are present:

- Error Dialog:** "You must put the host into maintenance mode before upgrading this datastore." (OK button)
- Warning Dialog:** "Upgrading from VMFS-2 to VMFS-3 is a non-destructive process, but it cannot be undone. Any virtual machines using this file system must be powered off before beginning. Do you want to upgrade your file system now?" (Yes/No buttons)

The 'Upgrade to VMFS 3...' button is highlighted with a red circle. Below it, a pie chart shows the storage usage for 'volume1 (1)':


Category	Capacity
Used	4.00 GB
Free	54.63 GB
Total	58.63 GB

The 'Extents' table shows:

Extent	Capacity
vmhba0:0:0:6	58.63


Upgrade Detail – Phase II

- **Example:** SAN access from 4 ESX hosts to 4 SAN-based LUNs
- An upgrade would be difficult without more shared storage available



Upgrade Detail – Phase II

- It is often helpful to align the Hosts, VMDKs and LUNs in a manner that facilitates the upgrade process
- VMware refers to this as HVL alignment (Host/VM/LUN)



Upgrade Detail – Phase III

■ Preparation

- Phase II complete
- Licenses
- All datastores VMFS-3
- All ESX hosts v3

■ Expected Downtime

- Virtual Machines: each must be rebooted
- ESX hosts: None
- VC Server: None

■ Complete

- VirtualCenter 2.x
- ESX 3.x
- VMFS-3
- Virtual Machines

■ Remaining

- VMware Tools

■ Upgrade Virtual Machines

- Enables new features, expanded capacity
- Legacy virtual machines will run; cannot be modified
- Clone VMs prior to virtual hardware upgrade!
- **Once upgraded, ESX 2 host cannot power it on**
- Requirements
 - VMDK stored on VMFS-3 or VMkernel NFS datastore
 - No legacy suspend (REDO) files exist
- May be batched with `vmware-vmupgrade.exe`
 - Default operation also upgrades the VMware Tools (this may be disabled with `-s`)
 - Supports powered-off Windows 2000+ and Linux virtual machines

Upgrade Detail – Phase IV

■ Preparation

- Phase III complete
- Licenses
- All VMs upgraded
- Supported Guest OS installed

■ Expected Downtime

- Virtual Machines: must be rebooted once
- ESX hosts: None
- VC Server: None

■ Complete

- VirtualCenter 2.x
- ESX 3.x
- VMFS-3
- Virtual Machines
- VMware Tools

■ Remaining

■ Upgrade VMware Tools

- Install one VM at a time
- Windows tools are digitally signed; can be installed unattended
- Downtime is one reboot per virtual machine
 - Can be scheduled as needed: 1-2 minutes per VM + application shutdown time
- Required to leverage VCB file-level backups

Post-Upgrade Tasks & Notes

- Client-based firewall enabled by default
- Configure `ntpd`
 - Open firewall for `ntpClient`
 - `esxcfg-firewall -e ntpClient`
- If SMB/CIFS mounts are used for ISO/floppy images
 - Open firewall for `smbClient`
 - `esxcfg-firewall -e smbClient`
 - Create symlinks of mount points in `/vmimages`
 - `ln -s /mnt/mySMBmount /vmimages/mySMBmount`

Post-Upgrade Tasks & Notes

- VirtualCenter Server
 - Relocate Microsoft sysprep components
- Virtual Infrastructure Client
 - Can mount local drives (CD-ROM and floppy) and images to virtual machines
- VMkernel owns hardware (no more `vmkpcidivvy`!)
- Command line management tools (`esxcfg-*`) are available
- Bond devices no longer exist
 - Bonds automatically created by assigning multiple physical NICs to a vSwitch
- `vmkusage` is gone; use charts in Virtual Infrastructure client
- VirtualCenter Agent
 - `vmware-serverd` and `vmware-ccagent` are gone
 - `vmware-hostd` and `vmware-vpxa` are new

Post-Upgrade Tasks & Notes

■ SSH

- Clean installs: **root** is not able to login via SSH
 - Create local users on each ESX host
 - Use Vintella Authentication Services to authenticate against AD
- Default firewall configuration blocks SSH from ESX hosts
 - **esxcfg-firewall -e sshClient**

■ If ESX 2 Host Upgraded In-place

- Continues to allow **root** to login via SSH – lock it down
 - **/etc/ssh/sshd_config** ⇨ **PermitRootLogin no**
- **/etc/fstab** replaced with new, old is in **/etc/fstab.save**
- Networking and Storage: Verify and Clean Up as Needed
 - Service Console – check DNS and routing
 - NIC mapping and bonding, vSwitch definition
 - Storage configuration (Paths and Policy: MRU vs. Fixed)

New Features – Resource Pools

- Boundaries around groups of VMs to guarantee or restrict resources
- Only CPU and Memory resources
- May be subdivided
- Created at ESX host, Cluster, or Resource Pool level
- Examples
 - Chargeback
 - Charge departments for reserved CPU and Memory capacity
 - Priority/Criticality/SLAs
 - Configure Resource Pools based on performance needs
 - Prevent Test machines from affecting Production
 - “Security”
 - Shares assigned are relative to Resource Pool, not ESX host

New Features – VMware DRS

■ Leverages VMotion

- Create clusters based on VMotion groups

■ Start with Partially Automated mode

- Manual mode requires a decision for **each** power-on operation

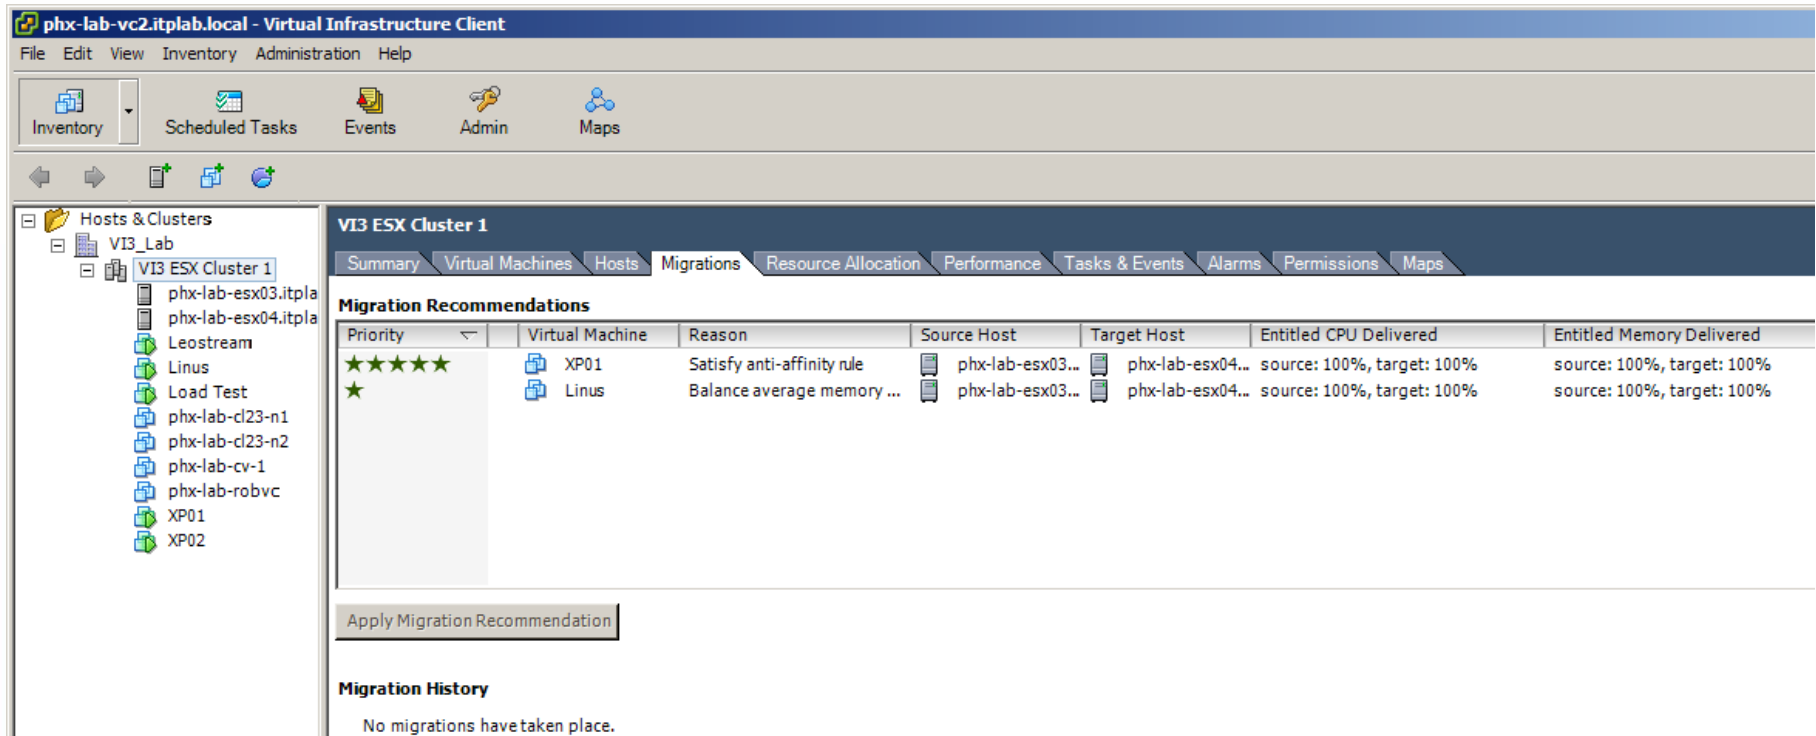
■ Balances load across the cluster; honors resource pool constraints

■ Cool Feature

- Enable Fully Automated DRS (Aggressive) and put a host into Maintenance Mode to have it automatically evacuate the host for maintenance

New Features – VMware DRS

- Recommendations assigned a 'star rating' (1-5 stars)
 - Displayed on the Migrations tab of the Cluster object



phx-lab-vc2.itplab.local - Virtual Infrastructure Client

File Edit View Inventory Administration Help

Inventory Scheduled Tasks Events Admin Maps

Hosts & Clusters

- VI3_Lab
 - VI3 ESX Cluster 1
 - phx-lab-esx03.itpla
 - phx-lab-esx04.itpla
 - Leostream
 - Linus
 - Load Test
 - phx-lab-cl23-n1
 - phx-lab-cl23-n2
 - phx-lab-cv-1
 - phx-lab-robvc
 - XP01
 - XP02

VI3 ESX Cluster 1

Summary Virtual Machines Hosts Migrations Resource Allocation Performance Tasks & Events Alarms Permissions Maps

Migration Recommendations

Priority	Virtual Machine	Reason	Source Host	Target Host	Entitled CPU Delivered	Entitled Memory Delivered
★★★★★	XP01	Satisfy anti-affinity rule	phx-lab-esx03...	phx-lab-esx04...	source: 100%, target: 100%	source: 100%, target: 100%
★	Linus	Balance average memory ...	phx-lab-esx03...	phx-lab-esx04...	source: 100%, target: 100%	source: 100%, target: 100%

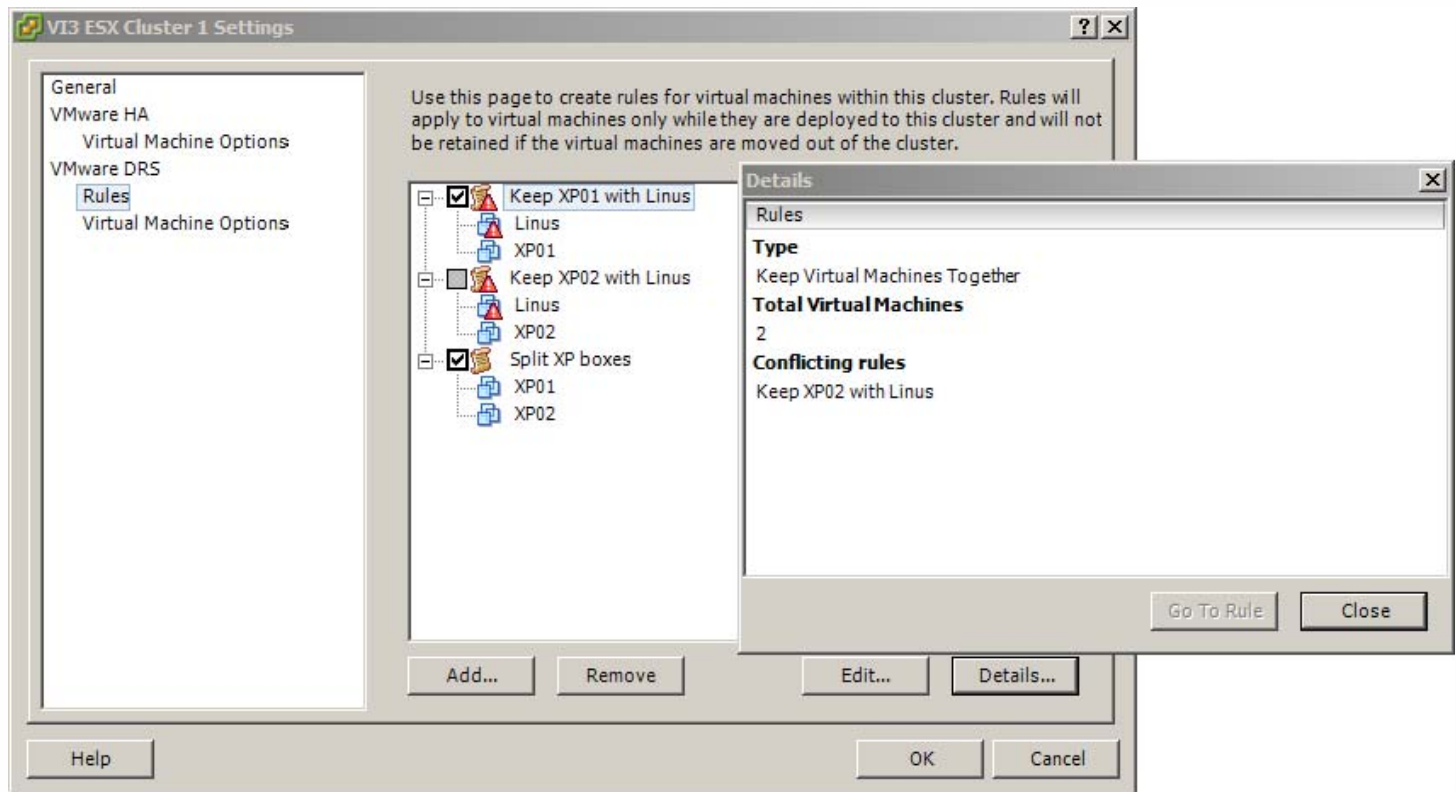
Apply Migration Recommendation

Migration History

No migrations have taken place.

New Features – VMware DRS

- Affinity/Anti-Affinity Rules
 - Virtual machines that should or should not be on the same ESX host
 - Logic issues (A+B, B-C, A-C) are flagged



New Features – VMware HA

■ What is it?

- If an ESX host fails, its hosted virtual machines are **powered up** on another host in the cluster
- Failures are detected by heartbeats on Service Console NICs
 - Failure to receive heartbeats triggers isolation detection
- Uses Legato AAM engine on Service Console

■ What isn't it?

- **Does not use VMotion**
- Virtual machines **become temporarily unavailable**
- **Does not** protect against failures within VMs
- **Does not** work between Clusters or Datacenters
- **Should not** be considered your DR solution

New Features – VMware HA

■ Requirements


- **Name resolution is required**
- ESX hosts must be able to resolve each other's names
- ESX hosts must be able to ping their default gateway(s)*
- `/etc/hosts` files may be used, but DNS is preferred

■ Notes from the field

- **Ensure that the Service Console has redundancy**
 - Prevents a switch failure from triggering Isolation Response (split-brain prevention) and powering off all virtual machines
- To reduce false positives
 - Enable portfast on physical switch ports
 - Enable “Rolling Failover” vSwitch policy
- Do not use with Beacon Probing network failure detection (BUG!)

New Features – VMware Consolidated Backup

- Virtual machines backed up from a centralized Windows 2003 Server (Backup Proxy) rather than from the ESX host
- Reduces load on ESX host and network
- Uses VMware Tools inside Guest OS to quiesce file system prior backup
- Single backup agent in a proxy server instead of per-VM agents
- File-level backups within Guest OS (Windows only)
- Full VM snapshot backup (Any Guest OS)



Q&A

- **Questions?**

- **Assistance**

- VI Upgrade Workshop

- This presentation highlights best practices from the detailed workshop

- http://www.vmware.com/pdf/vi_upgrade.pdf

- Contact VMware Professional Services or a VAC partner




Q&A

■ Questions?

■ Online Resources

- Planning an Upgrade to VMware Infrastructure 3
 - <http://www.vmware.com/vmtn/resources/539>
- VMware documentation:
 - <http://www.vmware.com/support/pubs/>
- Licensing
 - <http://www.vmware.com/download/licensing.html>
- Mike Laverick's (RTFM) Site:
 - <http://www.rtfm-ed.co.uk/>
- VMTN Forums
 - <http://www.vmware.com/community>



More Resources

- Vintella Authentication Services
 - http://www.quest.com/Vintella_Authentication_Services/
- Vizioncore esxMigrator
 - <http://www.vizioncore.com/esxMigrator.html>
- Microsoft Windows 2000 SP4 Update Rollup 1
 - <http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/rollup.mspx>

THANK YOU

Presentation Download

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

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

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

Some or all of the features in this document may be 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.

VMWORLD 2006



EXTRA SLIDES

Updating ESX 3.0

■ **esxupdate**: a new way to patch ESX hosts

> How?

- Download update package (3.0.1-32039-full.tgz)
- Copy to ESX server or host on NFS/SMB share and mount
- Migrate virtual machines; put host into Maintenance Mode

> Unpack a tarball and perform update with esxupdate

- `tar xfz 3.0.1-32039-full.tgz`
- `cd 32039`
- `esxupdate -n update`

> Tracks installed updates

- `esxupdate query`

Installed software bundles:

```
----- Name ----- --- Install Date --- --- Summary ---
      3.0.0-27701      14:54:27 07/28/06
      3.0.1-30353      09:43:50 08/29/06 Full 3.0.1 release of VMware ESX Server
      3.0.1-32039      10:25:54 10/06/06 Full 3.0.1 release of VMware ESX Server
```

Scripts and Snippets

- Location of VMDK files for all registered VMs on an ESX host

```
vmware-cmd -l | while read vm; do ID=`vmware-cmd "$vm" getconfig  
displayname`; ID=${ID/*= /}; DISK=`grep -i VMDK "$vm"`; DISK=${DISK/*= /};  
printf "$ID\t$DISK\n"; done
```

Storage Changes

- VMFS-3 supports directories and small file storage
- VI3 virtual machines store configuration files virtual disks, and per-machine swap files on VMFS-3
- Memory overcommitment relies on per-virtual machine swap files, not per-host swap files
- VI2 virtual machines with raw disks will fail on VI3 hosts
 - ESX Server 2.5.x VMs: convert a raw disk to a raw device mapping (RDM) prior to running on ESX Server 3.0
 - ESX Server 2.1.x VMs: remove a raw disk and add the disk back as a raw device mapping (RDM) when running on ESX Server 3.0