

BCT9552
VMware Infrastructure 3
Capabilities for Improving
Disaster Recovery

Jay Judkowitz
Sr. Product Manager, VMware

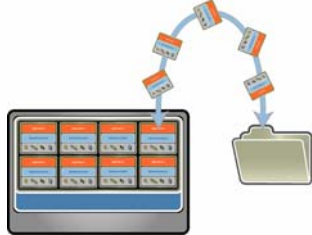



VMWORLD 2006

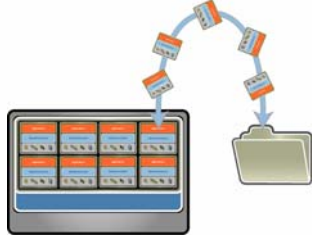
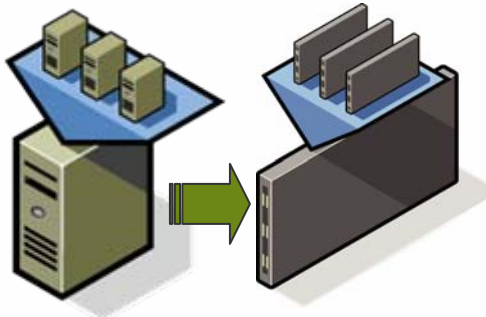
Agenda

- VMware and DR overview
- Scenarios and deployment options
 - Use cases by deployment type
 - Miscellaneous considerations

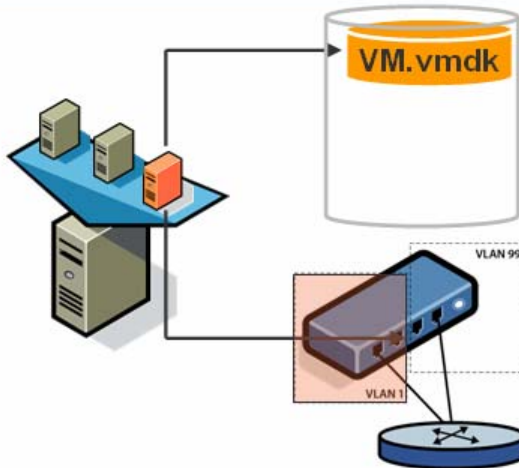
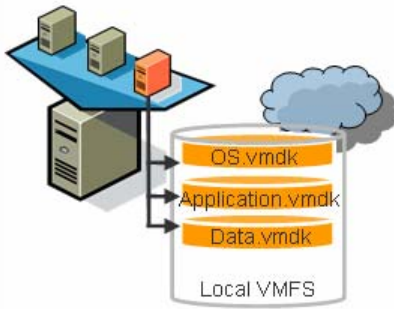
VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Encapsulation	<ul style="list-style-type: none">■ Enables simple replication of system state, applications, and data (VMFS3 now puts all necessary files for a VM in a single directory)	
Hardware Independence	<ul style="list-style-type: none">■ Allows waterfalling old equipment and/or buying new equipment■ Eliminates<ul style="list-style-type: none">➢ The need for identical HW at the primary and secondary sites➢ BMR process during recovery➢ System compatibility issues between the HW and OS exist at the DR site. Recovery is therefore much more reliable	

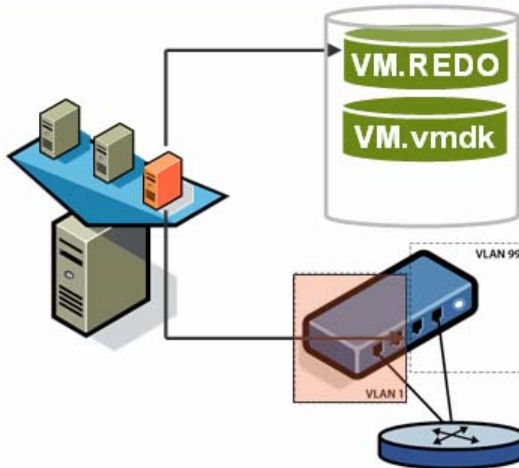
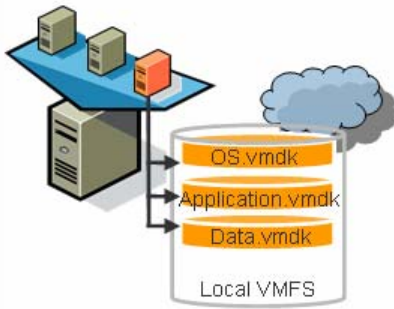
VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Encapsulation	<ul style="list-style-type: none">■ Enables simple replication of system state, applications, and data (VMFS3 now puts all necessary files for a VM in a single directory)	
Hardware Independence	<ul style="list-style-type: none">■ Allows waterfalling old equipment and/or buying new equipment■ Eliminates<ul style="list-style-type: none">➢ The need for identical HW at the primary and secondary sites➢ BMR process during recovery➢ System compatibility issues between the HW and OS exist at the DR site. Recovery is therefore much more reliable	

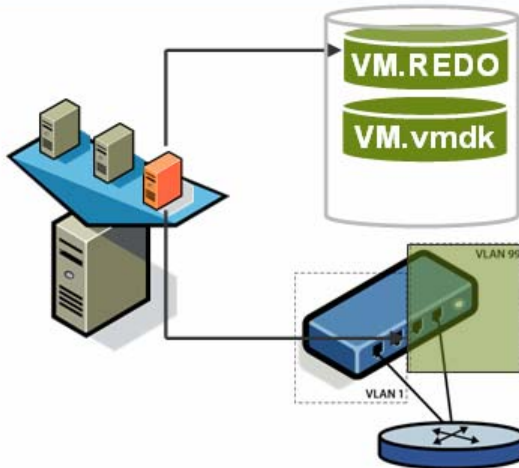
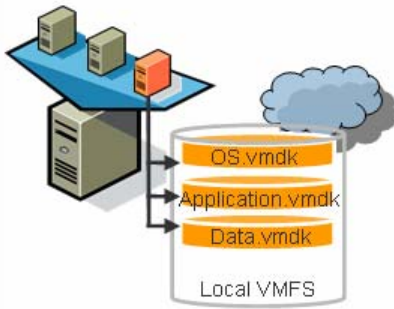
VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Snapshots and VLAN support	<ul style="list-style-type: none"> ■ Enables simple, reliable, and non-disruptive tests without extra hardware, OS, or application configuration 	 <p>The diagram illustrates a host with three VMs (represented by small server icons) connected to a network switch. A line connects one of the VMs to a storage unit labeled 'VM.vmdk'. The switch is shown with two VLANs: 'VLAN 1' and 'VLAN 99'. A router is also connected to the switch, indicating network connectivity.</p>
VMs Boot From Shared Storage	<ul style="list-style-type: none"> ■ Frees system disks from the host, making them available to array and network based replication tools, just like data ■ Eliminates the need to maintain OS and application patches and upgrades separately at source and destination sites ■ Provides the fastest recovery. Shut off replication, register the VMs and power them up 	 <p>The diagram shows a host with three VMs connected to a shared storage unit labeled 'Local VMFS'. The storage unit is depicted as a stack of three layers: 'OS.vmdk', 'Application.vmdk', and 'Data.vmdk'. A cloud icon is shown above the storage unit, representing network-based storage or replication.</p>

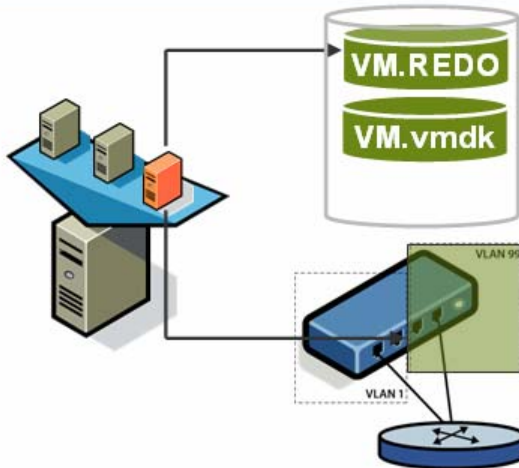
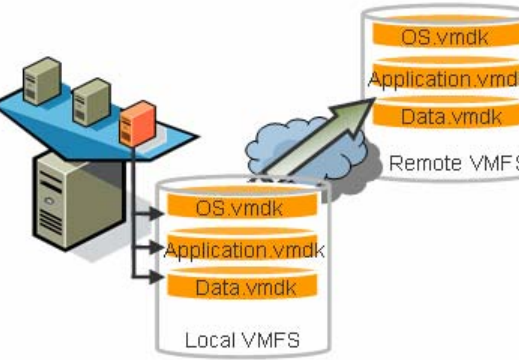
VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Snapshots and VLAN support	<ul style="list-style-type: none"> ■ Enables simple, reliable, and non-disruptive tests without extra hardware, OS, or application configuration 	 <p>The diagram illustrates a host connected to a storage array. The storage array contains two files: VM.REDO and VM.vmdk. The host is also connected to a network switch. The switch has two VLANs: VLAN 1 and VLAN 99. The switch is connected to a router.</p>
VMs Boot From Shared Storage	<ul style="list-style-type: none"> ■ Frees system disks from the host, making them available to array and network based replication tools, just like data ■ Eliminates the need to maintain OS and application patches and upgrades separately at source and destination sites ■ Provides the fastest recovery. Shut off replication, register the VMs and power them up 	 <p>The diagram illustrates a host connected to a storage array. The storage array contains three files: OS.vmdk, Application.vmdk, and Data.vmdk. The storage array is labeled Local VMFS.</p>

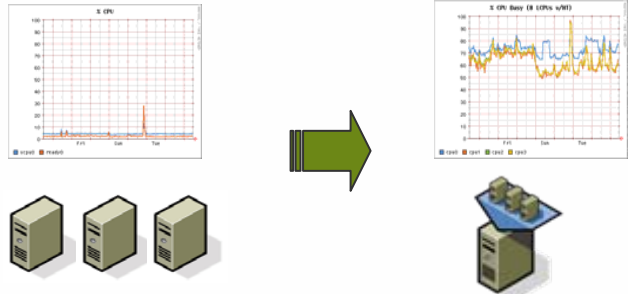
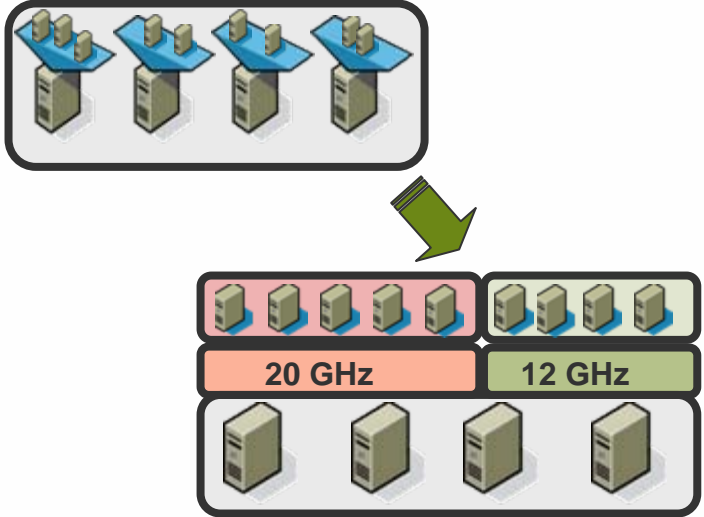
VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Snapshots and VLAN support	<ul style="list-style-type: none"> ■ Enables simple, reliable, and non-disruptive tests without extra hardware, OS, or application configuration 	 <p>The diagram illustrates a host environment connected to a storage array. The storage array contains two files: VM.REDO and VM.vmdk. The host is also connected to a network switch. The switch has two VLANs: VLAN 1 and VLAN 99. The switch is connected to a router.</p>
VMs Boot From Shared Storage	<ul style="list-style-type: none"> ■ Frees system disks from the host, making them available to array and network based replication tools, just like data ■ Eliminates the need to maintain OS and application patches and upgrades separately at source and destination sites ■ Provides the fastest recovery. Shut off replication, register the VMs and power them up 	 <p>The diagram illustrates a host environment connected to a storage array. The storage array contains three files: OS.vmdk, Application.vmdk, and Data.vmdk. The storage array is labeled Local VMFS.</p>

VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Snapshots and VLAN support	<ul style="list-style-type: none"> ■ Enables simple, reliable, and non-disruptive tests without extra hardware, OS, or application configuration 	 <p>The diagram illustrates a host (represented by a server rack) connected to a network switch. The switch is configured with two VLANs: VLAN 1 and VLAN 99. A VM disk is shown with two layers: VM.REDO and VM.vmdk. An arrow points from the host to the VM disk, indicating data flow or storage access.</p>
VMs Boot From Shared Storage	<ul style="list-style-type: none"> ■ Frees system disks from the host, making them available to array and network based replication tools, just like data ■ Eliminates the need to maintain OS and application patches and upgrades separately at source and destination sites ■ Provides the fastest recovery. Shut off replication, register the VMs and power them up 	 <p>The diagram illustrates a host (represented by a server rack) connected to a Local VMFS storage. The Local VMFS storage contains three layers: OS.vmdk, Application.vmdk, and Data.vmdk. An arrow points from the Local VMFS storage to a Remote VMFS storage, which also contains the same three layers: OS.vmdk, Application.vmdk, and Data.vmdk. The Remote VMFS storage is shown in a cloud-like environment, indicating network-based replication.</p>

VMs Are The Perfect Container From A DR Perspective

Feature	Benefit	
Consolidation	<ul style="list-style-type: none">Allows for smaller, less expensive DR sites	
Resource Pooling	<ul style="list-style-type: none">Eliminates need to preplan mapping between applications and hosts	

Why Do 63% Of Our Customers Use Us For DR?

■ Best RTO

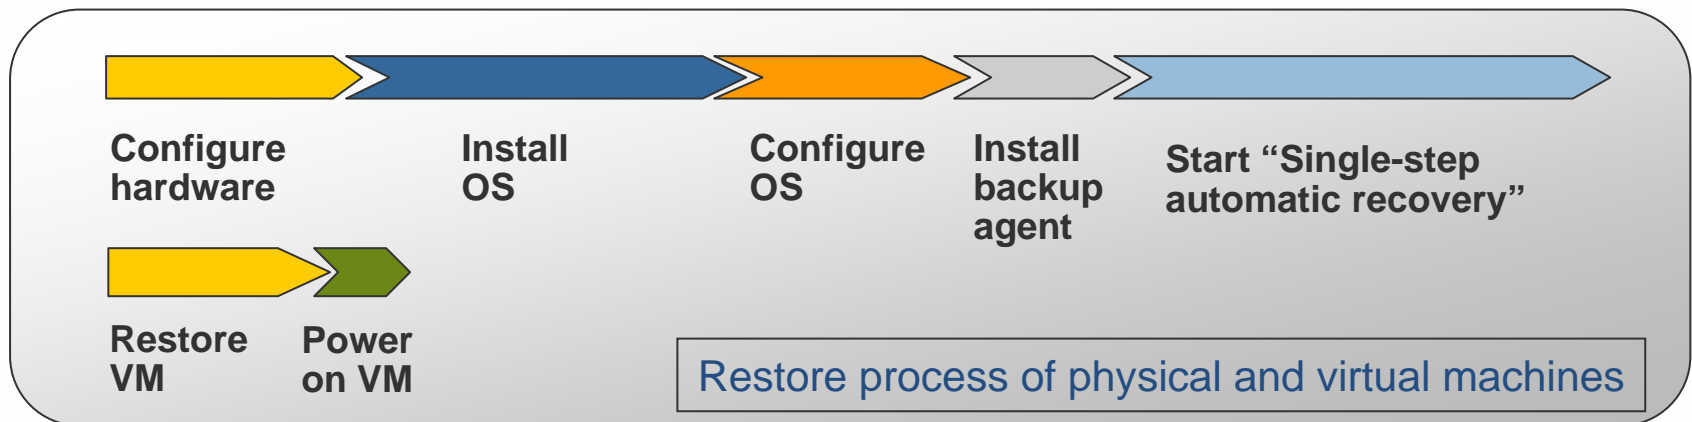
- Just replicate data however you'd like and boot the systems

■ Lowest cost

- Capital – Use hardware of choice at primary and secondary sites
- Operational – Simplicity of not worrying about the OS images

■ Most reliable recovery

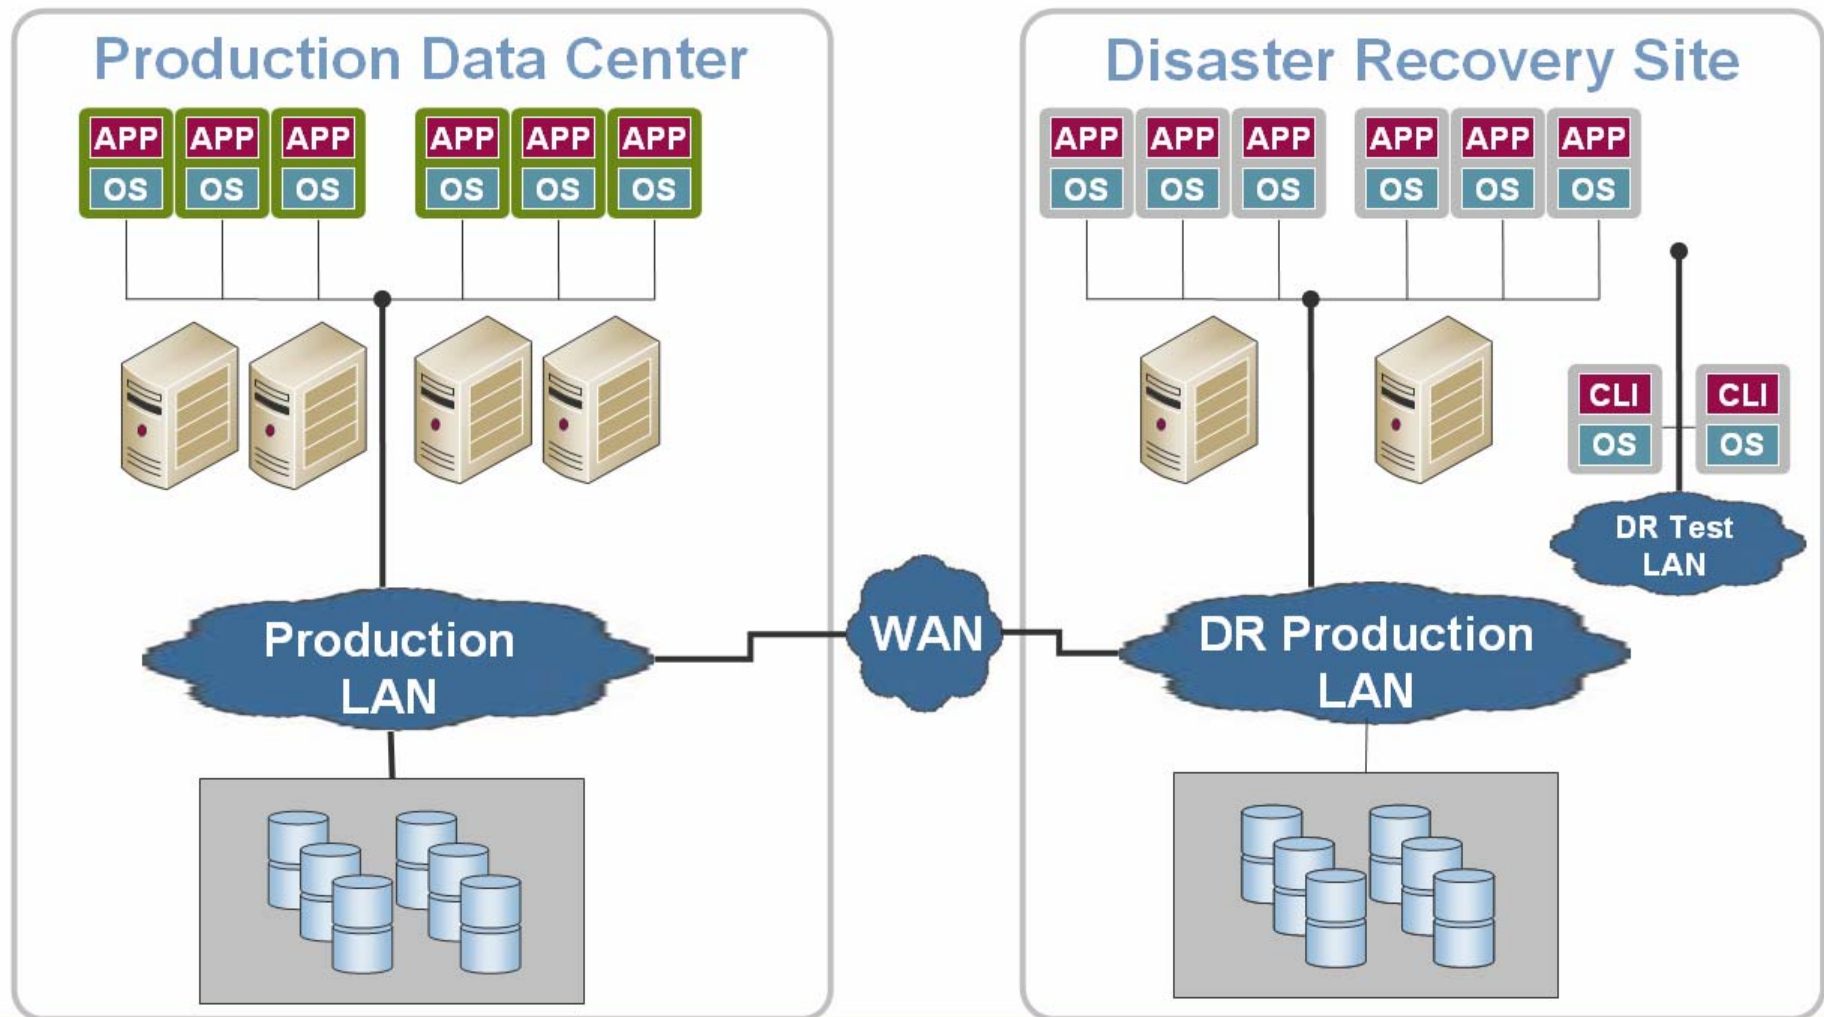
- HW dependencies as a source of failure is eliminated
- Regular testing becomes easier to perform



Agenda


- VMware and DR overview
- Scenarios and deployment options
 - Use cases by deployment type
 - Miscellaneous considerations

VM Failover – In Theory




VM Failover – In Theory

Production Data Center




Production
LAN




WAN

Disaster Recovery Site



DR Test
LAN

DR Production
LAN



VM Failover – In Theory

Production Data Center



Production LAN



WAN

Disaster Recovery Site



DR Test LAN

DR Production LAN



VM Failover – In Theory

Production Data Center



Production LAN



WAN

Disaster Recovery Site



DR Test LAN

DR Production LAN



Generalizations

	Enterprise Datacenters	Commercial Datacenters	Outsourced DR For SMB
DR Site	In-house, potentially across the country	In-house, potentially across the state	Outsourced to a regional DR hosting provider
Recovery Point (RPO)	Minutes	Hours	1 day
Recovery Time (RTO)	Minutes	1 day	1 week
Bandwidth To DR Site	Better than T3	T3	T1 or less
Storage Budget	High – Use all tier-1 storage	High – Use mostly tier-1 storage	Low – Using mostly tier-2 storage
Business Processes	Rigid processes, and controls	Some process, less restrictive controls	Everyone has root/administrator
Tolerance For Host Overhead	< 10%	~10%	~10%

Enterprise Data Centers

■ Replication

- Important workloads are most commonly protected by array-based and network-based continuous replication products
- Array and network vendors who support replication of ESX Server LUNs include:

- **EMC**
- **HDS**
- **HP**
- **IBM**
- **Network Appliance**

➤ Recommendations

- Use VMFS for most apps and physical compatibility RDMs for those apps that need app quiescing.
- Use consistency groups for multi-tier apps, either by putting the VMs on the same VMFS or by using RDMs in array consistency groups.
- Make sure volume resignaturing enabled for hosts at the remote site.

Enterprise Data Centers – Part 2

■ Workflow

- Many separate roles leveraging VC's fine grained access control
- Maybe some scheduling or workflow tool should be used to coordinate the startup of the VMs
 - Taking into account dependencies on physical systems
 - Taking into account dependencies on network and storage infrastructure
 - Taking into account dependencies between applications in the VMs
- Use DRS to simplify VM placement

■ Testing

- At least twice a year
- Think of having surprise tests
- Migrate production workloads instead of bringing up a parallel environment?

Commercial Data Centers

■ Replication

- Important workloads are most commonly protected by array-based, network-based, and host-based continuous replication products
- Replication providers that we work with today, on the host of VMware Server or in the guest OS on ESX Server, include:
 - **CA**
 - **Doubletake**
 - **EMC**
 - **SteelEye**
 - **Topio**
 - **VERITAS**
- Host based replication providers for ESX Server we work well with today include
 - **Vizioncore**
 - **PHD**
- Recommendations
 - No need for RDMs. Use VMFS
 - Not all OS based replication providers have consistency groups

Commercial Data Centers – Part 2

■ Workflow

- Generally there is an operations team that needs full VC/ESX access at the DR site
- Environment is small enough to not require 3rd party scheduling tools, but big enough that dependencies need tracking somewhere. Those dependencies lists need to be readily accessible in the remote site.
- Use DRS to simplify VM placement

■ Testing

- At least twice a year
- Think of having surprise tests
- Usually done with a test environment using test clients without actually migrating the production load

Outsourced DR For SMB

■ Replication

➤ Tape is most common transport for DR today. If you use tape, consider VCB. Current VCB integrations include:

- **CA Brightstor ArcServe**
- **CommVault Galaxy**
- **EMC NetWorker**
- **IBM Tivoli Storage Manager**
- **Symantec BackupExec**
- **VERITAS NetBackup**

➤ Recommendations

- Use image based backups to avoid BMR at remote site

Outsourced DR For SMB – Part 2

■ Scenarios

- Level 1 – you install ESX Server at DR site when you drive up
- Level 2 – Hosting provider installs ESX Server for you
- Level 3 – Hosting provider does your recovery for you end to end

■ Workflow

- Generally there is a operations team that needs full VC/ESX access at the DR site
- Usually you wait for servers unless you pay a premium. This prevents mass automation
- Workflow and recovery process should be negotiated with the outsourcer

■ Testing

- As contracted with the outsourcer
- Tests generally must be planned
- Tests may be partial to save costs with outsourcer

Outsourced DR For SMB – Part 3

- North American outsourcers we currently work with include:
 - > **Accenture**
 - > **Atos Origin**
 - > **Acxiom**
 - > **Capgemini**
 - > **CSC**
 - > **EDS**
 - > **IBM BCRS**
 - > **IBM Global Services**
 - > **Sungard**
 - > **Many smaller regional players.
Check with VARs/VACs**
- Encourage your outsourcer to have ESX Servers preinstalled as a service complete with VC. This will lower RTO at recovery time

Retail and Branch Offices

■ Considerations

- Value of the data depends on what is stored at the branch vs. what is stored centrally but accessed remotely
- Is the branch data valuable when/if the branch is destroyed?
- Based on the size of the branch, can you justify skilled IT personnel? Can you justify bandwidth? Balance b/w for replication vs. IT skills needed to manage tapes and backup SW

■ Recommendations

- Use host based replication if you can. (See vendors on slide 12)
 - Consider WAN accelerators and data deduplication technology such as:
 - **Avamar**
 - **Data Domain**
 - **Cisco**
 - **Riverbed**
- Use tape if you have to (See vendors on slide 14)

Agenda

- VMware and DR overview
- Scenarios and deployment options
 - Use cases by deployment type
 - Miscellaneous considerations

Application Considerations

	Enterprise Apps (Critical)	Multi-tier Apps (Non-Critical)	Infrastructure Applications
CPU and I/O load	<ul style="list-style-type: none">■ Ensure correct resource guarantees■ Ensure there are hosts in the cluster that can handle the SMP requirements	<ul style="list-style-type: none">■ Remember DRS affinity and anti-affinity rules	<ul style="list-style-type: none">■ Load is generally not an issue
Consistency groups	<ul style="list-style-type: none">■ Will be needed if multiple systems make up the app	<ul style="list-style-type: none">■ Definitely needed. Ensure the replication supports it	<ul style="list-style-type: none">■ Generally not needed
Special quiescing	<ul style="list-style-type: none">■ Almost always needed	<ul style="list-style-type: none">■ Needed usually for the DB tier only	<ul style="list-style-type: none">■ Generally not needed
Clustering	<ul style="list-style-type: none">■ Active/active geo-clustering can give 0 minute RPO and RTO■ Use active/passive clusters otherwise	<ul style="list-style-type: none">■ Front ends load balanced locally but need to be restarted in case of DR■ Perhaps use active/passive clusters for back end	<ul style="list-style-type: none">■ Usually inherently clustered (e.g. AD, DNS, LDAP)

Application Considerations – Part 2

	Enterprise Applications	Multi-tier Applications	Infrastructure Applications
Startup order	<ul style="list-style-type: none">■ Usually started after infrastructure apps but before other apps	<ul style="list-style-type: none">■ Usually needs a specific startup order inside the application	<ul style="list-style-type: none">■ Should usually be started first
Difficulty of changing IP and/or domain	<ul style="list-style-type: none">■ Check with ISV. This is very important	<ul style="list-style-type: none">■ Check with ISVs. This is very important	<ul style="list-style-type: none">■ The apps themselves may not have difficulty, but other applications need to be aware of the changes
Licensing issues	<ul style="list-style-type: none">■ Check with ISV about moving to and from hosts and/or VMs with different number of processors on source and destination sites	<ul style="list-style-type: none">■ Check with ISV about moving to and from hosts and/or VMs with different number of processors on source and destination sites	<ul style="list-style-type: none">■ Generally not an issue as long as the OS is taken care of


With in-house applications, anything can happen. Check with the application developer and make sure they leave great documentation before retiring. VMs eliminate HW interaction and legacy OS issues, but what else can go wrong?

Networking Considerations

- How should IPs be handled
 - Single IP range bridged across sites
 - VLAN transfer to DR site
 - DHCP for servers
 - Manual reconfiguration of IPs
- Think about DNS
 - Whatever choice you make has to be reflected in DNS

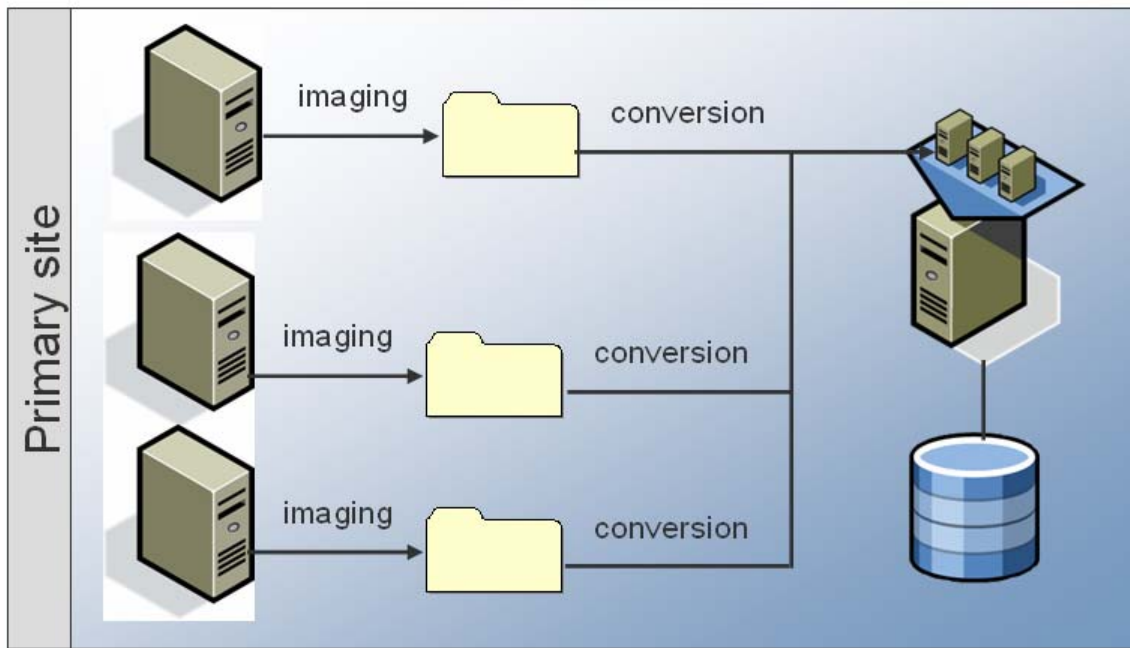
P2V DR

- There are two steps



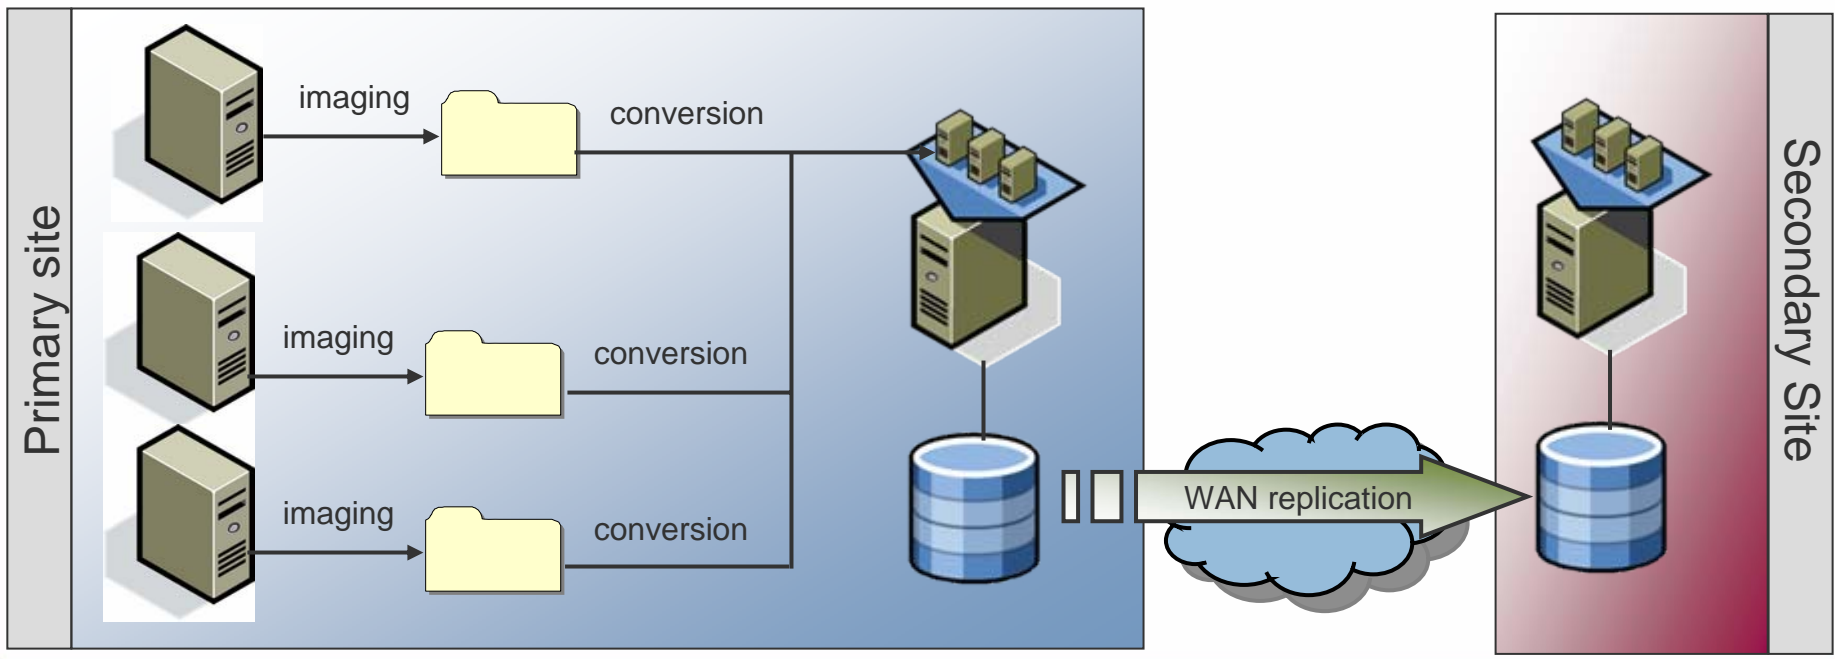
P2V DR

- There are two steps
 - Convert physical machines to VMs
 - Can use to protect against local server/storage failures



P2V DR

- There are two steps
 - Convert physical machines to VMs
 - Can use to protect against local server/storage failures
 - Use the local VMs as sources for V2VDR to remote site
 - Follow the same best practices in the last two sections



P2V DR – Part 2 – Workflow Options

- One time imaging and conversion
 - Good for legacy systems with little or no patching or upgrades
 - Offline imaging is best. Usually involves no software installation which would be risky and a single downtime is OK
- Regular imaging
 - Online imaging tool is generally required
 - Something with incremental image support is recommended
 - Can be run daily to handle all new patches and upgrades
 - Conversion can be done
 - Right after imaging – best RTO, but most work
 - At DR time – much less work, but much higher RTO

P2V DR – Part 3 – Partners

- Conversion and imaging vendors we work closely with include
 - > **VMware**
 - Look for our new P2V tool – VMware Converter. In beta now!
 - For now, VM Importer and P2V Assistant are great tools
 - > **Acronis**
 - > **HP**
 - > **Leostream**
 - > **Platespin**
 - > **Symantec**

Q&A

BCT9552 - VMware Infrastructure 3 Capabilities for Improving Disaster Recovery

Jay Judkowitz

Sr. Product Manager, VMware



VMWORLD 2006

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

