 STORAGE **Storage Decisions**

Virtual DR: Disaster Recovery Planning for VMware Virtualized Environments

VMware DR considerations

 STORAGE **Storage Decisions**

Virtual DR: Disaster Recovery Planning for VMware Virtualized Environments

VMware DR considerations

Raymond Lucchesi
President
Silverton Consulting, Inc.
Info@SilvertonConsulting.com
[Http://www.SilvertonConsulting.com](http://www.SilvertonConsulting.com)

© 2008 Silverton Consulting, Inc. 2 

 STORAGE **Storage Decisions**

Abstract


Virtual DR: Disaster Recovery Planning for VMware Virtualized Environments

VMware server virtualization is being touted as a near panacea for DR by making it easier than ever to recover servers to a secondary site. But, just as virtual servers complicate backup procedures, server virtualization has an impact on storage infrastructures and may require tweaking disaster recovery plans for storage systems.

STORAGE **Storage Decisions**

VMware DR advantages


- Testability
- Hardware independence
 - P2V, V2V, & V2P
- Data encapsulation

© 2008 Silverton Consulting, Inc. 4 

STORAGE **Storage Decisions**

DR testability


- VM DR can be tested on other ESX servers in same or remote site
 - Easy image copy availability
 - Easy configuration changes to run VM

© 2008 Silverton Consulting, Inc. 5 

STORAGE **Storage Decisions**

Hardware (HW) independence


- Primary <> DR site HW
- vmkernel isolates and virtualizes all CPU, networking and storage HW interaction
 - Except Raw Device Mapping (RDM)

© 2008 Silverton Consulting, Inc. 6 

STORAGE **Storage Decisions**

P2V, V2V, & V2P


- **P2V** - physical server hosted as a VM at DR site via VM converter and 3rd party tools
- **V2V** - DR site ESX server HW different than primary site HW
- **V2P** - VM hosted as a physical server at DR site, requires compatible HW, 3rd party tools

© 2008 Silverton Consulting, Inc. 7 

STORAGE **Storage Decisions**

Data encapsulation


- VMware encapsulates all VM data in few files under one directory
 - Except RDM data
- File data can be restored to re-start a VM on a remote site ESX server

© 2008 Silverton Consulting, Inc. 8 

STORAGE **Storage Decisions**

VMware Datastores

- VMware clustered file system (VMFS)
 - DAS
 - iSCSI
 - FC storage
- NAS/NFS volumes


© 2008 Silverton Consulting, Inc. 9 

VMFS capabilities

- Distributed, sharable file system, shared across ESX servers and VMs
- Flexible block sizing
- Flexible file system/volume sizing
- File system change journal


VMFS datastore

- VMware cluster file system



Files for each VM (.vmx, .vmdk & others)

- File(s) that encapsulate the config, O/S, application, and data for a Virtual Machine (VM)



Raw Device Mapping (RDM)

Two modes

- Virtual compatibility mode
 - uses VMFS mapping file
 - virtualizes physical device I/O
- Physical compatibility mode,
 - uses VMFS mapping file
 - I/O directly to physical device bypassing ESX I/O virtualization
 - VM snapshots not supported

Recap VMware

- Provides easy DR testability
- Removes H/W dependencies
- Encapsulates all VM data
 - RDM data lone exception

VMware backup alternatives


- Backup SW agents on VMs
 - File level backup & restores
- Backup SW agents on ESX service console
 - Image level backup & restores
- VMware consolidated backup
- CDP

STORAGE **Storage Decisions**

VM backup agents

- Backup agents on VM read files
- VM backup agent transfers file data to backup server over LAN
- Backup server writes data to backup target

→ No .vmdk or .vmx images for DR
 → Ability to do file level restores
 → Performance considerations


© 2008 Silverton Consulting, Inc. 16 

STORAGE **Storage Decisions**

ESX service console backup agents

- Backup agents on ESX service console reads .vmdk and .vmx files
- ESX backup agent transfers file data to backup server over LAN
- Backup server writes data to backup target


→ No file level restore
 → Performance considerations

© 2008 Silverton Consulting, Inc. 17 

STORAGE **Storage Decisions**

VM consolidated backup (VCB)


- "LAN-free" backup of VMs
 - Suspend VM
 - Uses VM snapshot to replicate datastores
 - Catalog's VM state on snapshot
 - Resume VM
 - Snapshots mounted or streamed to VCB proxy
 - Snapshots backed up to target media via other backup SW
 - VCB releases snapshots after backup
- Requires
 - VCB backup proxy (Windows 2003) server
 - Other Backup SW
 - Shared SAN access to VMFS and NAS datastores and snapshots
- VI3 supports all datastore types for VCB

© 2008 Silverton Consulting, Inc. 18 

STORAGE **Storage Decisions**

VMware Snapshot

- Copy of disk, memory, CPU state taken point-in-time
- Parent-child relationship
- .REDO bitmap files
- Activity state of VM?


© 2008 Silverton Consulting, Inc. 19 

STORAGE **Storage Decisions**

VCB

Works well for

- Offloading ESX server backup cycles
- Image copies of VMs
- File level backups for Windows


© 2008 Silverton Consulting, Inc. 20 

STORAGE **Storage Decisions**

VCB

Works poorly for

- VMs with large .vmdk - need file level backups
- OLTP or always on VMs
- RDM physical mode


© 2008 Silverton Consulting, Inc. 21 

 STORAGE **Storage Decisions**

VCB integrated S/W Tools


- VizionCore VrangerPRO
- esXpress
- Vmts.net - vmbk.pl
- Symantec Backup Exec system recovery option
- Also supported by Symantec Backup Exec, Net Backup, EMC Networker, CommVault Galaxy and others


© 2008 Silverton Consulting, Inc. 22 

 STORAGE **Storage Decisions**

VMware backup recap


- How you backup VMs impacts DR
- VCB can help DR
 - For the right VMs

© 2008 Silverton Consulting, Inc. 23 

 STORAGE **Storage Decisions**

VMware CDP

- VM based write splitters protect VM files
- Network and storage array write splitters protect disk images
- Some CDPs support ESX write splitters via Veritas Volume manager also protecting

© 2008 Silverton Consulting, Inc. 24 

VMware local clustering

- DRS for performance optimization
- HA for fault tolerance
- Both depend on
 - VMotion to migrate active VMs
 - Shared access to datastores

© 2008 Silverton Consulting, Inc. 25

VM Dynamic Resource Scheduler (DRS)

- User defined pools of resources and VM prioritization
- Automatic or manually balances VM load across defined resource pools
- Also used for service outages
- Also supports power management

© 2008 Silverton Consulting, Inc. 26

VM High Availability (HA)

- Local fault tolerant cluster of ESX servers using heartbeat to detect failed VM/ESX server
- Reserves resources for fail over
- In combination with DRS selects optimal placement for restart

© 2008 Silverton Consulting, Inc. 27

STORAGE **Storage Decisions**

VMotion

- Running VM quiesced, snapped and terminated
- Activates new VM on another ESX sever
- Requires
 - Shared access to datastores
 - Compatible H/W
- VM data stays in place

© 2008 Silverton Consulting, Inc. 28

STORAGE **Storage Decisions**

Storage VMotion

- Moves .vmdk files to other datastores
- Used primarily for technology upgrades
- Also used for balancing storage access

© 2008 Silverton Consulting, Inc. 29

STORAGE **Storage Decisions**

Storage VMotion

Local data migration

- Allows live VM directory migration to other VMFS datastores
- Multi-step process that clones VM directory and then re-directs all writes to new datastore
- Uses shared access to all VMFS datastores

© 2008 Silverton Consulting, Inc. 30

Why DRS, HA and VMotion for DR

- Can operate both at local and remote site to optimize performance
- Resource pools and VM prioritization also needed for DR

VMware DR alternatives

- Non-RDM DR
- RDM DR
- VMware Site Recovery Manager
- Non-VMware replication-clustering

Non-RDM VM DR

VM data is .vmx, .vmdk(s), & other files

- Image copies available
- H/W & S/W to run ESX server
- Configure ESX server to run VM
- Configure datastore(s)
- Restore VM files
- Re-IP network
- Start ESX server
- Restart VM

RDM VM DR

- Same as Non-RDM VM DR but
- RDM mapping files in VMFS needed
 - Physical copy of RDM LUN(s) required to be backed up and restored at DR site before VM restart
 - Storage HW configuration matches primary site



VMware Site Recovery Manager (SRM)

- Maps LUNs to VMFS to (.vmdk & .vmx) files
- Automates
 - Procedures to invoke SAN-LUN replication
 - Processes to failover to hotsite
 - Procedures to re-IP networking at hotsite
- Storage supported replication agents
 - 3PAR, Compellent, Dell, EMC, FalconStor, Hitachi, HP, IBM, LeftHand Networks, NetApp, and Xiotech



SRM Storage Replication Agents (RA)


- RA is VMware defined API
 - Provides consistent SRM services across storage vendors
 - To initiate, monitor, and terminate storage replication
- Can support asynch and/or synchronous disk mirroring
- Can take advantage of storage snapshots



STORAGE **Storage Decisions**

Non-VMware replicator-clustering products

- VizionCore Vreplicator
 - Uses software replication
- Double-take Server Recovery Option
 - uses software replication
- Veritas Cluster Services for VMware
 - uses storage HW replication


© 2008 Silverton Consulting, Inc. 37 

STORAGE **Storage Decisions**

Other Cluster software

Operate at VM level and only at Local site


- MSCS
- Livelink SUSE
- WanSyncHA
- Veritas Cluster Services for VMware also support local clustering

© 2008 Silverton Consulting, Inc. 38 

STORAGE **Storage Decisions**

Remote data replication alternatives

- SAN data mirroring
- Storage Virtualization
- Replication appliances
- Remote CDP
- Software replication

© 2008 Silverton Consulting, Inc. 39 

SAN mirroring considerations

- Dedicated networking
- Active storage at remote DR site
- Types of SAN mirroring
 - Synch
 - Semi-synch
 - Asynch
- Requires matching vendor storage HW

VMware SAN mirroring considerations

- Insure all LUNs for VMFS datastores are single consistency group
- Datastore mirrored LUNs/volumes must contain all .vmdk and .vmx files for VM needed for DR
- RDM LUNs also need to be mirrored
 - Should be included in consistency group

Storage virtualization


Storage HW independence

- Heterogeneous data mirroring
 - Primary <-> DR site storage hardware
- Also provides storage tiering and pooling across subsystems

STORAGE **Storage Decisions**

Replication appliances

- Usually network attached dedicated HW providing remote replication
- Can use TCP/IP over WAN for data transfer
- Support for Asynch mirroring
- VMware considerations same as SAN mirroring


© 2008 Silverton Consulting, Inc. 43 

STORAGE **Storage Decisions**

CDP replication appliances

Similar to replication appliances but


- Provide any point-in-time recovery
- Requires write-splitter

© 2008 Silverton Consulting, Inc. 44 

STORAGE **Storage Decisions**

Software Replication

- Double-take replication software for Windows
- Softek Replicator
- VizionCore Vreplicator


© 2008 Silverton Consulting, Inc. 45 

STORAGE **Storage Decisions**

RTO: 8d..?

Coldsite DR considerations

- Backup data offsite
 - Could be image or file data
- H/W contracts for guaranteed delivery timeframes
- S/W agreements for service and licensing in case of disaster


© 2008 Silverton Consulting, Inc. 46 

STORAGE **Storage Decisions**

RTO: 1d..8d

Warmsite DR considerations

- Owned DR site or Outsourced DR site
 - For outsourced - SunGard, IBM, HP, etc.
- Data can be at 3rd site or located at DR site
 - Outsourced sites can be multiple locations


© 2008 Silverton Consulting, Inc. 47 

STORAGE **Storage Decisions**

RTO: 3h..1d

Hot site DR considerations

- Asynch replication via SW, appliance or storage subsystem
 - Need to replicate datastores and RDM data
 - S/W support - Vreplicator & Double-Take SRO
 - No support for RDM data
- Dedicated networking
- Active data replication at DR site

© 2008 Silverton Consulting, Inc. 48 

RTO:0h..3h

Mirror site DR considerations

- Dedicated SAN or appliance based replication
 - Need to replicate datastores and RDM data
- Dedicated networking
- Active storage at DR site



For more information

Ray Lucchesi
 +1-720-221-7270
 Info@SilvertonConsulting.com



Some terminology

vm disk	VM disk file
vmx	VM configuration file
CDP	Continuous data protection
Datastore	Storage that holds VM disk, configuration, and other files
DRS	Dynamic Resource Scheduler
ESX	VMware enterprise server virtualization
HA	High Availability
Orphaned Server	Non-running VM, disk files exist, but no ESX server runs VM
RDM	Raw device mapping
SAN	Storage area network
Storage Vmotion	VM directory live migration from one VMFS datastore to another
VCB	VMware consolidated backup
VI 3	Virtual Infrastructure 3
VM	Virtual machine
VMFS	VMware file system
vmkernel	VMware hyper-visor which virtualizes hardware
Vmotion	VM live migration from one ESX server to another