



Education

## **SERVER VIRTUALIZATION AND STORAGE DISASTER RECOVERY**

Ray Lucchesi, Silverton Consulting  
[Ray@SilvertonConsulting.com](mailto:Ray@SilvertonConsulting.com)

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - ◆ Any slide or slides used must be reproduced without modification
  - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
- The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

**NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

- **Server Virtualization (SV) and Storage Disaster Recovery (DR)**
  - ◆ Server virtualization demands increasingly sophisticated storage subsystems. To support disaster recovery (DR) in virtual server environments storage admins need to understand how I/O and storage layouts in these environments differ from non-virtualized server storage. In addition, some server virtualization systems in order to automate DR can make use of special storage subsystem capabilities unique to these environments.

# Server Virt. (SV) DR advantages

- Testability
- Hardware independence
  - ◆ P2V, V2V, & V2P
- Performance optimization
- Data encapsulation

# DR testability

- DR can be tested at VM level at local site on other SV servers
- DR can be tested at VM level at remote site with data replicas
- Virtual machine (VM) DR can be easily tested on other servers
  - ◆ Image copies needed
  - ◆ Configuration changes needed to run VM

# Hardware (HW) independence

- Primary <> DR site HW
- SV hypervisor/kernel (HV/K) isolates and virtualizes all CPU, networking and storage HW interaction
  - ◆ Except physical/raw device mode (PDM)
- DR site can use re-purposed HW

# P2V, V2V, & V2P

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

# Performance optimization

- VMs can easily be migrated to other SV servers at DR site to optimize performance or to balance resource utilization
- SV VM migration services supported
- Data migration services may also be used



# Data encapsulation

- SV encapsulates all VM data in few files under one directory
  - ◆ Except PDM data
- File data can be restored to re-start a VM on a remote site virtualized server

# SV Data Mapping


- Most but not all SV products support clustered file systems for data storage (CFS) using
  - ◆ iSCSI
  - ◆ FC storage
- Can also use NAS/NFS volumes
- Some SV products only need shared LUN access
- Most SV products also support physical disk access bypassing IO virtualization

# SV CFS capabilities

- Most CFS's support distributed,, high performance, scaleable, sharable file system
- Some CFS's support flexible block sizing, flexible file system/volume sizing, snapshots, and file system change journals

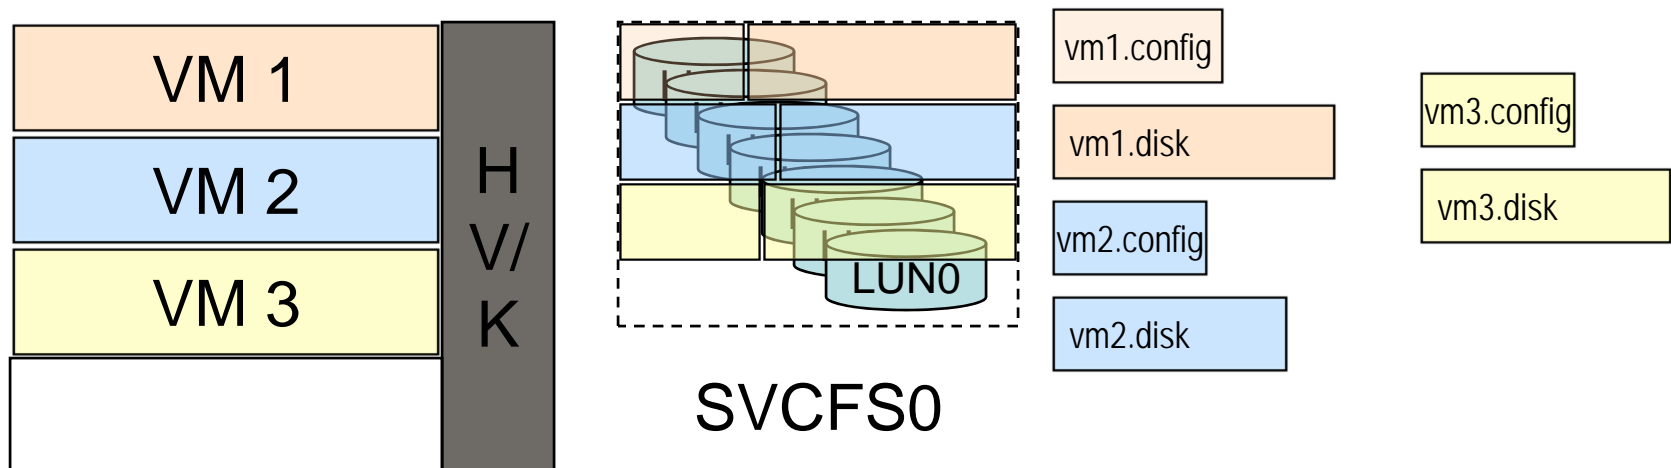
# SV CFS data stores

## ➤ SV CFS data stores




# Files for each VM

- File(s) that encapsulate the config, O/S, application, and data for a VM
  - ◆ Typically a .config file and at least a .disk file that corresponds to all physical disks needed by the VM



# Shared LUN data stores

➤ Shared LUNs are partitioned to support VMs



# Physical Device Mapping (PDM)

## Two modes

### ➤ Virtual compatibility mode

- ◆ Simulates physical device I/O

### ➤ Physical access mode,

- ◆ I/O directly to physical device bypassing HV/K I/O virtualization
- ◆ SV CFS level snapshots not supported

# Recap SV DR advantages

- Provides easy DR testability
- Removes H/W dependencies at DR site
- Performance optimization at DR site
- Encapsulates all VM data
  - ◆ PDM data lone exception



# SV backup alternatives

- Backup SW agents on VMs
  - ◆ File level backup & restores
- Backup SW agents at HV/K service manager (HV/K-SM)
  - ◆ Disk image level backup & restores
- Specialized backup services
- Continuous data protection (CDP)

# VM backup agents

- No change to current backup procedures
- 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 .disk or .config images for DR
  - Ability to do file level restores
  - Performance considerations

- Backup agents on HV/K-SM reads .disk and .config files
- HV/K-SM backup agent transfers file data to backup server over LAN
- Backup server writes data to backup target
- Backups all VMs on SV server
- No file level restore
- Performance considerations

# Specialized backup

- “LAN-free” backup of VMs
  - ◆ Suspend VM
  - ◆ Uses snapshots to replicate data stores
  - ◆ Catalog’s VM state on snapshot
  - ◆ Resume VM
  - ◆ Snapshots mounted or streamed to backup server
  - ◆ Snapshots backed up to target media via other backup SW
  - ◆ Releases snapshots after backup
- Requires shared access to data stores

# SV-CFS level snapshots

- Copy of disk, memory, CPU state taken point-in-time
- Parent-child relationship
- Change journal logs
- Activity state of VM?

# Specialized backup services

Works well for

- Offloading HV/K-SM backup cycles
- Image copies of VMs
- File level backups for some operating systems

# Specialized backup services

Works poorly for

- VMs with large .disk files - need file level backups
- OLTP or always on VMs
- PDM storage

- VM based write splitters protect VM files
- Network and storage array write splitters protect disk images
- Some CDPs support HV/K write splitters
- Depending on splitter location can have file or disk level rollback



# SV backup recap

- How you backup VMs impacts DR
- Specialized backup services can help DR
  - ◆ For the right VMs
- Disk image files similar to bare metal restore files

# SV local clustering

- For performance optimization
- For fault tolerance/high availability (FT/HA)
- Both depend on
  - ◆ Active VM migration to migrate VMs between SV servers
  - ◆ Shared access to datastores

# SV performance clustering


- 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

# SV FT/HA clustering

- Local fault tolerant cluster of SV servers using heartbeat to detect failed server
- Reserves resources for fail over
- In combination with performance clustering selects optimal placement for restart

# Active VM migration

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



# Why clustering and migration

DR requires similar processes

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

# SV DR alternatives

- Non-PDM DR
- PDM DR
- Specialized DR software/services

# Non-PDM VM DR

VM data is .config, .disk(s), & other files

- Disk image copies available at remote site
- H/W & S/W to run SV server
- Configure SV server to run VM
- Configure data store(s)
- Restore VM files
- Re-IP network
- Start SV server
- Restart VM



Same as Non-PDM VM DR but

- PDM configuration files needed
- Physical copy of PDM LUN(s) required to be backed up and restored at DR site before VM restart
- Storage HW configuration matches primary site

# Specialized DR services 1

- Maps LUNs to data stores to (.disk & .config) files
- Can support HW asynch and/or synchronous disk mirroring
- Can take advantage of HW storage snapshots
- Automates
  - ◆ Procedures to invoke SAN/IPSAN-LUN replication
  - ◆ Processes to failover to hotsite
  - ◆ Procedures to re-IP networking at hotsite

# Specialized DR services II

- Software packages that support geographically distributed clustering
- Software packages that support disk data replication using WAN
- Specific to SV solution

# Other cluster software

Operate at VM level and only at local/single site

- Specific to VM O/S
- Available for SVs and VMs at one site alone
- Features specific to each product and each SV solution

# Remote data replication

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

# SAN and IPSAN mirroring

- Dedicated networking for SAN, WAN used for IPSAN
- Active storage at remote DR site
- Types of SAN mirroring
  - ◆ Synch
  - ◆ Semi-synch
  - ◆ Asynch
- Requires matching vendor storage HW

# SV SAN and IPSAN mirroring

- Insure all LUNs for data stores are single consistency group
- Data store mirrored LUNs/volumes must contain all .disk and .config files for VM needed for DR
- PDM 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



# Replication appliances

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

# CDP replication appliances

- Similar to replication appliances but
- Provide any point-in-time recovery
  - Requires write-splitter at VM, at HV/K, network, or storage HW

# Software Replication

- Often specific to VM O/S being used
- Uses WAN for replicating data
- Typically tied in at the HV/K level
- Can be tied to software clustering products specific to each SV solution

- Please send any questions or comments on this presentation to SNIA: [trackvirtualization@snia.org](mailto:trackvirtualization@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

**- SNIA Education Committee**

**Ray Lucchesi**