# Developing a disaster recovery plan for virtual machines: A tutorial

06 Apr 2009 | SearchDisasterRecovery.com

Enterprise IT tips and expert advice

*By Ray Lucchesi*

Virtual machine (VM) disaster recovery (DR) is a multifaceted activity that fails over a VM from a primary site to a remote location. There are a few approaches to facilitating disaster recovery in a virtual machine environment: VMware Inc. introduced its vCenter Site Recovery Manager (SRM) software that automates virtual machine failover.

Alternatively, there are geographically disbursed clustering (geoclustering) services that support automatic failover, but can also recover more than just VMs. There are also standard data protection packages available that support varying levels of VM DR. While these packages are more manual than Site Recovery Manager or geoclustering, they cost substantially less.

Learn about look at the various approaches to virtual machine disaster recovery in this tutorial.

## VMware Site Recovery Manager

Facilitating recovery with VMware Site Recovery Manager automation depends heavily on array or storage area network (SAN) replication to copy datastore data between sites. SRM software executes on a SRM server or virtual machine at both the protected and DR sites, but also requires a Virtual Center to run at the remote site.

Once Site Recovery Manager is executed, an administrator should:

- Establish datastore replication
- Identify replicated datastores
- Select protected virtual machines
- Remap VM hardware
- Create a data recovery plan

Re-IP networking refers to the fact that the IP addresses at the remote site can't be the same as the primary site. Some of these are associated with the application and operating system running in the virtual machine and some are associated with VMware Hypervisor interfaces like the server running vCenter Server, Site Recovery Manager, etc. As the VMs are brought up at the remote site, the IP addresses must be changed in order to run.

Moreover, multiple recovery plans can be defined and administrators may select which one to use for a specific failover. Alternative recovery plans such as these provide varying failover capabilities and supply recovery options for partial failures, e.g., a single datastore or ESX host failure at the protected site.

As such, SRM requires at least one manual step. In addition, SRM also

supports DR testing at the local site and an administrator may modify an already existent recovery plan to support this testing.

The nice thing about VMware Site Recovery Manager is that you can have as many or as few recovery plans as you need. It's entirely conceivable that one would have a recovery plan for a total site failure and one or more for separate infrastructure failures.

VMware High Availability (HA) provides for ESX failover, but only to the local site. SRM is only involved when you want to failover to a remote site. Not every infrastructure failure would warrant a "disaster" being invoked, which would require SRM automated failover to a remote site.

VMware SRM currently has some limitations, including:

Supporting Raw Device Mode data
Supporting multi-LUN datastores
Supporting multi-site DR automation
Supporting automated failback

VMs can access Fibre Channel storage in at least two ways. The first way is through normal VM hypervisor SCSI data access, which is virtualized to a VMware-defined VM cluster file system (VMFS datastore). The second is through Raw Device Mode, whereby the VM actually owns the Fibre Channel port hardware and controls that link, and likely the storage attached at the other end of the link.

What non-support for Raw Device Mode data means is that failover for virtual machines that have this data are more complex and less automated. SRM will not monitor replication of this data and will not automatically promote this data to active VM accessibility on failover. All of these steps have to be done manually or via data center scripting.

Raw Device Mode is normally used by performance-intensive virtual machines. These are typically high-profile applications, but are least likely to be virtualized. However, due to their criticality, they are very likely to warrant the highest form of disaster recovery.

Whether this is a concern for system admins/data centers depends on how much of their infrastructure and servers are virtualized. As more data centers move to 100% virtual machines, this will become more of a concern.

As a side note, VMware does supply support for Raw Device Mode in a beta version of SRM. Failback can still be accomplished, but an administrator would need to reconfigure SRM to perform the failback as an SRM failover.

While failover is typically unscheduled, failback is typically a scheduled activity, once you have failed over. It takes time to bring the primary site back online, repair the infrastructure and power up the data center. These time consuming activities can be scheduled to occur, so you would also be able to schedule the failback process.

It's possible that the recovery plan for failback could be in place beforehand, but Site Recovery Manager interrogates storage replication activity to validate that a protected datastore is being replicated. So a failback process

identifying protected datastores and VMs and remapping inventory steps for SRM, might have to wait until the failover actually takes place before it starts, particularly when:

Reestablishing datastore replication
Re-identifying protected datastores
Re-selecting protected VMs
Remapping site inventories
Creating a failback recovery plan

## Geoclustering for virtual machine disaster recovery

Many geoclustering products are available that provide even more sophisticated cross-site recovery. In fact, geoclustering can support automated failback and failover, multi-destination DR sites and raw device mode data, and may not require a Virtual Center.

Symantec Corp. Veritas Cluster Services (VCS) allows for physical server to VM, VM to physical server and VM to VM failover.

For instance, VCS can failover a physical server at the protected site to a VM at the remote site, or vice a versa. Such capabilities go well beyond what VMware SRM was intended to support, but depending on data center needs, may be worthy of consideration. Also, VCS executes at the ESX service console level when supporting VM failover.

Windows HPC Server 2008 is another geoclustering product, but only supports server to server or VM to VM failover. As such, HPC Server must be executing in the Windows server at both the local and remote site, and only supports Windows-to-Windows failover.

## SAN or array replication for VM disaster recovery

Most failover automation depends heavily on SAN or array replication, and with this in place, automating failover can be accomplished with any number of approaches.

Once datastore replication is in place, administrators can build their own scripts using native VMware or other software to semi-automate virtual machine failover. However, this custom scripting must do all the work required to reconfigure the ESX servers to run the VMs, re-IP the VMs and promote replicated datastore copies.

## Data protection software for virtual machine disaster recovery

Data protection packages such as EMC Corp. NetWorker, CommVault Galaxy, IBM Corp. Tivoli Storage Manager (TSM) and Symantec BackupExec and NetBackup all support DR at varying levels. Such support may consist of bare-metal restore options and/or sophisticated independent backup data replication.

Tivoli Storage Manager supports a disaster recovery manager option that can be used to automatically replicate TSM protected data to a remote site. Once TSM is recovered at the remote site, data can be restored and VMs can be reconfigured via manual operator activity or hand scripted automation.

Alternatively, other backup packages support a bare-metal restore option. Such functionality can provide a one step, restore-able version of all the data required by a server or VM. Once the VM data has been restored, one would need to reconfigure the VM to run at the remote site and re-IP its networking. After this is done, the VM can be powered on and recovered from its backup.

Furthermore, any backup package can be used to recover VM file data at a remote site. Without a bare metal restore option, it may take more steps to recover all the VM data, but once it is restored, the rest of the DR process would be similar.

VMware DR can be supported in multiple ways. But any failover automation will depend heavily on the data replication used and the software selected, specifically:

- VMware SRM can easily automate most VM failover, but has some current limitations
- Geo-clustering software provides automatic failover functionality, but except for VCS, is limited to only a single operating system
- SAN or array replication can also be used, but requires hand customized scripting to semi-automate failover
- Most data protection packages support DR, but require customized scripting to semi-automate failover

VM DR does not have to consist of only one approach alone. Due to replication expenses, automated failover may well be limited to only a few critical virtual machines, with the rest relegated to less automated recovery. Such a multi-tier DR plan can easily be supported with combinations of the above products to support fully automated recovery for critical virtual machines and more manual recovery for the rest.

*About the author: Ray Lucchesi is president of Silverton Consulting, a storage, strategy and systems consulting services company, based in the USA offering products and services to the data storage community.*