SearchDisasterRecovery.com

Tutorial

VMware vCenter Site Recovery Manager disaster recovery best practices

VMware Inc. released Site Recovery Manager (SRM) in June 2008 to provide an automated solution for failover of virtual environments to a recovery site. Site Recovery Manager was developed to help automate and simplify the recovery process to a disaster recovery (DR) site. It allows you to create recovery plans using vCenter Server; extend recovery plans with custom scripts; perform non-disruptive testing; automate execution of recovery plans with a single command and to reconfigure virtual machine (VM) networking at the recovery site.

In this tutorial on Site Recovery Manager disaster recovery best practices, learn what you need in place to use VMware Site Recovery Manager, how to test and execute your recovery plans, and if Site Recovery Manager would be a good fit for your company's disaster recovery strategy.

VMware vCenter Site Recovery Manager disaster recovery best practices

By Eric Seibert

Benefits of Site Recovery Manager for disaster recovery

Executing recovery plans is never an easy process, and vCenter Site Recovery Manager may help eliminate some of the complexity and manual processes, such as utilizing run-books and executing scripts to perform a failover. However, once you fail over, you have to eventually failback, and Site Recovery Manager also provides a mechanism for that. Testing recovery plans is also a critical and often disruptive process, but SRM allows for non-disruptive testing by isolating recovery virtual machines on virtual networks that are not part of production networks.

While Site Recovery Manager requires vCenter Server, keep in mind that it's not a complete solution for disaster recovery and relies on a supported third-party data storage replication application to handle the replication of virtual machine data to a recovery site. VMware works with storage vendors to certify that their storage arrays are supported and integrated with SRM. This includes storage vendors such as 3PAR, Compellent, Dell, EMC Corp., FalconStor Software, Hitachi Data Systems, Hewlett-Packard (HP) Co., IBM Corp., NetApp, Sun Microsystems Inc. and Xiotech Corp.

A Site Recovery Manger timeline

The first release of Site Recovery Manager was version 1.0 and was followed up with Update 1 that was released six months later. In October 2009, VMware released SRM version 4.0 that provided support for vSphere. The reason for the jump from version 1.0 to 4.0 was to keep it in line with the vSphere version numbers. VMware also added vCenter to the product name as it was considered a vCenter management and automation product. In addition to support for vSphere, SRM 4.0 also introduced many new features including the following:

- Support for NFS datastores. Previously, only iSCSI and Fibre Channel (FC) datastores were supported.
- Support for shared recovery sites. This allows for multiple protected sites to recover to a single recovery site, previously a separate recovery site had to be created for every protected site.
- Support for protecting VMs that are using the Fault Tolerance (FT) feature.
- Support for Distributed Power Management (DPM) to power off hosts at the disaster recovery site when not in use.

What you need to run Site Recovery Manager

Site Recovery Manager consists of an application server that runs on a Windows server and has its own database and a client plug-in for the vSphere client. An SRM server and a vCenter server must be deployed at each site. Physical servers can be used for these sites, but in many cases they are deployed on virtual machines instead. The new linked mode feature in vCenter Server allows you to have a single pane of glass that allows for easier management of the environment. SRM has some additional requirements that include the following:

- A reliable high-speed network connection between sites, preferably a dedicated connection, because the use of VPN connections over the internet is not recommended.
- Both the recovery site and protected site must have array-based replication between them using a storage array replication adapter that is supported by SRM.
- The recovery site must have sufficient hardware, storage and network resources to support the VM workloads at the protected site.
- The recovery site should have access to the same IP networks (public & private) as the protected site.

For additional operating system and database requirements you should consult the Site Recovery Manager compatibility matrix. SRM will not work with just any storage array that supports data replication as a specific storage replication adapter (SRA) must be developed by each vendor for use with SRM. For a list of SRM storage arrays, you can consult the

Storage Partner Compatibility Matrix; you can also download the specific vendor SRA's from VMware's website. Once you download the appropriate SRA, you can then install it on each SRM server. Keep in mind that you must use the same SRA at both the protected and recovery sites and you cannot use different SRA's or versions of SRA's at each site. For example, if you are using an EMC Celerra storage area network (SAN) at your protected site, you must also have an EMC Celerra SAN at your recovery site.

Site Recovery Manager Components

Site Recovery Manager consists of three components. The first component is the storage component, which interfaces with the SRAs and manages the array replication. To do this, configure array managers in SRM which discover your SRAs, select replicated datastore groups and begin managing storage operations.

The second component is protection groups, which specify which virtual machines at the protected site are to be included in the failover to the recovery site. A protection group encompasses a whole datastore and all virtual machines that are located on it are protected. When you create a protection group, select a replicated datastore group containing the virtual machines that you want to protect. Then select a non-replicated datastore at the disaster recovery site that is used to create placeholder virtual machines on it. Placeholders are small files used to identify a VM at the recovery site and are used until testing or failover occurs were they are replaced with a new VM that is created from the replicated storage. If you don't want to protect a VM, be sure to move it off the datastore you selected for your protection group, and onto another datastore that is not part of a protection group. All virtual machines in the protected group must have valid folder, network and resource pool configurations at the recovery site. You can optionally configure inventory mappings to specify specific resources to use at the recovery site.

The final components are the disaster recovery plans which are used to specify how VMs are migrated to the recovery site. Disaster recovery plans contain one or more protection groups and consist of the steps needed to complete the VM recovery such as the following:

- Powering on/off or suspending VMs
- Changing network settings of VMs
- Waiting for OS heartbeats or guest OS shutdown
- Execute applications or scripts

Disaster recovery plans can either be assigned to isolated test networks that do not allow the recovered virtual machines to communicate on the internet or production networks, or to recovery networks that are used in case of an actual recovery.

Testing and executing recovery plans

While SRM provides automated recovery, its real strength is its ability to easily test recovery without disrupting existing production environments, something that is very difficult to do in traditional physical DR implementations. With SRM, when you perform a recovery test, you are using an isolated test virtual network and also a temporary copy of replicated data at the recovery site. When a test runs it goes through each step in the recovery plan except for those that are marked recovery only and involve powering down virtual machines at the protected site. The steps can all be monitored in the vSphere client while the simulated failover test runs. Tests can be paused, resumed or canceled at any time while they are running.

Implementing recovery plans is simple. Users just need to push a button in the vSphere client to begin the failover process. A confirmation message will display a warning that the process cannot be undone and that running the recovery plan will alter virtual machines at both the protected and recovery sites. Then the following occurs when you execute a failover:

- Array replication between the protected site and recovery site stops.
- All the virtual machines at the protected site will be powered down.
- All the placeholder virtual machines at the recovery site are replaced by powered on VMs that will be added to the vCenter Server's inventory.

At that point, the virtual machines at the recovery site will have taken over for the ones at the protected site. Failback to the original protected site is not an automated process; you essentially have to repeat the process in the reverse order. Since array replication was halted during the failover it needs to be re-enabled in the reverse order to replicate data from the recovery site to the original site. But not all arrays support this, so check your SRA documentation first. The original protected site needs to be configured as a recovery site and new protection groups need to be created at the recovery site. Then a failback recovery plan needs to be executed to reverse the roles of the two sites. Once failback completes replication and is disabled again, it must be re-enabled. Plus, the array managers, inventory mappings, protection groups and recovery plans in SRM must all be re-configured at the original protected site.

Site Recovery Manager leverages both server virtualization and storage virtualization technology to make disaster recovery implementation, testing and execution a much simpler and automated process.

If you wish to know more about Site Recovery Manager, Mike Laverick has written an excellent book called ""Administering VMware's Site Recovery Manager" based on SRM 1.0 that is available as a free download and he is in the process of writing another updated for SRM 4.0.

About this author: *Eric Siebert is an IT industry veteran with over 25 years experience covering many different areas but focusing on server administration and virtualization. Eric has been heavily involved with VMware ESX over the last three years and enjoys dealing with the opportunities and challenges that virtualization offers. He is also a VMTN user moderator and maintains his own VMware VI3 information website, http://vmware-land.com. In addition, he is a regular blogger and feature article contributor on TechTarget's http://searchservervirtualization.com and http://searchvmware.com websites.*