ORACLE

ZFS STORAGE APPLIANCE

An Oracle Technical White Paper
March 2014; v2.1

Best Practice Guide for Implementing VMware vCenter Site Recovery Manager 4.x with Oracle ZFS Storage Appliance

ORACLE

Introduction 1
Overview 2
Prerequisites 3
Application Prerequisites 3
Operating System Prerequisites 3
Storage System Prerequisites 3
General Requirements 3
Supported Layouts 4
Configuring Oracle ZFS Storage Appliance Projects for Replication 8
Configuring Site Recovery Manager With Oracle ZFS Storage Appliance Storage Replication Adapter 4.2 8
Configuring Site Recovery Manager Pairing of Protected and Recovery Sites 9
Configuring Array Managers 9
Configuring Inventory 10
Configuring Protection Groups 11
Configuring a Recovery Plan 13
Running a Test Failover 13
Preparing for Failover in a Disaster Recovery Situation 14
Running a Failover 15
Manually Failing Back to the Primary Site 15
Detailed Steps for Running a Manual Failback 16
Conclusion 18
Recommended Resources 19
Oracle ZFS Storage Appliance and VMware ESX Server Resources on Oracle.com 19
VMware Resources 19
Oracle ZFS Storage Appliance Resources 19

Introduction

This white paper details the configuration and deployment of VMware vCenter Site Recovery Manager (SRM) with Oracle ZFS Storage Appliance. Topics include the installation and configuration of Site Recovery Manager as well as running a disaster recovery (DR) plan or DR test plan.

Disaster recovery is more complicated than just failing over the infrastructure of a virtual environment. Disaster recovery procedures must be part of a larger Business Continuity plan. This BC plan should cover all business processes of a company, identifying risks to their continuity as well as risk mitigation strategies for avoiding or minimizing disruption of those business processes. The IT environments are a critical part of such plans, but the plans should not be limited to them. Identifying and defining the following elements are key input requirements for Business Continuity plans:

- Recovery Time Objective (RTO) — How long it takes to recover from the disaster or how long it takes to execute the recovery plan to make critical services available again.

- Recovery Point Objective (RPO) — How far back in time the data will be after the recovery plan has been completed.

The relevant RPO and RTO factors as identified in the BC plan for the business processes should match your architected solution.

The intended audience for this paper is virtual environment administrators, system administrators, storage administrators, and anyone who would like to understand or deploy Site Recovery Manager with an Oracle ZFS Storage Appliance. This paper assumes readers have familiarity with both configuring replication on Oracle ZFS Storage Appliance products and deploying them with VMware ESX. Some understanding of DR solutions is also expected.

Overview

Proper integration of the Oracle ZFS Storage Appliance and the VMware vCenter Site Recovery Manager presents an effective solution for disaster recovery. Understanding this solution's functioning, as well as important considerations in its setup, helps you realize its full benefits.

Site Recovery Manager creates a protected site and a recovery site and enables the automatic recovery of grouped virtual machines (VMs) residing on the Oracle ZFS Storage Appliance products' replicated shares. The virtual machines are grouped into Site Recovery Manager protection groups on the protected site, and the protection groups are placed into recovery plans at the remote recovery site.

Once a recovery plan is executed, Site Recovery Manager clones the replicated projects on the Oracle ZFS Storage Appliance at the recovery site and mounts the project's shares within VMware ESX as network attached storage (NAS) data stores. Site Recovery Manager also reconfigures the virtual machines' networking to work at the recovery site and then powers the virtual machines on.

In the case of recovery test plans, Site Recovery Manager also connects the virtual machines to a private test bubble network to allow for isolated testing.

A primary benefit of this solution is that replication is supported across the entire Oracle ZFS Storage Appliance product line and across storage profiles. No dedicated link is necessary for replication and any network can be used. Additionally, to provide faster and more efficient target site catch-up, only changes are replicated (except during the initial replication).

Site Recovery Manager is implemented on the Oracle ZFS Storage Appliance through the plug-in software called Oracle ZFS Storage Appliance Storage Replication Adapter v4.2 for VMware vCenter Site Recovery Manager (abbreviated in this document to 'Oracle ZFS Storage Appliance Storage Replication Adapter' or 'Oracle ZFS Storage Appliance SRA').

Some special considerations exist:

- Synchronous mode is not supported, so a Zero Data Loss (ZDL) requirement cannot be met. However, the continuous replication mode can provide an alternative with minimal data loss.
- Discovery and disaster recovery failovers of replicated iSCSI and Fibre Channel (FC) LUNs have been added in the 4.2.0 release of Oracle ZFS Storage Appliance Storage Replication Adapter. This release supports all three protocols: NFS, iSCSI, and FC.

NOTE: References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same family of Oracle ZFS Storage Appliance products. Some cited documentation or screen code may still carry these legacy naming conventions.

Prerequisites

Note the following prerequisites for the Oracle ZFS Storage Appliance Storage Replication Adapter v4.2 for VMware vCenter Site Recovery Manager 4.x.

Application Prerequisites

Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 supports VMware vCenter Site Recovery Manager 4.x and above only. Older versions of Site Recovery Manager (1.x) are not supported with Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0.

Operating System Prerequisites

Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 has been tested with the following VMware software releases:

- VMware ESX 3.5 update 4
- VMware ESX 4.x
- VMware vCenter Server 4.x
- VMware vCenter Site Recovery Manager 4.x

Storage System Prerequisites

Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 has been tested with the 2010.Q3.3 and later Oracle ZFS Storage Appliance software releases. Earlier software releases are not supported.

It is recommended that no other replication than the configured VMware Site Recovery Manager setup should exist on the Oracle ZFS Storage Appliance at the protected and recovery sites.

General Requirements

Note the following installation and configuration requirements that must be met:

- VMware vCenter Server is installed and configured at both protected and recovery sites.
- Site Recovery Manager is installed on both sites.
- Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 is installed on both sites (see the installation guide included with the Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software).
- NAS/NFSv3 shares, FC LUNs, or iSCSI LUNs are configured to ESX servers.
- A small Virtual Machine File System version 3 (VMFS3) device is configured at the recovery site as a VM placeholder.
- NAS data stores contain configured VMs or vdisks to be discovered by the Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software.

- Replication of the required projects is configured prior to configuration of the Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software.

Supported Layouts

The following diagrams list the supported layouts as well as images of unsupported layouts. The main rule is that a VM must not have virtual disks (vmdk or Raw Device Mapping [RDM]) that reside on two different Oracle ZFS Storage Appliance products. Figure 1 shows two VMs being replicated. Each VM is on a separate Oracle ZFS Storage Appliance.

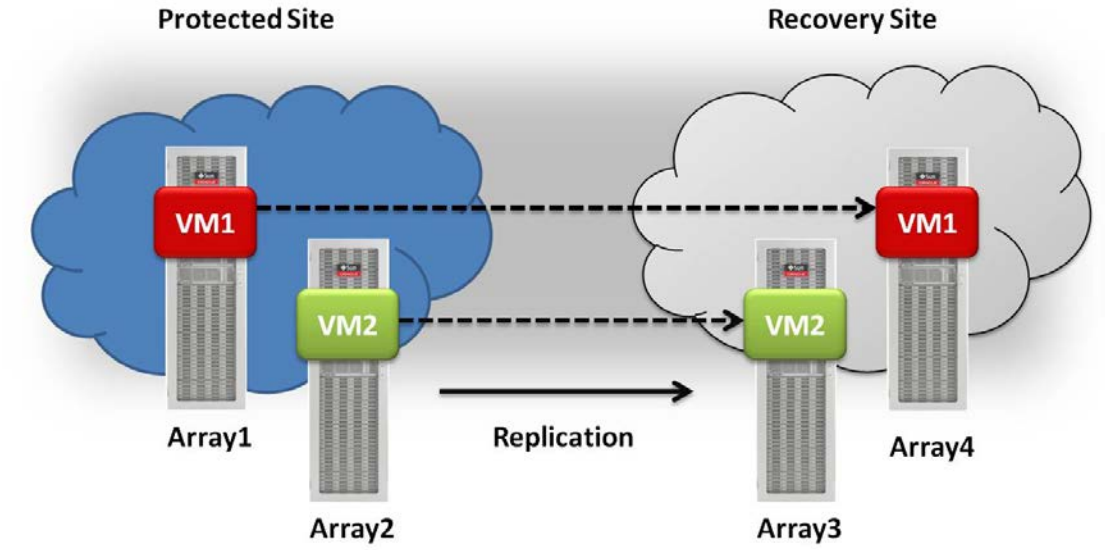

Figure 1. Supported layout with two VMs replicated on separate Oracle ZFS Storage Appliance products

Figure 2 shows two VMs on one Oracle ZFS Storage Appliance at the protected site replicating to two appliances at the recovery site.

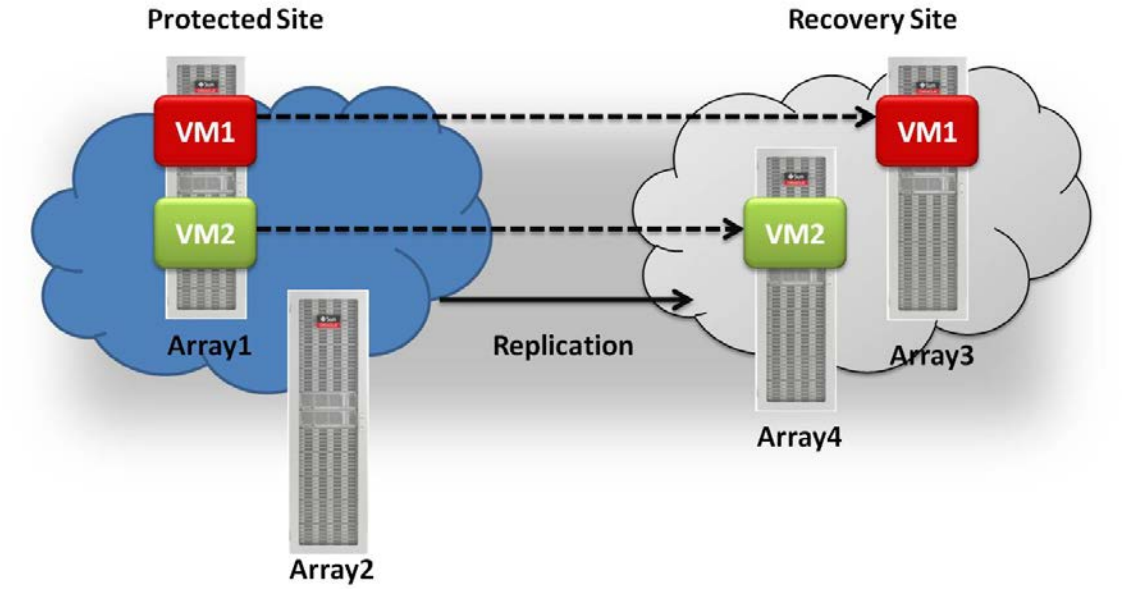

Figure 2. Two VMs on one Oracle ZFS Storage Appliance replicated on separate Oracle ZFS Storage Appliance products

Figure 3 shows VMs on two appliances at the protected site replicating to a single appliance at the recovery site.

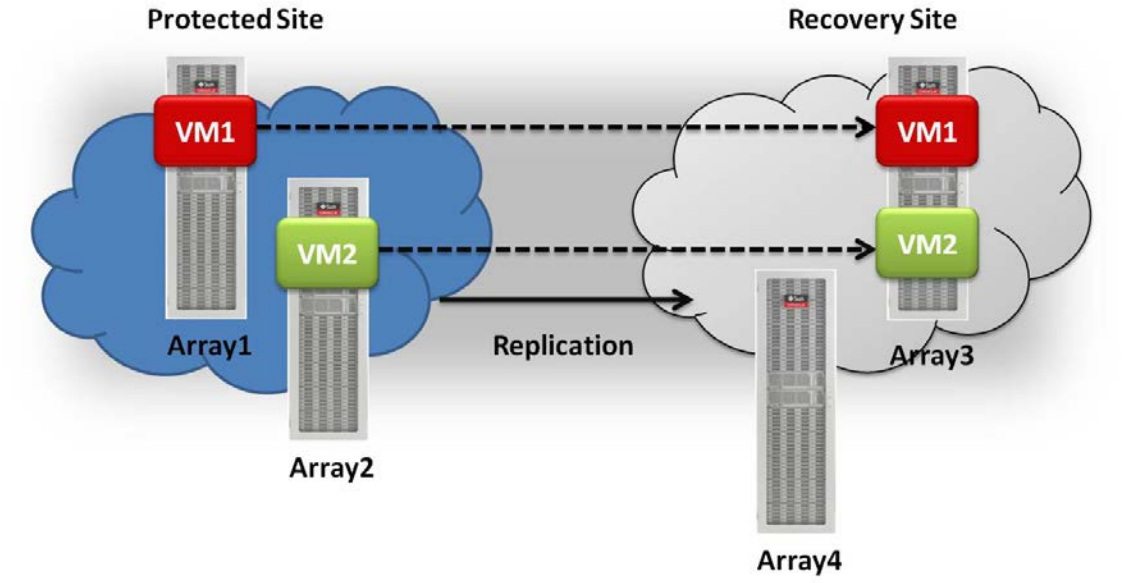

Figure 3. Two VMs on separate Oracle ZFS Storage Appliance products replicated on one Oracle ZFS Storage Appliance

Figure 4 shows an unsupported layout. The VM has virtual disks residing on multiple appliances at the protected site.

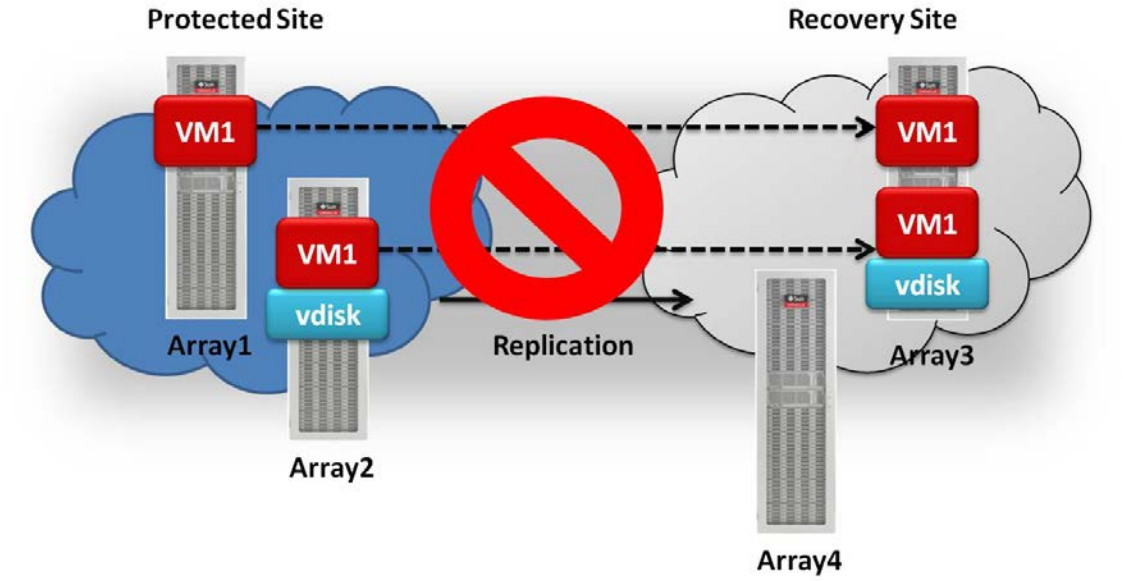

Figure 4. Unsupported: virtual disks residing on multiple Oracle ZFS Storage Appliance products

Figure 5 shows another unsupported layout. The VM is replicating virtual disks to multiple appliances at the recovery site.

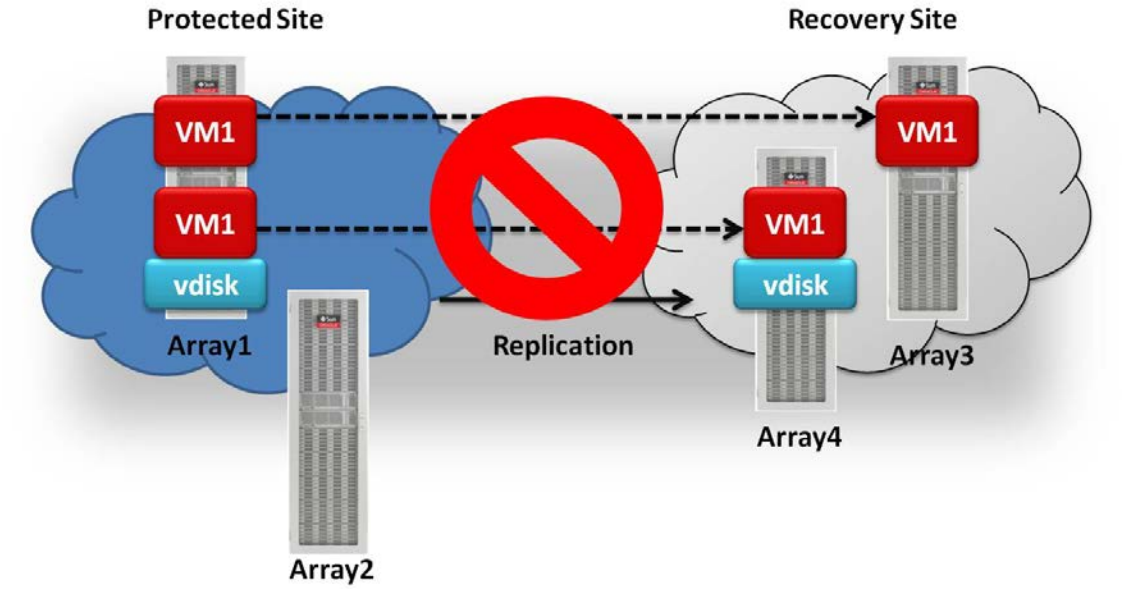

Figure 5. Unsupported: virtual disks replicated on multiple Oracle ZFS Storage Appliance products

Figure 6 shows an unsupported layout in which the same VM is being replicated to multiple appliances at the recovery site.

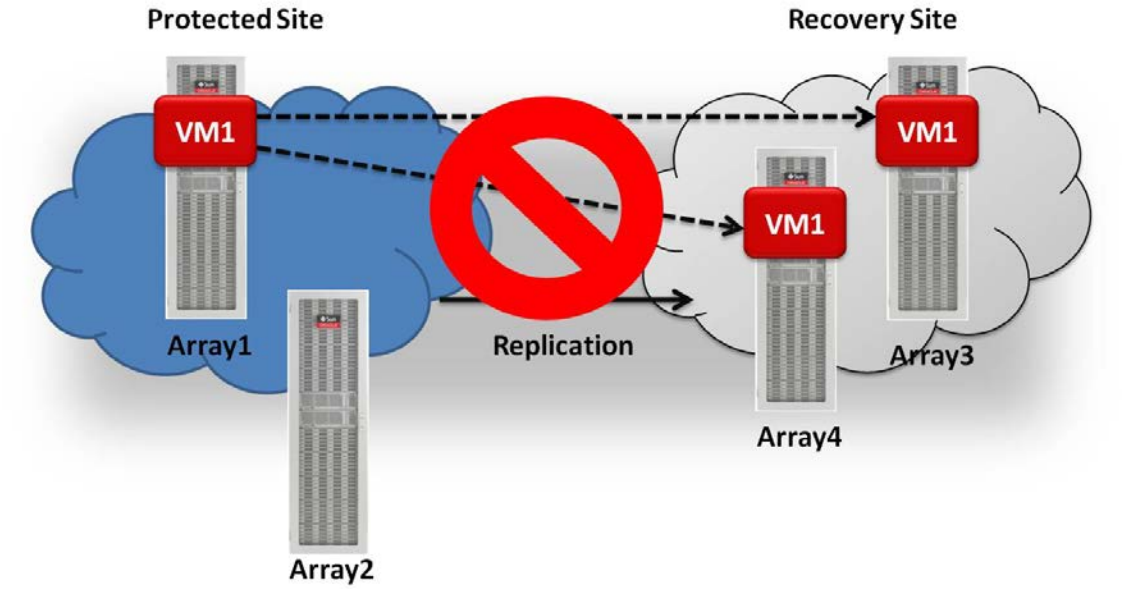

Figure 6. Unsupported: same virtual disk replicated on multiple Oracle ZFS Storage Appliance products

Configuring Oracle ZFS Storage Appliance Projects for Replication

The Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software requires Oracle ZFS Storage Appliance projects to be configured for remote replication. Configuring project-level replication therefore enables the automatic replication of all constituent NFS shares, FC LUNs or iSCSI LUNs. Each project in a protected Oracle ZFS Storage Appliance can be replicated to only one recovery Oracle ZFS Storage Appliance.

A consistency group is a set of shares that are replicated in a consistent fashion with the write order preserved across all the devices in this group. Each Oracle ZFS Storage Appliance project serves as a specific consistency group. Therefore, the ordering of writes to a replicated project's constituent NFS shares or FC or iSCSI LUNs is always preserved.

The Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software itself does not configure replication between the protected or recovery sites. Actions such as configuring replication targets, selecting appropriate projects for replication, and initiating remote replication must be performed prior to configuring Site Recovery Manager to discover replicated Oracle ZFS Storage appliances.

At the recovery Oracle ZFS Storage Appliance, the Storage Replication Adapter does not alter project names, share names, or mount points following test failover operations.

After a DR failover, all replication between the protected and recovery sites should be halted, if possible.

Each VMware data store utilizing NFS, FC or iSCSI protocols should reside on an Oracle ZFS Storage Appliance belonging to a replicated project.

Configuring Site Recovery Manager With Oracle ZFS Storage Appliance Storage Replication Adapter 4.2

The configuration of Site Recovery Manager with Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 and an Oracle ZFS Storage Appliance requires the following high-level tasks:

1. Install the Site Recovery Manager plug-in on each vCenter Server.

   Installation of the Site Recovery Manager software and plug-in is documented in the *VMware vCenter Site Recovery Manager Administration Guide*.

2. Install the Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 software on each vCenter Server.

Information on installation of Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 is provided in the *Sun ZFS Storage 7000 Storage Replication Adapter for VMware Site Recovery Manager Administration Guide* (see the Recommended Resources section at the end of this document). The guide is installed to the installation location when the executable is run on the vCenter server. For Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0, pay close attention to the steps detailing the installation of the Crypt SSLeay Perl modules.

3. Configure the Site Recovery Manager pairing between the protected site vCenter Server and recovery site vCenter Server.
4. Configure the Storage Array manager to communicate with the Oracle ZFS Storage Appliance products.
5. Create protection groups at the protected site.
6. Create a recovery plan at the recovery site.
7. Test the recovery plan.

Configuring Site Recovery Manager Pairing of Protected and Recovery Sites

After enabling the Site Recovery Manager plug-in at both the protected and recovery sites, pair the sites together by browsing to the Site Recovery Manager plug-in GUI on the protected vCenter Server.

1. Click Connection: Configure in the Protection Setup screen and enter the address of the recovery site vCenter Server.
2. Enter the user name and password.

The following screen confirms a successful pairing:

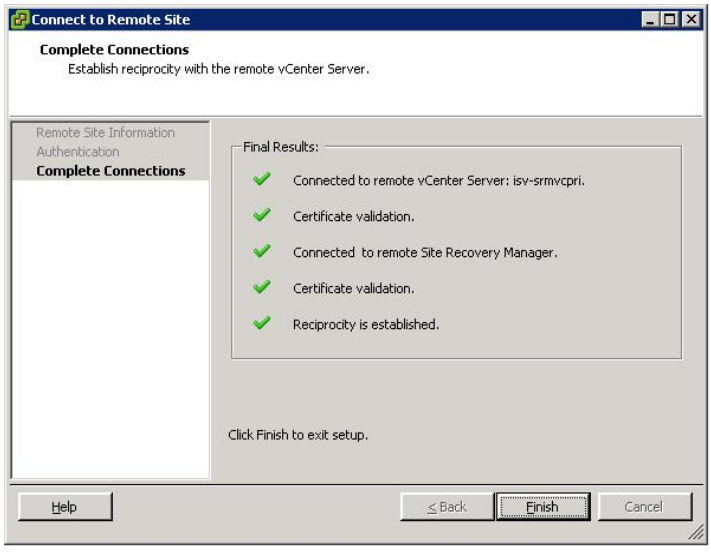

Figure 7. Successful site pairing

Configuring Array Managers

1. Click the Configure link next to Array Managers. Follow the wizard to add a protected site array.
2. For the Manager Type, enter Sun ZFS Storage Appliance.

3. For the URL, enter the management URL of the Oracle ZFS Storage Appliance that is being configured, for example: `https://172.20.100.214:215`.
4. Click OK to continue. The next screen will show the protected site array and recovery site array along with the number of replicated devices.
5. Click Next, and follow the same steps to configure the recovery site array.

After configuring the recovery site array, the final step is to review the replicated data stores, as shown in Figure 8.

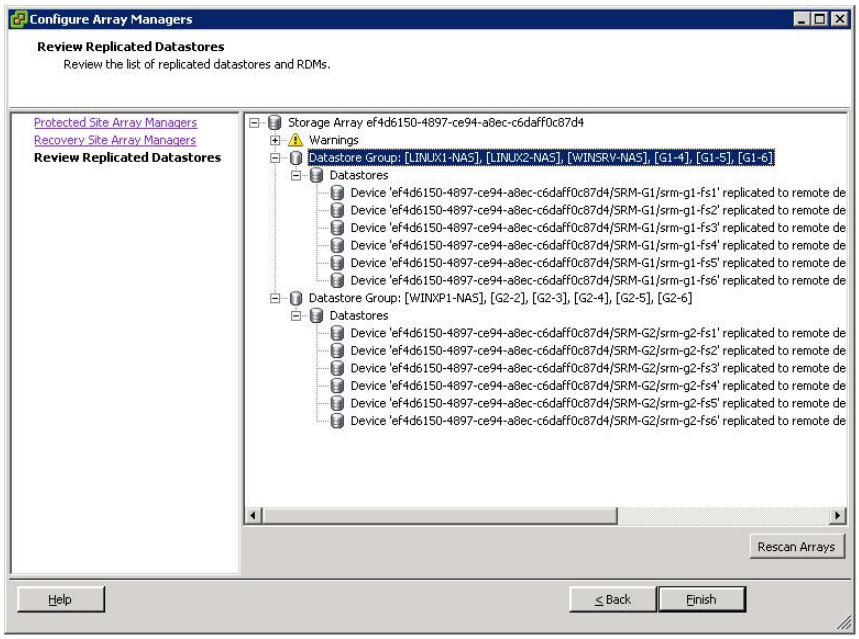

Figure 8. Replicated data stores

This screen lists all replicated data stores and their replication targets. If more NAS data stores are added later, click Rescan Array to discover the new replication pairings.

Perform these steps on the recovery site's Site Recovery Manager plug-in, but with the recovery site array added first. The replicated data stores might appear under the Warnings tag; however, this is normal if the installation is only a unidirectional configuration. Oracle ZFS Storage Appliance Storage Replication Adapter 4.2.0 also supports bidirectional Site Recovery Manager installations, so the recovery site could replicate data stores in the reverse direction with the original protected site acting as a recovery site as well.

Configuring Inventory

To match up resources on the protected site to the recovery site, click the Inventory Mappings: Configure link. Click each resource, and then click Configure. Select a corresponding resource at the recovery site.

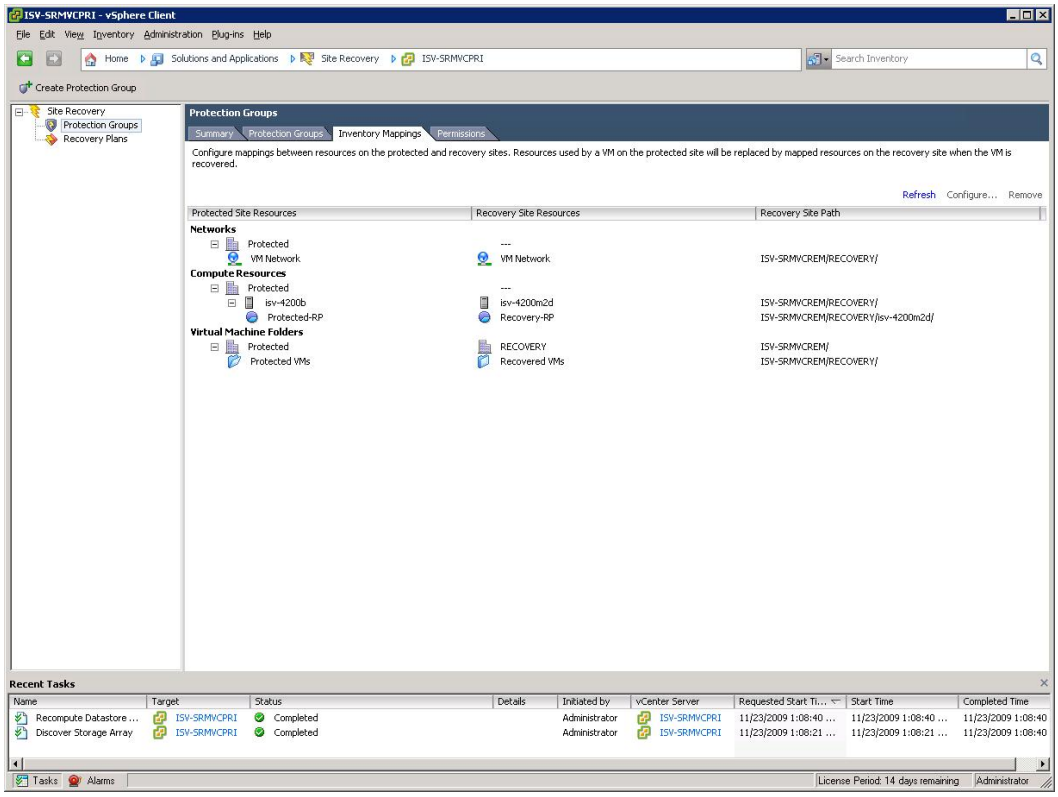

Figure 9. Inventory mappings

Configuring Protection Groups

A protection group is a collection of VMs that are recovered together. The protection group is configured at the protected site and identifies which data stores and which VMs will be recovered on the recovery site in the event of a failover with Site Recovery Manager.

The Oracle ZFS Storage Appliance replicates data at the project level; therefore, all NFS shares, FC LUNs or iSCSI LUNs in a project (and all VMs in the project) are recovered together. If more than one project is being replicated, a separate protection group needs to be configured to hold the VMs in that project.

Use the following steps to define and configure a protection group:

1. At the protected site, click Create Protection Group.
2. Enter a name for the protection group.
3. Select the data stores (collected in Data Store Groups) to be protected by this protection group. After clicking a data store group, the VMs that are in the data store group are shown.

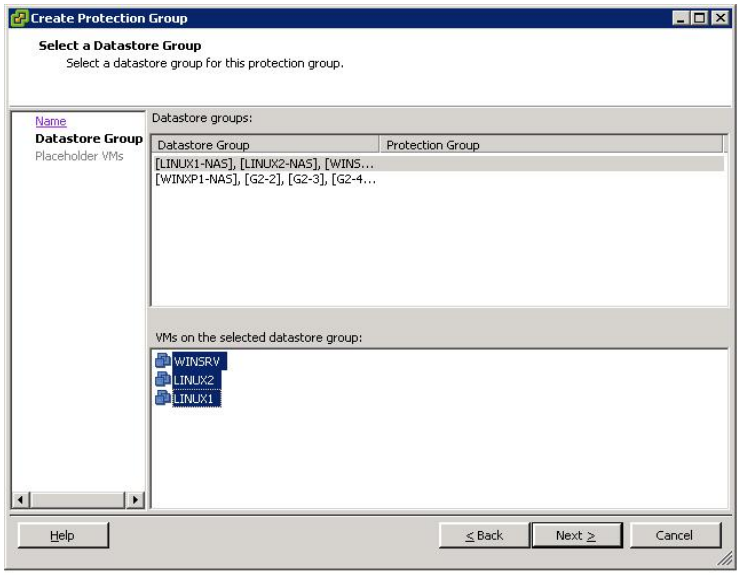

Figure 10. Data store group

4. Select the placeholder data store. The placeholder data store must be pre-configured prior to creating the protection group. This data store can be small because it only needs to hold the `.vmx` file for each VM.

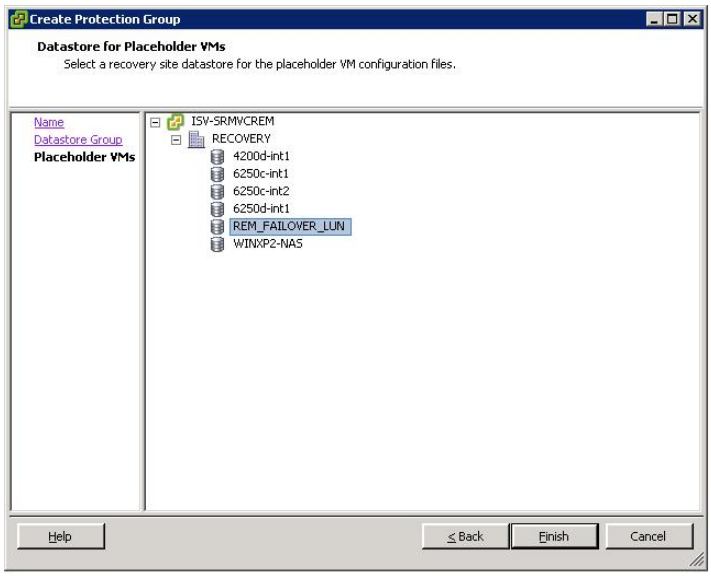

Figure 11. Data store placeholder

5. Configure any specific settings for each VM.

6. If multiple projects (data store groups) are being replicated, repeat these steps for each project to ensure the VMs are protected.

Configuring a Recovery Plan

Recovery plans are built at the recovery site. Each plan can contain one or more previously configured protection groups that were added at the protected site. When a recovery plan is started, all protection groups within that plan are recovered.

1. At the recovery site, click Create Recovery Plan.
2. Enter a name for the recovery plan.
3. Select the protection groups that will be included in the recovery plan.

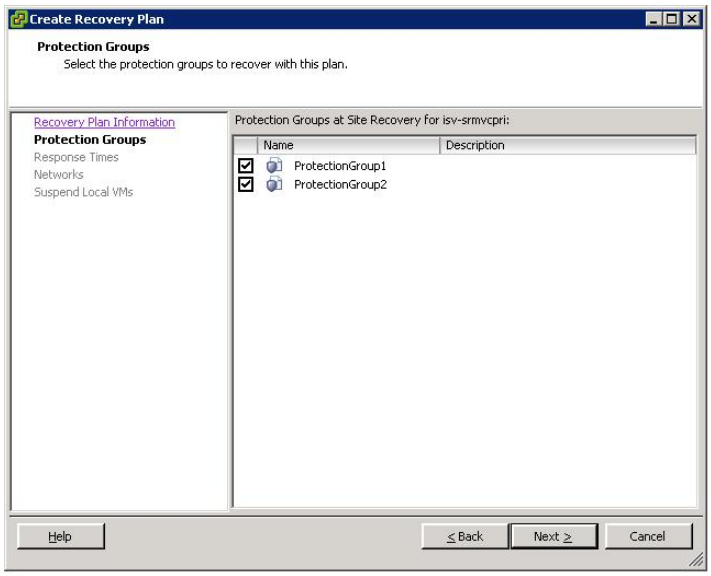

Figure 13. Recovery plan

4. Adjust VM response times, if necessary.
5. Configure the test network settings. If the test network is set to Auto, a test bubble network is created to isolate the network traffic during a recovery test failover.
6. Select any VMs to suspend, if necessary.
7. Click Finish.
8. Review the configuration. The status of the recovery plan should be "OK."

Running a Test Failover

Running a test failover allows you to execute the recovery plan in a controlled and easily rolled-back manner and ensure that proper VM operation will occur in the event of a true failover.

To execute a test failover, highlight the recovery plan to test in the recovery site vCenter Site Recovery Manager plug-in GUI and click Test.

Site Recovery Manager creates the needed clones of the replicated projects and NAS shares, mounts them to the recovery site VMware ESX servers, configures the VMs for the test bubble network, and boots the VMs.

The test failover execution can be monitored in the Recovery Steps tab.

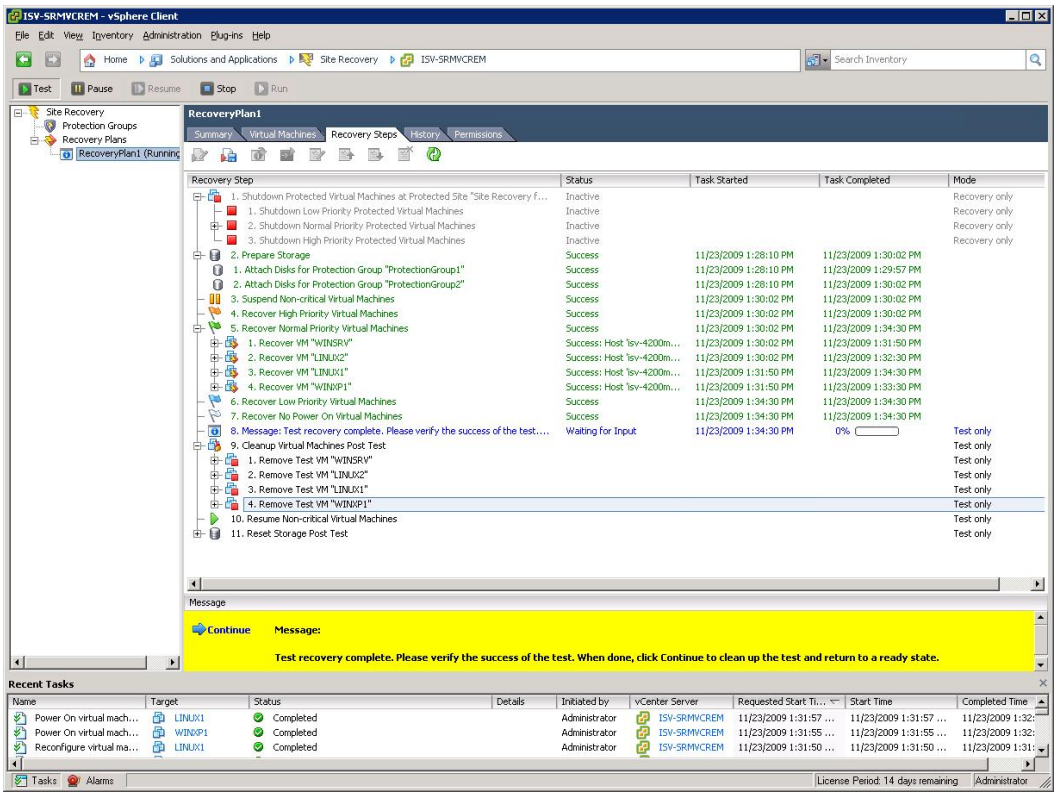

Figure 14. Test failover recovery steps tab

Once the test failover is complete, you can check the recovery VMs for basic functionality. Connect to one of the VMs through the vCenter Console. Perform a basic ping test to ensure that all machines in the isolated test bubble network can be accessed.

After testing is complete, click the Continue link in the Recovery Steps tab to return the configuration back to the original nonfailover condition. Site Recovery Manager powers down the VMs at the recovery site, removes them from the inventory (the placeholders will remain), unmounts the NAS shares from the recovery site VMware ESX servers, and removes the NAS share clones.

Preparing for Failover in a Disaster Recovery Situation

An actual disaster recovery situation involves more than simply failing over VMs. The range of operational, business, and security considerations that must be considered will vary from installation to installation.

Describing a full disaster recovery scenario is, therefore, beyond the scope of this one document. However, there are a few high-level tasks to consider when using VMware vCenter Site Recovery Manager for a DR scenario:

- Ensure that all Site Recovery Manager/Storage Replication Adapter installations have been correctly configured. Test Site Recovery Manager with the Test Failover feature to ensure a correct setup and to verify that the system functions.
- Ensure network access to the recovery site by any system administrators, security administrators, storage administrator and virtual environment administrators that need to work on the DR situation.
- Ensure that adequate infrastructure is in place at the recovery site. This might include DNS servers and Active Directory servers.
- Monitor Oracle ZFS Storage Appliance replication to ensure proper functionality.
- Establish the proper chain of command. Who decides that a DR situation is underway? Who is involved in actually running the DR failover?

Running a Failover

Assume a true DR situation has been declared and a failover needs to occur. To run a failover, log in to the recovery site's vCenter Site Recovery Manager plug-in GUI and perform the following steps:

1. Click the recovery plan that needs to run and click Run.
2. A verification window pops up to confirm that a true failover is requested. Respond accordingly and click Run Recovery Plan.

   The recovery plan executes. If the protected site is down (power outage), failure messages might appear as Site Recovery Manager tries to power down the protected site VMs. This is normal. As with the test failover, Site Recovery Manager creates the needed clones of the replicated projects and NAS shares, mounts them to the recovery site VMware ESX servers, configures the VMs for the mapped network resources, and boots the VMs.
3. Continue with any site-specific procedures needed to bring applications and other infrastructure servers online.
4. If possible and if necessary, perform any steps to isolate the protected site from the recovery site to prevent any problems if the protected site comes online.

Manually Failing Back to the Primary Site

After the protected site has been brought back to service, the most common action is to fail the recovery site back to the protected site to let it handle the normal operations.

The recommended way to do this is to replicate back to the protected site from the recovery site and then use Site Recovery Manager to reverse the protection group and recovery plan flow. Site Recovery

Manager provides no automatic way to do this, so the procedures outlined in this guide to configure normal protected-to-recovery site failover must be done manually.

Detailed Steps for Running a Manual Failback

The following details the specific steps to run the manual failback:

1. Recover any protected site infrastructure.
   a. Ensure that the VMs do not start.
   b. Create and attach a small placeholder data store at the protected site VMware ESX server.
2. Remove outdated VMs from the protected site ESX inventory.
3. Unmount outdated NAS shares from the ESX server.
4. Re-establish network connectivity between the protected and the recovery sites.
5. The Site Recovery Manager failover action performs a 'role-reversal' of the replication setup on the Oracle ZFS Storage Appliance. The role-reversal creates a 'manual' replication on the recovery site Oracle ZFS Storage Appliance, pointing the replication back to the original protected site. Manually update the replication to ensure any updates that have occurred at the recovery site are replicated back to the protected site and to complete the role reversal.
6. The Reverse-Replication update in Step 5 moves any shares at the original protected from a local project (eligible to be mounted read-write) to a replica project (read-only); however, the original local project may still be present, even though it is empty (no shares in the project). Confirm the project is empty, and then remove it.
7. Remove the original Primary Site Protection Group(s) at the original protected site. Also remove the original DR site recovery plan.
8. Remove any leftover virtual machine placeholder entries in the recovery site placeholder data store.
9. Configure the array manager at the recovery site. For a failback, the recovery site Oracle ZFS Storage Appliance system is the protection site and the original protected site is the recovery.
10. Configure the Site Recovery Manager inventory settings in reverse.
11. Create protection groups at the recovery site. A small placeholder data store is required at the protected site.
12. Create a recovery plan at the protected site.
13. Ensure replication is up to date.
14. Perform a test failback.
15. Schedule an outage to perform a controlled failback.
16. Perform manual shutdown of VMs at the recovery site to ensure a proper shutdown of all applications and data consistency.

17. Perform one final replication update to ensure all data is replicated.
18. Execute a recovery plan at the protected site.
19. Verify correct operations at the protected site.
20. Re-create the original protected site to the recovery site replication and Site Recovery Manager relationship.

Conclusion

This white paper described the configuration and deployment of VMware vCenter Site Recovery Manager (SRM) with Oracle ZFS Storage Appliance, including the installation and configuration of Site Recovery Manager and the steps for running a disaster recovery plan.

Recommended Resources

Consult the following resources for further information.

NOTE: References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same family of Oracle ZFS Storage Appliance products. Some cited documentation or screen code may still carry these legacy naming conventions.

Oracle ZFS Storage Appliance and VMware ESX Server Resources on Oracle.com

- "Using Sun Storage 7000 Unified Storage Systems with VMware ESX Server" http://www.oracle.com/technetwork/systems/articles/storage-vmware-jsp-138864.html
- Oracle ZFS Storage Appliance Reference Architecture for VMware vSphere4: http://www-content.oracle.com/content/groups/public/@otn/documents/webcontent/354079.pdf
- *Oracle ZFS Storage Appliance Storage Replication Adapter for VMware Site Recovery Manager Administration Guide* http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html

VMware Resources

- VMware web site: http://www.vmware.com/
- VMware vCenter Site Recovery Manager resources: http://www.vmware.com/products/site-recovery-manager
- VMware ESX Server documentation: http://www.vmware.com/support/pubs/
- VMware communities: http://communities.vmware.com/home.jspa

Oracle ZFS Storage Appliance Resources

- Oracle ZFS Storage Appliance Web site: http://www.oracle.com/us/products/servers-storage/storage/nas/overview/index.html
- Documentation wiki for the Oracle ZFS Storage Appliance: http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html
- Blog of the Fishworks engineering team: https://blogs.oracle.com/fishworks/

Best Practice Guide for Implementing VMware Site Recovery Manager 4.x with Oracle ZFS Storage Appliance
July 2010, v1.0
Author: Ryan Arneson
October 2012, v2.0
March 2014, v2.1
Contributing Author: Anderson Souza

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Oracle is committed to developing practices and products that help protect the environment

SOFTWARE. HARDWARE. COMPLETE.