



Disaster Recovery Best Practices for Microsoft SQL Server 2012 with Dell EqualLogic Auto-Replication and VMware vCenter Site Recovery Manager

Dell Storage Engineering
November 2013

Revisions

Date	Description
December 2013	Initial publication

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

1	Introduction.....	7
1.1	Audience.....	7
1.2	Terminology	8
2	Overview of the DR solution with SRM and EqualLogic replication	10
3	Test infrastructure	12
3.1	Solution configuration - Hardware components.....	13
3.2	Solution configuration - Software components	14
3.3	EqualLogic storage configuration	14
3.3.1	Scenario 1: SRM DR using VSM	14
3.3.2	Scenario 2: SRM DR using ASM/ME.....	15
3.4	SQL Server configuration.....	16
3.5	VMware configuration.....	16
3.6	Network configuration.....	16
3.7	EqualLogic iSCSI storage connectivity	17
4	Solution deployment.....	18
4.1	General environment requirements.....	18
4.1.1	EqualLogic replication configuration.....	18
4.1.2	Configure protected and recovery sites.....	19
4.2	Application consistency and recovery options.....	19
5	Scenario 1 - SRM using EqualLogic Replication with VSM.....	20
5.1	Pairing SRM sites.....	20
5.2	Configuring access to the EqualLogic storage.....	21
5.3	Use VSM to configure smart replicas.....	23
5.4	Configure smart replicas of datastores using VSM	24
5.5	Build Protection groups	27
5.6	Build recovery plans.....	29
5.6.1	Create a new recovery plan	29
5.7	Execute the DR plan in test mode	30
5.8	Execute a recovery plan in a true DR scenario	31
5.8.1	Initiate recovery on the existing SRM plan.....	31
5.9	Resynchronize primary site after recovery	33



5.9.1	Resynchronizing the environment after the primary storage has been recovered.....	33
5.9.2	Resynchronizing the environment after the primary storage is lost and rebuilt.....	35
5.9.3	DR Failback	35
5.10	SQL test results	36
6	Scenario 2 - SRM using EqualLogic replication with VSM and ASM/ME.....	38
6.1	SRM testing with ASM/ME	39
6.2	ASM/ME configuration for SQL Server.....	39
6.2.1	ASM/ME with SQL Server overview	41
6.2.2	Snapshot Smart Copies.....	41
6.2.3	Clone Smart Copies.....	42
6.2.4	Replica Smart Copies.....	42
6.3	Execute recovery plan in a true DR scenario	49
6.3.1	Steps to recover SQL Database.....	49
6.4	DR test results using ASM/ME.....	51
6.4.1	DR failback using ASM/ME	51
7	Supported storage replication topologies with SRM.....	52
7.1	Additional best practices and considerations	53
7.1.1	TempDB and OS page file considerations for replication.....	53
7.1.2	Protect VMs with non-replicated Windows page files.....	53
7.1.3	Customize IP addresses at the DR site	54
7.1.4	Behavior of VMware datastore names during SRM failover.....	54
8	Conclusions and findings	56
A	Additional resources	57



Acknowledgements

This best practice white paper was produced by the following members of the Dell Storage team:

Engineering: Srikanth Nandigam

Technical Marketing: Omar Rawashdeh, Magi Kapoor

Editing: Margaret Boeneke

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to SIFeedback@Dell.com.



SIFeedback@Dell.com



Executive summary

Online transaction processing (OLTP) systems are usually at the core of business operations running business-critical workloads that require increased levels of performance, security, ease of management and availability. Disaster recovery (DR) and business continuity (BC) planning are processes that help organizations prepare for unforeseen and disruptive events — whether an event might be a major disaster or simply a power outage. When failing over business operations to recover from a disaster or disruptive event, there are usually several steps that are manual, lengthy and complex. Often custom scripts are written and used to simplify some of these processes. These processes can affect the recovery time objective (RTO) that any DR solution can deliver as well as increase operational expenditure (OpEx) costs.

Implementing a virtualized environment using VMware vCenter Site Recovery Manager (SRM) on Dell EqualLogic PS Series storage provides the infrastructure with unique capabilities. Such capabilities include the implementation of automated DR processes that are quick to implement, easy to test, and significantly reduce RTO and OpEx costs. The EqualLogic PS Series Auto-Replication feature coupled with Virtual Storage Manager (VSM) for VMware and EqualLogic Auto-Snapshot Manager/Microsoft Edition (ASM/ME) offers differentiated SQL application-consistent DR solutions.

This technical paper covers high-level How-To steps and details two validated configurations of a DR reference architecture leveraging EqualLogic VSM and ASM/ME using VMware SRM.



1 Introduction

With the explosion of data, database and IT administrators are challenged to provide availability and protection against data corruption or loss. OLTP systems running on Microsoft SQL Server 2012 are among the most commonly seen in today's business applications. The data availability and protection of these systems is very critical to the success of any business, and SQL Database administrators (DBAs) are concerned about SQL server support in virtualized environments. DBAs want to design and deploy an SQL-based OLTP infrastructure that provides high availability, DR and enterprise level performance. DBAs and system administrators require a robust automated DR solution that can leverage the existing storage based replication to achieve both business continuity and DR.

This paper provides the detailed overview of the DR solution for SQL server(s) running in a VMware virtualized environment using the VMware Site Recovery Manager and the EqualLogic Auto-Replication feature. This white paper offers guidance on the design and validation of the VMware SRM and EqualLogic solution for protecting Microsoft SQL servers. In this reference architecture, the virtualized Microsoft SQL servers are protected and recovered by VMware SRM, EqualLogic Auto-Replication with EqualLogic VSM for VMware, and EqualLogic ASM/ME to achieve application consistency during site failover and failback. This reference architecture addresses many challenges that DBAs and IT administrators face every day to protect SQL databases and how they can benefit from using EqualLogic storage replication with EqualLogic ASM/ME:

- VMware SRM in conjunction with EqualLogic Auto-Replication improves the RTO by fully automating the DR process.
- This solution helps quickly restore business operations by providing instant recovery of virtual machines (VMs) using SRM.
- SRM allows tests of DR functionality without disturbing the normal business operations by leveraging EqualLogic clones on the remote site and by bringing VMs in a bubble network.
- Using EqualLogic VSM and/or ASM/ME, SRM allows you to recover entire VM(s) with SQL application consistency.
- VSM plugin for vCenter allows the creation of hypervisor consistent snapshots, clones and replicas for data protection and DR.
- ASM/ME provides Volume Shadow Copy Service (VSS) application consistent snapshots, clones and replicas for NTFS, Microsoft SQL Server, Microsoft Exchange, and Microsoft SharePoint applications

1.1 Audience

This reference architecture is intended for storage administrators, storage architects, VMware administrators, and Microsoft SQL DBAs who are already familiar with implementing and administering EqualLogic storage in VMware infrastructure and are considering deploying VMware SRM for protecting SQL servers using the EqualLogic Auto-Replication feature. It is assumed that the reader is familiar with various components of this solution.



1.2 Terminology

Application Consistent: This is the state in which all databases are in-synch and represent the true status of the application. Database application responds to its VSS writer being triggered by flushing all of its memory and I/O operations to disk so that the database is completely consistent. In doing so, there is nothing in memory and no pending I/O to be lost. A VSS writer should effectively place all the data for an application in the same state as it would be if the application were properly closed.

Crash Consistent: This is the state after a system failure or power outage. A crash-consistent database image is consistent with a copy of a database image after the database instance, server, or storage system has crashed. If a database system is restored to a crash consistent state, then it is necessary to follow some manual recovery procedures.

Datastore: Virtual representations of combinations of underlying physical storage resources in the datacenter. A datastore is the storage location (for example, a physical disk, a RAID, or a SAN) for virtual machine files.

Delegated Space: The amount of space on the secondary group that is delegated to a replication partner to be reserved for retaining replicas.

EqualLogic ASM/ME: EqualLogic Auto-Snapshot Manager/Microsoft Edition (ASM/ME) enables you to create fast, space-efficient, point-in-time copies of EqualLogic volumes attached to Microsoft hosts.

EqualLogic VSM: The EqualLogic Virtual Storage Manager (VSM) for VMware is an application that you install as a plug-in to VMware vCenter Server.

Local Reserve: The amount of space reserved on the local or primary group for holding temporary snapshots and failback snapshots of the source volume.

Protected Site: The datacenter containing the VMs for which data is being replicated to the recovery site.

Recovery Site: The datacenter containing the recovery VMs performing work while the protected site is unavailable.

Replica: A time synchronized copy of an EqualLogic volume stored in a secondary EqualLogic group.

Replica Reserve: The space allocated from delegated space in the secondary group to store the volume replica set for a specific volume.

Replica Set: A collection of all time-synchronized replicas for a specific source volume.

Replication: A service that produces a point-in-time copy of a production volume and updates the remote volume to be consistent with the source volume.

RPO: The recovery point objective is the amount of data loss that is acceptable and defined by application in case of disaster.



RTO: The recovery time objective is the amount of time it takes to recover the lost or corrupted data from backup.

Snapshots: A logical point-in-time view of a volume.

SRA: A Storage Replication Adapter allows VMware Site Recovery Manager (SRM) to integrate with 3rd party storage array technology. SRA allows VMware SRM to communicate with EqualLogic storage array to perform certain replication tasks.

Virtual Machine: A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. Multiple virtual machines can operate on the same host system concurrently.

VMware SRM: A DR management and automation software that accelerates recovery by automating the recovery process in a VMware virtualized infrastructure.



2 Overview of the DR solution with SRM and EqualLogic replication

VMware vCenter SRM provides business continuity and DR protection for virtual environments. It can protect deployments from individually replicated datastores to an entire virtual site.

In an SRM environment there are two sites, a protected site and a recovery site. Protection groups that contain protected VMs are configured on the protected site and these VMs can be recovered by executing the recovery plans on the recovery site. This DR solution leverages two different scenarios to achieve SQL application consistency, one using EqualLogic VSM for VMware and the second one by using ASM/ME. This white paper validates both scenarios and provides detailed test methodology and steps involved to achieve DR using VMware SRM.

Key highlights of this validated solution architecture:

- VMware SRM works in conjunction with EqualLogic Auto-Replication technology through the EqualLogic Storage Replication Adapter (SRA). The EqualLogic SRA is a software package that is installed on each VMware SRM server at each site.
- This reference architecture leverages the EqualLogic Auto-Replication feature, which is snapshot based point-in-time replication that offers group to group volume replication over extended distance.
- This validated configuration helps to orchestrate an entire site failover using vCenter SRM and EqualLogic Auto-Replication.
- EqualLogic VSM or ASM/ME can be used to achieve application consistency during the failover of VMs using VMware SRM.
- In the first scenario, SRM leverages EqualLogic VSM to schedule smart replicas and to recover entire VMs, quiesce the disk I/O by putting VM in a VMware snapshot mode and flushing I/O to the disk.
- In the second scenario, the SRM leverages EqualLogic ASM/ME to provide SQL consistent replicas for DR, snapshots and clones for local database protection and distribution. ASM/ME used in conjunction with SRM provides more granular database recovery. ASM/ME allows for recovery of individual objects of an application rather than the entire VM.

The following diagram shows various components in the solution and communication between them.



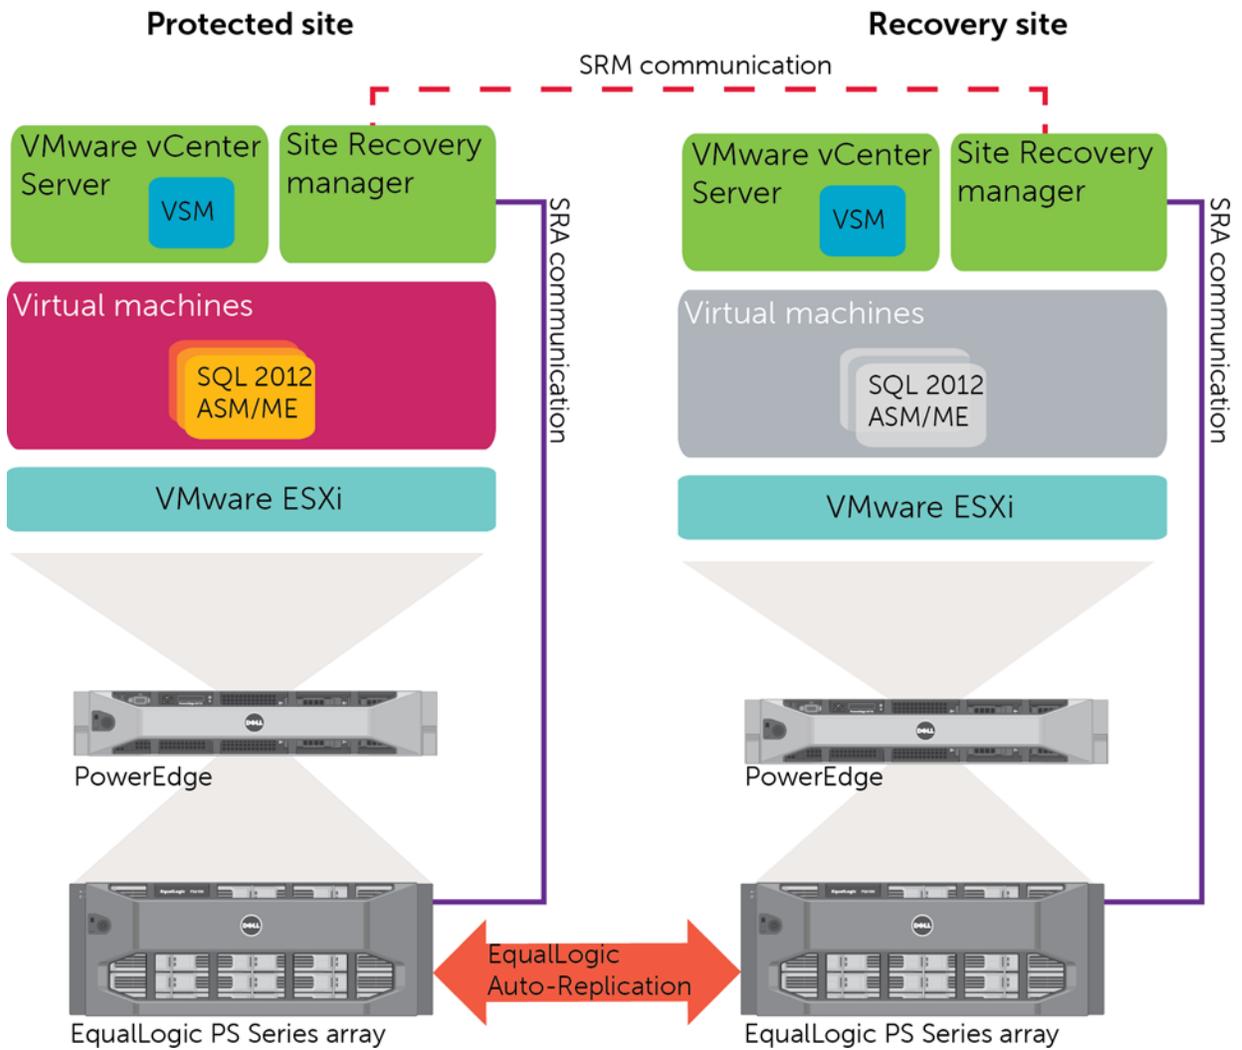


Figure 1 High-level solution architecture SRM components



3 Test infrastructure

This solution consists of a fully virtualized SQL Server 2012 server environment, running on VMware ESXi 5.1 virtualization layer, with vCenter SRM used to provide the automated site failover.

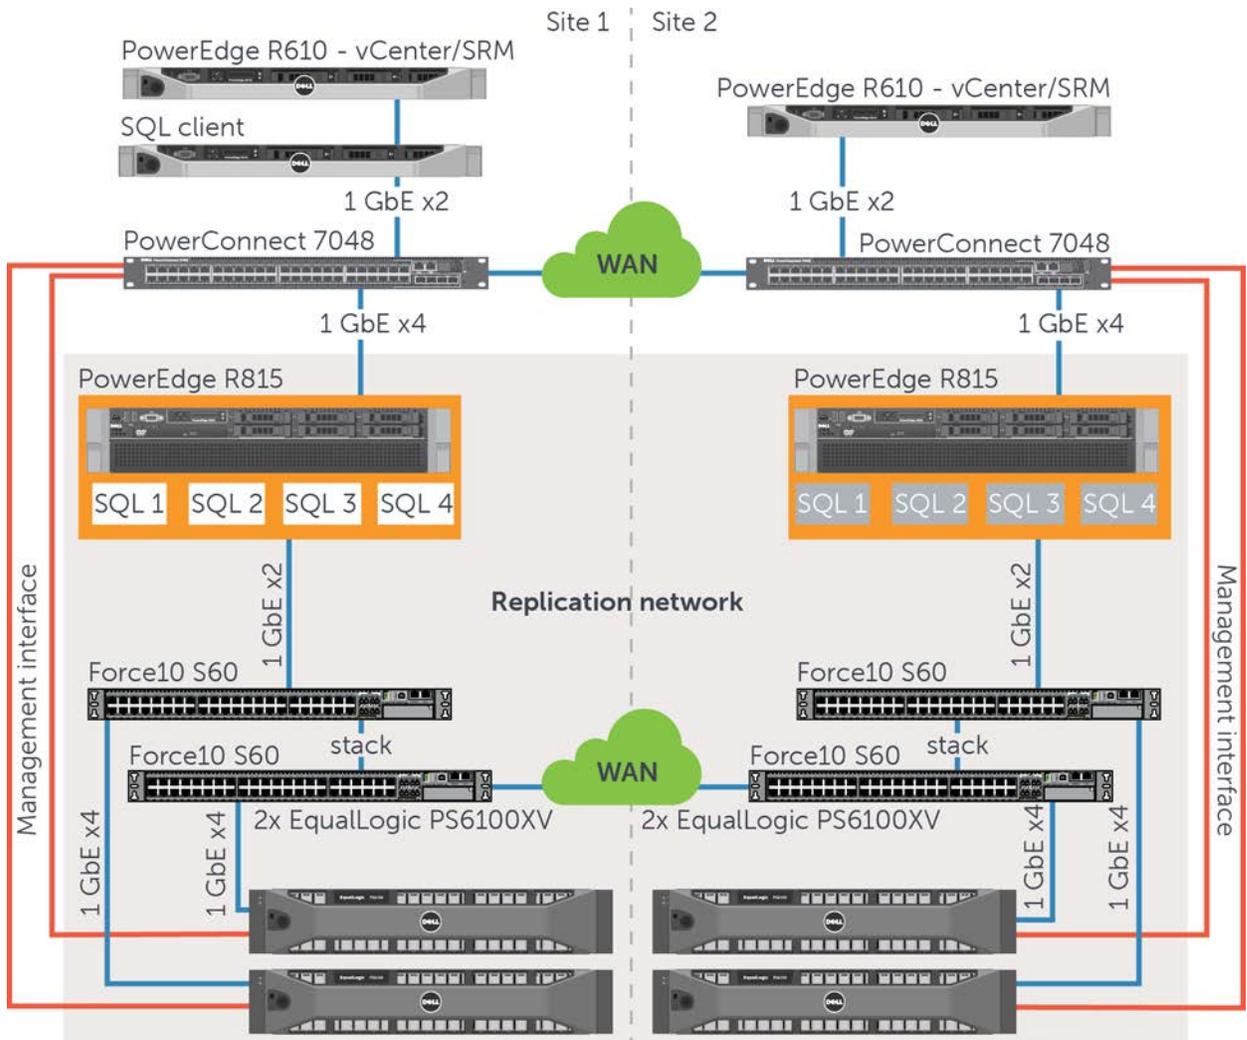


Figure 2 SRM reference architecture

The layout of the environment tested in the lab for implementing VMware SRM on EqualLogic storage consisted of the following:

- Production site or Primary Site, referred to as Site1 contains the following:
 - VMware ESXi 5.1 host configured to access EqualLogic storage via iSCSI Software initiator.
 - One EqualLogic PS Series storage group consisting of two 6100XV arrays to provide shared storage for Virtual Machine File System (VMFS) datastores and external volumes for SQL server database.



- A physical Windows Server 2008 R2 server running VMware vCenter Server and SRM with the EqualLogic SRA installed.
- Four VMs running Windows Server 2008 R2 with Microsoft SQL Server 2012 installed and populated with a test database created using Quest Benchmark Factory for Databases tool.
- DR site or secondary site, referred to as Site2 contains the following:
 - VMware ESXi 5.1 host configured to access EqualLogic storage via iSCSI Software initiator.
 - One EqualLogic PS Series storage group consists of two 6100XV arrays to provide storage for VMFS datastores and external volumes for SQL Server databases.
 - A Physical Windows Server 2008 R2 server running VMware vCenter Server and SRM with the EqualLogic SRA installed.

EqualLogic Auto-Replication is used to replicate volumes containing VMFS datastores from the primary site to the DR site. All VMs at the primary site are protected by EqualLogic auto replication and SRM.

In an SRM environment, the site in which a VM is currently running is referred to as the protected site for that VM. The site to which the data of the VM is replicated is referred to as the recovery site for that VM. When using SRM to manage a failover, it happens at the same granularity as the EqualLogic Volume/Datastore relationship. That is, if you have configured an EqualLogic volume as a datastore, all VMs in that datastore will be part of the same SRM protection group and therefore part of the same SRM recovery plan.

3.1 Solution configuration - Hardware components

Table 1 Hardware component design details

Equipment	Qty	Description
Dell PowerEdge R815 Server	2	One VMware ESXi Server hosting four SQL VMs at each site (Primary and Secondary).
Dell PowerEdge R610 Server	2	A server at each site to run VMware vCenter and SRM and EqualLogic SRA.
Dell Networking S60	4	Two S60 network switches stacked together at each site to provide network for local iSCSI traffic.
EqualLogic PS6100XV	4	Two PS6100XVs at each site form a single group.



3.2 Solution configuration - Software components

Table 2 Software component design details

Software	Version	Description
VMware ESXi	5.1	VMware ESXi Server hosting SQL VMs at each site (Primary and Secondary).
VMware vCenter Server	5.1	To Manage vSphere and its plug-ins.
VMware SRM	5.1	SRM software for automating the DR activities.
EqualLogic SRA	2.2	Enables SRM to communicate with EqualLogic array to perform replication related tasks such as promoting replicas.
EqualLogic VSM	3.5.2.1	EqualLogic Virtual Storage Manager for VMware
EqualLogic Multipathing Extension Module	1.1.1-262227	Multipath driver for enhanced path load balancing.
Microsoft Windows Server	2008 R2	Operating system running on VMs.
Microsoft SQL server	2012	SQL Database.
EqualLogic Array Firmware	V6.0.5	Storage arrays with built in Auto-Replication feature.
EqualLogic ASM/ME	V4.6.0	SQL host integration tools for EqualLogic.

3.3 EqualLogic storage configuration

In this solution testing, two EqualLogic PS6100XV arrays were configured in a single EqualLogic group at each site. The storage was provisioned in two different configurations to address the storage needs for two different DR methods outlined in the document.

3.3.1 Scenario 1: SRM DR using VSM

In the first scenario, the solution uses VSM to achieve application consistency. The VM boot volume is created from the VMFS datastore, while all other volumes are also accessed from the same VMFS datastore. While this approach does not provide granular database recovery, it uses VSM to provide application consistency during database failover and failback. This allows a fully automated failover and failback using SRM. VMware Data Recovery uses Microsoft Windows Volume Shadow Copy Service (VSS) quiescing, VSM provides a mechanism for creating application consistent point-in-time copies of guest VMs running on Windows 2008 R2. VSS produces consistent shadow copies by coordinating with business applications, file-system services and storage hardware. VSS support is provided with VMware Tools, which runs in the guest operating system. VSM acts as a VSS requestor that requests creation of smart copies of datastores on the PS EqualLogic array

Each VM is created from a separate datastore residing on different EqualLogic volumes.



Table 3 shows the EqualLogic Disk Layout and how it is mapped on each of the four VMs:

Table 3 EqualLogic disk layout used in Scenario 1

EqualLogic Storage Volumes	VMware VMFS datastore	VMware VMDK Files	Windows NTFS File systems	User Capacity (GB)
Eql-esx1-vol(xx*)	vm-boot(xx*)	Boot Disk	C:\	200
Eql-esx1-vol(xx*)	vm-boot(xx*)	SQL data	E:\	300
Eql-esx1-vol(xx*)	vm-boot(xx*)	SQL Logs	F:\	50
Eql-esx1-vol(xx*)	vm-boot(xx*)	SQL TempDB	G:\	50

*xx- Each VM is mapped with a number starting 01.

3.3.2 Scenario 2: SRM DR using ASM/ME

In Scenario 2, the EqualLogic storage layout differs from the previous one. For each VM running SQL Server 2012, this table provides the storage layout on EqualLogic that uses ASM/ME to achieve application consistency. The VM boot volume is created from the VMFS datastore, while all other volumes are accessed through the guest Windows Server iSCSI initiator. This approach provides a mechanism to protect the database at a more granular (database component) level by using ASM/ME for providing local Smart Copies (snapshots and clones) for recovering the database upon corruption or human errors. This approach, in addition to the site failover of SQL servers, provides local protection of databases running on individual VMs using snapshot and clone Smart Copies. ASM/ME uses the Microsoft Volume Shadow Copy Service (Microsoft VSS) to provide a framework for backing up and restoring data in the Windows Server environment. ASM/ME creates point-in-time copies of application database volumes on PS Series group, ensuring that the backed-up data is easy to restore and recover. The VSS Writer for SQL server prepares the application for the backup or restore operation, ASM/ME functions as a requestor that requests creation of smart copies via VSS provider which is installed on the guest operating system (Windows Server 2008 R2) that interacts with PS Series EqualLogic array for creating smart copies on the storage array.

Table 4 shows the EqualLogic Disk Layout and how it is mapped on each of the four VMs.

Table 4 EqualLogic disk layout used in Scenario 2

EqualLogic Volumes	VMware VMFS datastore	Windows NTFS File systems	User Capacity (GB)
eql-esx1-vol (2TB)	vm-boot	C:\	200
sql-systemDB(xx*)	N/A	E:\	20
sql-data-(xx*)	N/A	F:\	300
sql-log-(xx*)	N/A	G:\	50
tempDB-(xx*)	N/A	H:\	50

*xx- Each VM is mapped with a number starting 01.



Note: The EqualLogic configurations used in this solution do not necessarily reflect the performance best practices for running SQL Server. It uses this particular configuration to outline how SQL servers can be protected using SRM. For optimal performance, please refer to the white paper, [Configuring EqualLogic for SQL OLTP workloads](#).

3.4 SQL Server configuration

Each SQL server is configured with 24GB of memory and the Dell Quest Benchmark Factory for Databases tool was used to simulate SQL OLTP transactions by running a TPC-E like workload. TPC-E is an industry standard benchmark for OLTP. The database is populated with 200GB of data using the tool. Each SQL server has four different volumes for storing different sets of data, SQL data files and transaction logs are separated on two different volumes, and there are dedicated volumes for storing TempDB and SystemDB. A unique test database was created on each SQL server instance.

3.5 VMware configuration

The following table shows the VMs configured on the ESXi host of the primary site. All VMs are loaded with Windows Server 2008 R2 with SQL Server 2012 running on them.

Table 5 VM configuration details

Virtual machine name	vCPU	Memory (GB)
SQL1	8	24
SQL2	8	24
SQL3	8	24
SQL4	8	24

3.6 Network configuration

Both the protected site and the recovery site consist of a single ESXi host. The network design in this solution for both the protected site and the recovery site is the same. The design is based on Dell Networking switches for the back-end iSCSI storage traffic.

- Dell Networking S60 switches are used for back-end storage traffic and two physical NICs are connected to this network from each ESXi hosts for iSCSI VMkernel traffic.
- Dell PowerConnect 7048 switches are used to connect ESXi hosts at each site to handle management network traffic and VM network traffic.
- For SRM DR testing, one additional VM port group has been added to the networking configuration for the ESXi hosts at the recovery site and protected site.
- Both the protected site and the recovery site EqualLogic arrays are replicated over a dedicated 1Gbps connection.



3.7 EqualLogic iSCSI storage connectivity

The iSCSI storage network used for connecting the ESXi host to the EqualLogic PS Series storage arrays is often a private backend network that is physically separated from the VM network. In this solution, two physical NICs have been connected to the backend storage network and the software iSCSI initiator is configured with the EqualLogic Multipathing Extension Module software to provide intelligent path routing. This is recommended to increase bandwidth and perform automatic connection management between the ESXi host and the PS Series array. On the ESXi host, datastores are formatted using VMFS-5 file system.



4 Solution deployment

This section describes high level requirements and steps involved in configuring the SRM for building a SQL DR solution using two different scenarios for application consistency.

4.1 General environment requirements

The following are required before configuring the SRM for building an SQL DR solution and are applicable to both scenarios described in this document.

- vCenter must be installed at the primary and DR sites. VMware SRM requires two independent ESXi environments, each managed by its own vCenter server.
- vSphere client available at primary and DR sites.
- SRM installed at primary and DR sites.
- EqualLogic Storage Replication Adapter (SRA) installed on the server running SRM servers at the primary and DR sites.
- VM Boot volumes (C:\) must reside on VMFS datastores created from EqualLogic volumes.
- A small volume configured as a datastore at the DR site for storage (local) for placeholder VM.
- EqualLogic Multipath Extension Module (optionally) be installed and configured on both sites for providing optimized data access to the EqualLogic volumes. iSCSI initiators on both ESXi hosts must be configured with their corresponding group EqualLogic IP address at each site.
- EqualLogic Auto-Replication must be configured for volumes containing datastores (which contain VM boot volumes) to be protected.
- EqualLogic VSM must be installed, configured on both sites and must be aware of each other.
- EqualLogic volumes containing datastores should be configured and scheduled for replication using VSM.
- If SQL DB volumes are part of the VMFS datastore, then it is required to use the VSM plugin in the vCenter to provide hypervisor application-consistent replicas.
- If SQL DB volumes are not part of the VMFS datastore, then it is required to use ASM/ME to provide SQL application-consistent replicas.

4.1.1 EqualLogic replication configuration

EqualLogic Auto-Replication is used to replicate volumes between different groups as a safeguard to protect against data loss. The two groups must be connected through a TCP/IP based network. In practice, the actual bandwidth and latency characteristics of the network connection between replication groups must be able to support the amount of data that needs to be replicated and the time window in which replication needs to occur.

IT administrators can enable replication through the Group Manager GUI or the Group Manager CLI. Once enabled, volume replication functions can be managed using the GUI, the CLI, or ASM tools. ASM tools are included as part of the Host Integration Toolkits (HIT) for Microsoft and VMware. A replica can be created from a single volume or a volume collection.



Auto-Replication initially creates a copy on a secondary storage system and then synchronizes the changed data to the replica copy. A replica represents the contents of a volume at a point in time at which the replica was created. This type of replication is often referred to as “point-in-time replication” because the replica copies the state of the volume at the time the replication is initiated. The frequency at which replication occurs determines how old the replica becomes relative to the current state of the source volume.

Replication setup between the two sites must be established via the EqualLogic Group Manager or the CLI for all volumes hosting the VM datastores and SQL database related volumes. Once the primary/secondary relationship is established on EqualLogic arrays, configure all volumes for replication. Creation of replicas or replication schedules is not required as these are managed outside the EqualLogic Group Manager by VSM or ASM/ME. For sizing guidelines on how much delegated space and local reserve is required on EqualLogic, refer to the [Dell EqualLogic Auto-Replication Best Practices and Sizing Guide](#) and the EqualLogic Product Administrator Guide. The amount of bandwidth required for replication also depends on the change rate and RPO requirements.

SRM allows testing the DR process without actually failing over to the remote site. During this test process instead of promoting the EqualLogic volume, this clones the replica set on the remote EqualLogic array. Appropriate consideration should be made on the available capacity on the remote site for a successful SRM DR test procedure.

4.1.2 Configure protected and recovery sites

Configuring SRM to protect VMs replicated by EqualLogic Auto-Replication involves these steps:

1. Install and enable the SRM plug-in on each of the vCenter servers.
2. Install the EqualLogic SRA on each of the vCenter servers. The SRA adapter must be installed on each SRM server prior to configuring arrays in SRM.
3. Pair the SRM sites. This enables the SRM communication between the sites.
4. Configure the storage array managers. This enables the SRM software to communicate with the EqualLogic Groups.
5. Build SRM protection groups at the primary site. Protection groups define groups of VMs that will be recovered together and the resources they will use at the DR site.
6. Build SRM recovery plans at the DR site. Recovery plans identify the startup priority of VMs.
7. Test recovery plans.
8. Failover to the protected site to simulate a planned migration or unplanned disaster.

4.2 Application consistency and recovery options

Data integrity is one of the biggest concerns while deploying SQL Server in a virtualized environment because the hypervisor sits between the VM operating system and the storage subsystem. To ensure SQL database is application-consistent and maintains data integrity upon failover, use one of the two reference architectures outlined in this document to achieve SQL DR protection in a VMware environment running on EqualLogic storage.



5 Scenario 1 - SRM using EqualLogic Replication with VSM

This section describes the recovery process of SQL servers using SRM and VSM. Some steps such as pairing SRM sites and configuring access to the EqualLogic storage arrays from SRM are identical for both methods whether you are using VSM or ASM/ME to achieve SQL application consistency.

In this configuration, four 2TB EqualLogic volumes are presented to the ESXi host. Datastores were created using VMFS-5 file system for storing the boot disks for VMs running Windows Server 2008 R2. Additional disks are presented to the VM for storing DB volumes (logs, temp, DB) from each datastore for their corresponding VM. You can create the datastore from the VSM plugin or the vCenter configuration. Once the datastore is created, Windows VMs are configured to boot from this datastore, with SQL 2012 database installed on other virtual disks.

Figure 3 shows the layout of VMs running Windows Server 2008 R2 where VSM is configured to protect SQL databases.

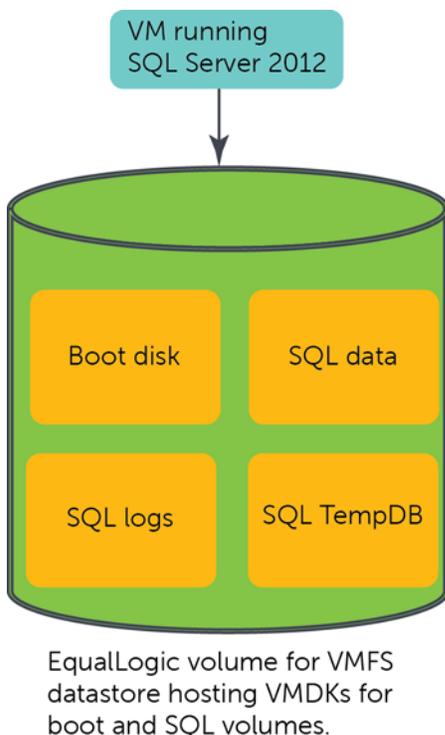


Figure 3 VM datastore layout for Scenario 1 - VSM

5.1 Pairing SRM sites

Each site must have a separate vCenter server with SRM installed and enabled. After installing SRM and enabling it, each vCenter server SRM plug-in indicates there is no paired site, and none of the DR setup options have been configured. Pair the sites together beginning on the vCenter server at the primary site. Basic SRM configuration is done by working through the options in the

Setup section of the main site Recovery Summary tab. To begin pairing the sites, complete the following steps:

1. Click the Configure link next to the Connection label and enter the address of the DR site vCenter server.
2. Enter the DR site vCenter username and password.
3. At the end of the configuration dialogs, verify that the connection was made successfully.

After the site pairing is complete, the primary site vCenter server will display the paired site addresses and "Connected" in the Setup box. Note that the SRM plug-in at the DR site will also display the sites paired in the opposite direction.

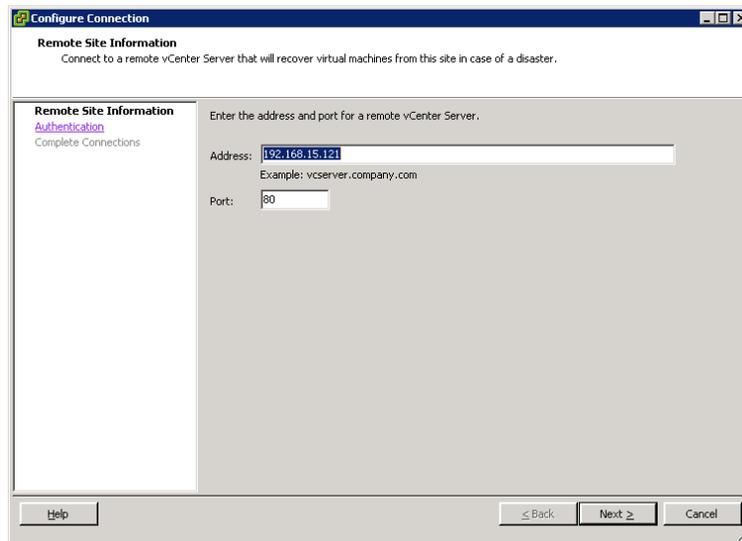


Figure 4 Pair protected and recovery sites using SRM

5.2 Configuring access to the EqualLogic storage

From the vCenter server in the primary site, click **Configure** next to Array Managers. The Configure Array Managers dialog appears. Click the **Add...** button to add an array.

Pre-requisite: Before configuring the array managers, it is recommended that the EqualLogic Auto-Replication is setup between the sites using the EqualLogic Group Manager and replication is configured for volumes.



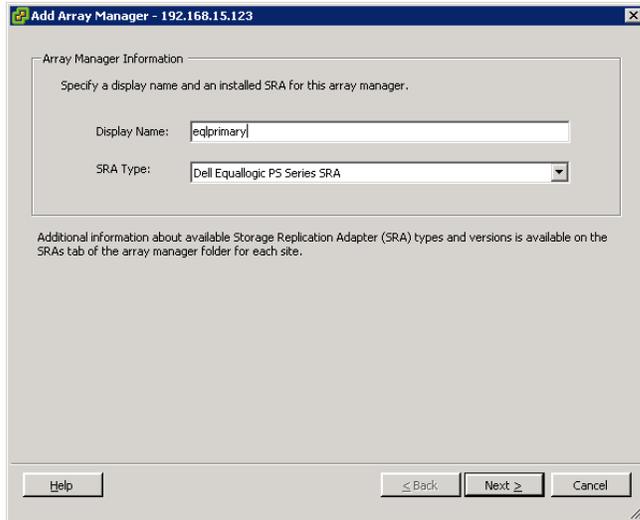


Figure 5 Configure SRA using SRM

Figure 6 shows how to add an EqualLogic group – If the vCenter Server (server with SRA installed) does not have access to the EqualLogic storage network (Group IP address), use the management interface IP address as shown below.

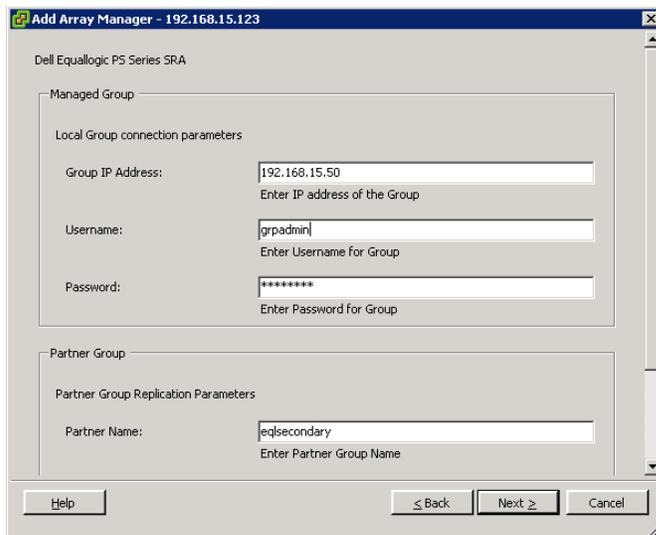


Figure 6 Adding Primary EqualLogic array using SRA

The figure below illustrates how to add the remote EqualLogic array (using the management IP).



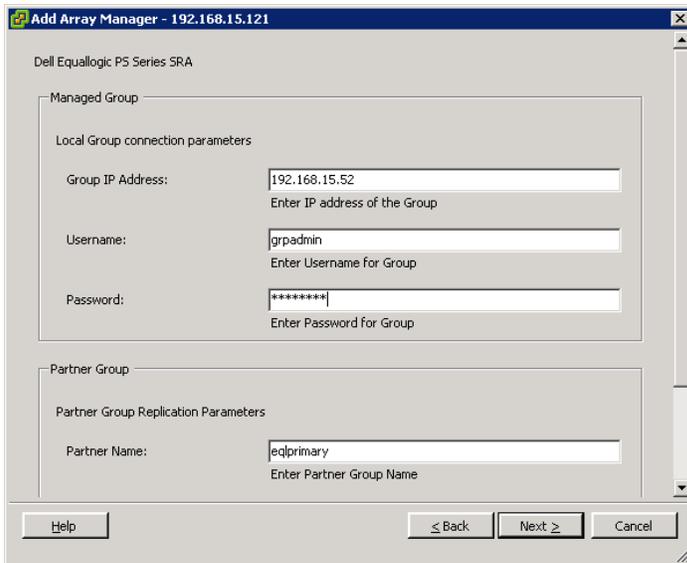


Figure 7 Adding remote EqualLogic array using SRA

Add an array on each site and enable the array pair on each site for use with SRM. This must be done once from any of the two sites. After enabling the array pair, the devices for enabled pairs are listed below.

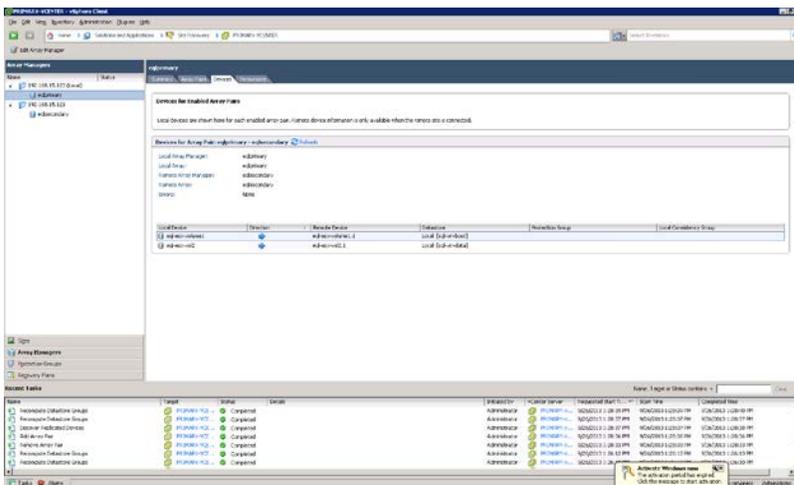


Figure 8 Displaying EqualLogic array pairs and devices

5.3 Use VSM to configure smart replicas

From the primary site vCenter, launch Dell EqualLogic Virtual Storage Manager to add a PS Series group (Use management IP addresses). Once VSM peers are configured and EqualLogic partners are defined in the VSM, application consistent point-in-time Smart Copy replicas on the partner EqualLogic storage array can be created.



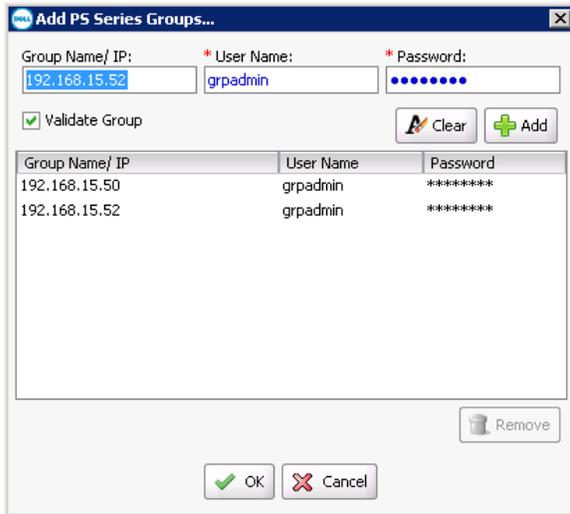


Figure 9 Adding primary remote EqualLogic Groups on VSM

5.4 Configure smart replicas of datastores using VSM

Once the VSM relationship is established between the two sites, configure datastores for replication using VSM. This is the datastore where the VM boot disks and SQL database volumes are present. By creating a smart replica of the datastore to the remote site, VSM quiesses the I/O on the VM to take a snapshot on the primary EqualLogic array and sends all changed data to the remote EqualLogic array.



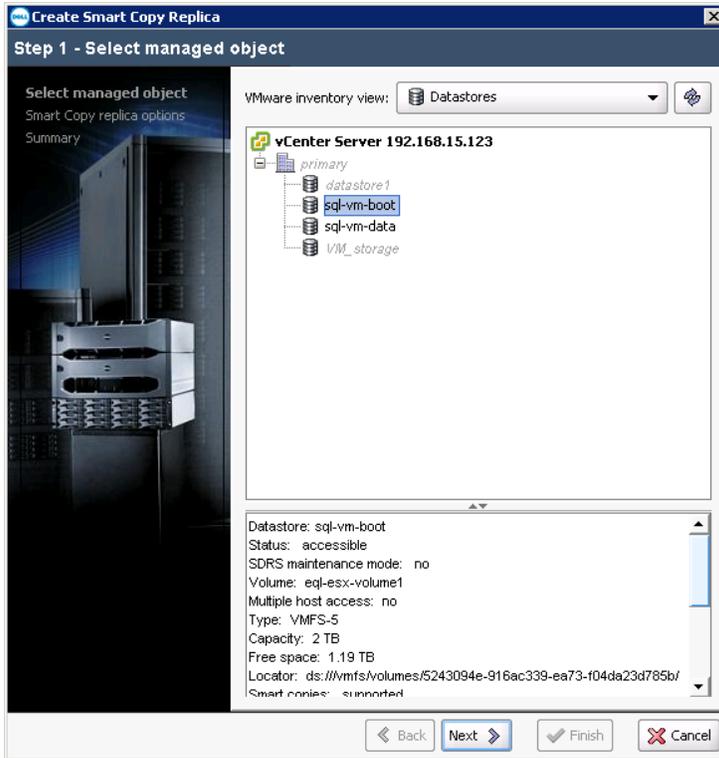


Figure 10 Selecting datastore for replication using VSM

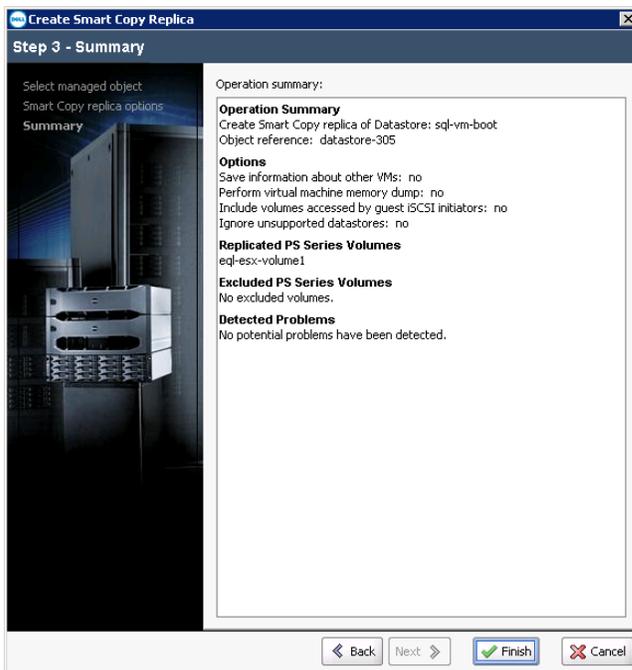


Figure 11 Creating Smart Copy replicas using VSM



This creates hypervisor-consistent smart replicas on the remote EqualLogic array. Create a schedule based on the change rate and the available WAN link between the two sites. This will ensure that consistent replicas are created at a point-in-time using VSM. There is no need to create any schedules from the EqualLogic Group Manager because the smart replicas are managed from the VSM. The following figures show how to create a schedule from the VSM plugin.

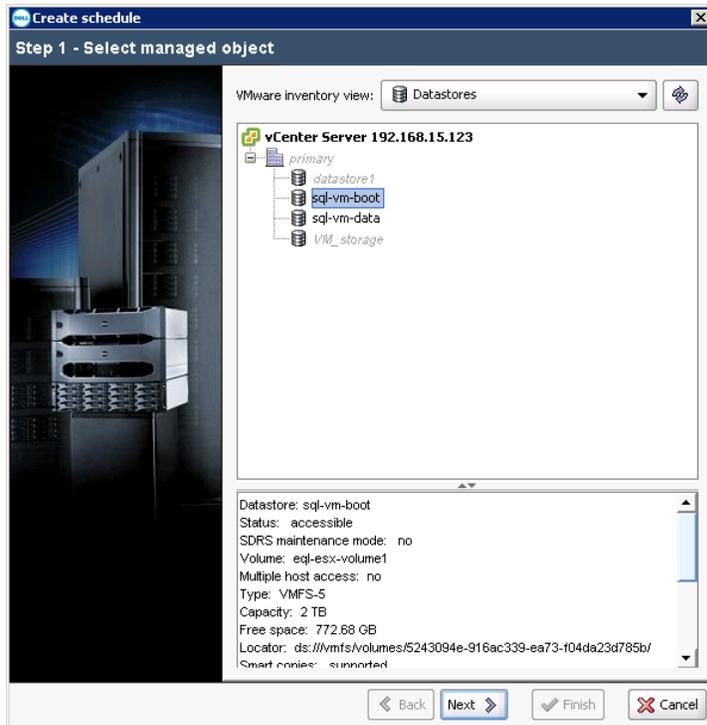


Figure 12 Creating schedule for Smart Copy replicas using VSM





Figure 13 Configuring schedule for Smart Copy replicas using VSM

Once smart copy replica schedules are created, VSM creates hypervisor consistent smart replicas on the remote EqualLogic Group. These replicas can be used to failover the SQL Server VMs using the SRM.

5.5 Build Protection groups

Protection groups allow VMs to be collected in groups that will be recovered together. Protection groups are created on the primary site vCenter server and identify which datastores and VMs are protected and will be recovered with SRM. In this solution, EqualLogic volumes are scheduled to be replicated using VSM. VMFS datastores should be the basis for grouping VMs because they can be easily configured with appropriate replication schedules.

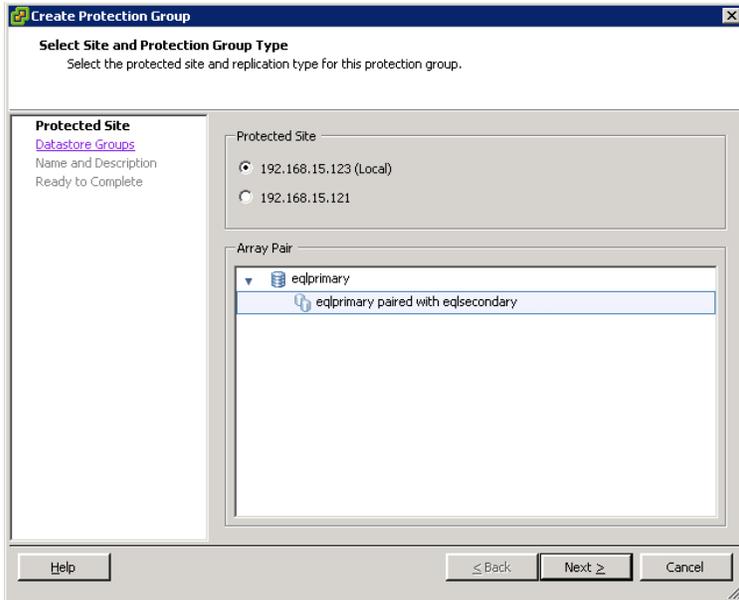


Figure 14 Building protection groups for VMs

Select the datastore where the SQL Server boot and data volumes are located in case of VSM configuration. If VM spans multiple datastores, they all will be included in the protection group.

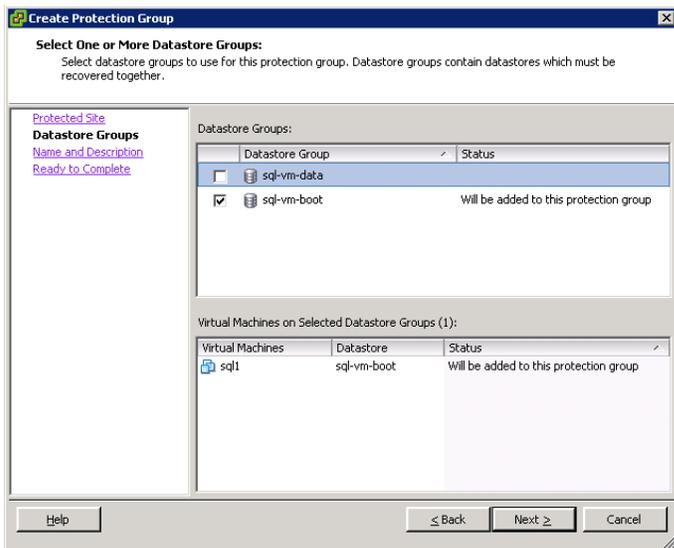


Figure 15 Selecting datastores for protection groups



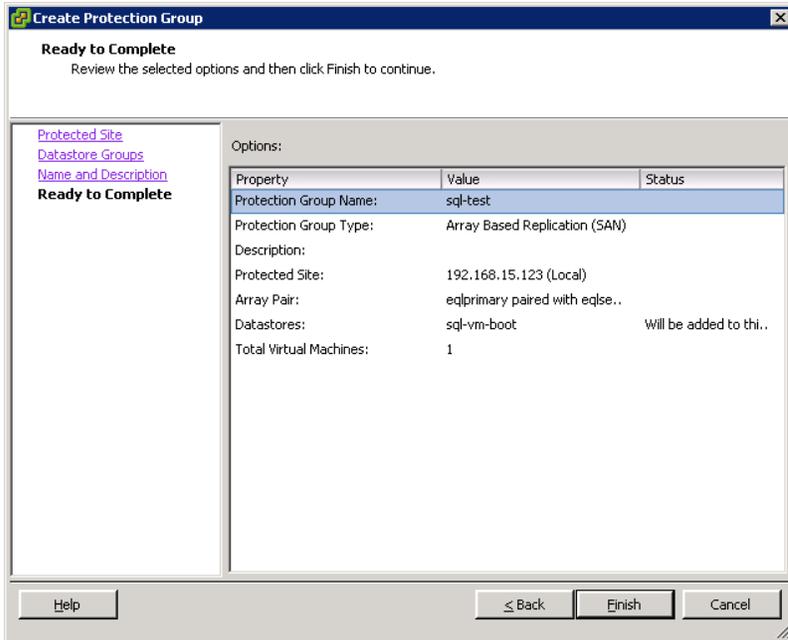


Figure 16 Creating protection groups for VMs

5.6 Build recovery plans

Recovery plans are created in the vCenter server at the DR site. Each recovery plan contains one or more of the protection groups created at the primary site. All of the protection groups contained in a recovery plan will be recovered by SRM when that plan is executed.

5.6.1 Create a new recovery plan

Select the recovery site from the protected site, and select the protection group for the recovery plan which contains all of the VMs to be failed over or failed back.



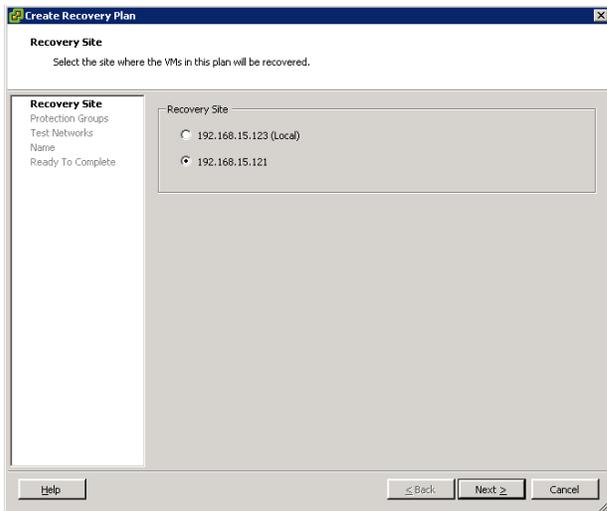


Figure 17 Creating a new recovery plan using SRM

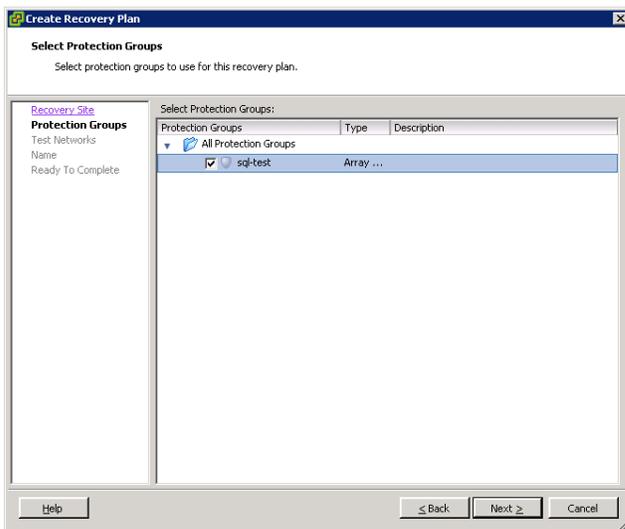


Figure 18 Selecting protection groups in the recovery plan using SRM

5.7 Execute the DR plan in test mode

When running the recovery plan in a test mode, the EqualLogic array on the remote site creates a clone of the replica that is presented to the other ESXi server on the recovery site. This volume is mounted on the remote ESXi server and the VMs are powered on. If a recovery plan consists of more datastores (volumes), a clone on the remote EqualLogic array will be created for every volume. This may not be feasible in all environments due to the additional space requirements required for creating clones on the remote EqualLogic array. To overcome this, we recommended creating multiple recovery plans with fewer VMs and test them one after another. During the test phase of the SRM plan, the VMs are connected to the DR test network rather than the public VM network and it does not impact the primary site operations or the replication



between the two EqualLogic sites. Clean up the test mode, once the DR testing is completed, this will delete the clones and the free space can be reclaimed on the EqualLogic storage pool.

5.8 Execute a recovery plan in a true DR scenario

Because different environments have different requirements for DR, this requires proper planning, documentation and testing of a DR process. Each environment can have a different DR plan, and there can be many variables in each DR plan such as DNS and Active Directory. This only shows an example of SQL Server failing over in a disaster to the remote site using SRM.

A disaster was simulated in the test by physically powering off the ESXi server and the EqualLogic storage on the primary site (site1), while running a TPC-E like workload using the Quest Benchmark Factory for Databases tool from the client to SQL Server. In the example, all resources running on site1 have been physically powered off. It was also verified that the connection to the primary ESXi server was no longer there and the VMs were no longer accessible in the network. Depending on the RPO that has been designed into the solution (five minutes in this test), some amount of in flight data loss can be expected during failover.

The RPO of five minutes was tested in this lab controlled environment, where both sites are connected in the same rack with a 1 Gbps dedicated connection. In most environments, this is not practically feasible as latency and other factors impact the WAN link and thereby the RPO. For better understanding on how EqualLogic replication behaves under various WAN conditions, refer to the [Dell EqualLogic Auto-Replication Best Practices and Sizing Guide](#).

When a real failover operation is initiated, SRM automates the following tasks:

- If the EqualLogic storage at the primary site is available and communicating to the remote site, it synchronizes (changed data since last replication) to the remote site EqualLogic array before failing over VMs.
- It shuts down the affected VMs at the primary site as assigned in the protection group configured in the SRM plan.
- Remote replicas are promoted to recovery volumes on the DR site EqualLogic array.
- SRM recovery connects the replicated datastores to the ESXi hosts at the DR site.
- It connects the VM network adapters to the appropriate recovery site network through the resource mappings.
- VMs will be powered on in the order (priority) defined in the SRM recovery plan.
- Finally, it executes any custom commands specified in the recovery plan.

5.8.1 Initiate recovery on the existing SRM plan

From the recovery plan, select "Recovery" to initiate a failover of all VMs being protected in the plan.



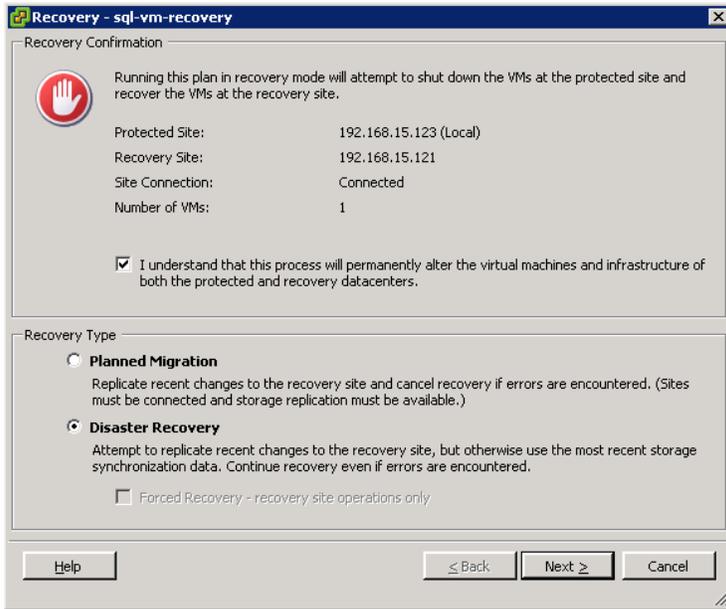


Figure 19 Executing the recovery plan in a DR scenario

Upon a successful failover, the EqualLogic replica set on the remote site is promoted to a R/W volume and presented to the ESXi host on the remote site. This also pauses the EqualLogic replication for that volume because the volume is promoted and mounted as a datastore on the remote ESXi server. The VMs in the protection group are restarted on the remote ESXi server. A planned migration can also be done when both of the EqualLogic systems are available and communicating. Connectivity to the remote site VMs and SQL server instances running on them was verified. The failover process in this test took seven minutes and 10 seconds to failover all four VMs to the recovery site. The failover times depend on various factors such as the number of VMs, the number of protection groups, WAN link latency, and other factors. The following is a sample output of SRM failover:

```

Plan Summary
Name:      sql-farm
Description:
Protected Site:      192.168.15.123
Recovery Site:192.168.15.121
Run Summary
Operation:      Recovery
Recovery Type:Disaster recovery
Started By:      Administrator
Start Time:      2013-09-19 15:13:26 (UTC 0)
End Time:        2013-09-19 15:20:36 (UTC 0)
Elapsed Time:    00:07:10
Result: Success
Errors: 0
Warnings:      0

```



5.9 Resynchronize primary site after recovery

Once the original primary site is back online, it is usually necessary to return operations back to the primary site. The first step is to re-protect VMs using the SRM. Re-protect involves setting up the replication in the reverse direction from the remote site to the original primary site.

From the standpoint of EqualLogic storage, when resynchronizing the environment back to the primary site, there are two high-level scenarios to consider:

1. Resynchronizing the environment if the primary storage has been recovered. This involves replicating the changes to the partner and establishing the replication process in the reverse direction by keeping replicas on the original primary site. This requires that the original primary EqualLogic and the DR EqualLogic volumes have a common snapshot copy (fast failback must be enabled for the volume on the original primary while setting up the replication) that uses a consistent point from which to resynchronize all changes. This helps faster failback of the volumes as only changes are replicated.
2. Resynchronizing the environment if the primary site storage was lost. In this case, the primary storage might have been destroyed or a common Snapshot copy no longer exists between the primary and DR volumes. This requires that the new volumes are created on the EqualLogic storage, and the entire volume must be replicated back to the primary site as if a new replication for that volume was being set up. This obviously takes longer because all the data should be sent to the primary site EqualLogic storage.

Resynchronizing after a DR requires the following processes:

- Recover or Replace/Rebuild the infrastructure (storage) at the primary site.
- Reestablish network connectivity between the sites and make sure SRM can see both sites.
- Build SRM relationship in the reverse direction.
- Perform a recovery to the original primary site (this requires additional downtime).

5.9.1 Resynchronizing the environment after the primary storage has been recovered

If a disaster did not cause complete loss of data from the primary EqualLogic PS Series array (for example, an extended power outage), recovering back to the original primary site is less time consuming. The first step is to build an SRM relationship in the reverse direction.

This example below shows how to re-protect the VMs, the DR site becomes the new protected site and the original primary site becomes the new recovery site.



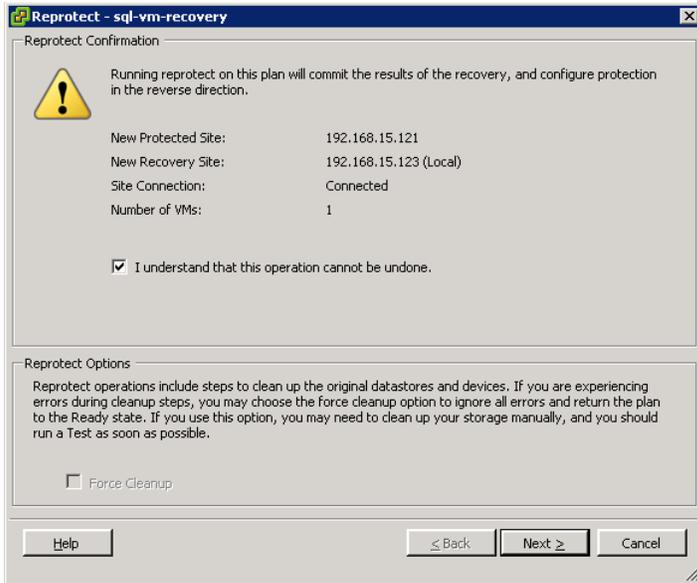


Figure 20 Re-protecting VMs

Re-protect operation replicates the recovery volume in the other direction – from the secondary EqualLogic array to the primary EqualLogic array.

Once the replication direction is changed and the VMs are running from the original DR site, all I/O operations to the SQL server can be resumed. Use the VSM on the remote site to setup a schedule for replicating the hypervisor consistent replicas to the original primary site.

Configure VSM on the remote site to schedule replicas to the primary site.



Figure 21 Create smart replicas using VSM



Using VSM scheduler, schedule the replication from the new protected site to the new recovery site.

5.9.2 Resynchronizing the environment after the primary storage is lost and rebuilt

In this disaster scenario, where there was complete loss of data on the primary EqualLogic PS Series array or complete loss of the whole infrastructure at the primary site, the recovery from the SRM perspective is identical except that new volumes must be created on the primary EqualLogic array and replicate the entire data from the DR site to the primary. This is more time consuming and complex as new volumes need to be created first on the primary site storage. If the complete data transfer takes longer with the existing WAN link, refer to the [Dell EqualLogic Auto-Replication Best Practices and Sizing Guide](#) for using the Manual Transfer Utility (MTU) for uploading the data without using the WAN link.

5.9.3 DR Failback

Re-establishing normal operations involves reversing the EqualLogic replication and SRM relationships again to establish the original primary-to-DR site replication and protection.

To return operations back to the primary site, SRM recovery is needed from the current protected site to the recovery site (original primary site). As new schedules are created from the VSM to replicate back to the original primary site, these replicas on the EqualLogic storage are promoted on the primary site and all VMs are turned on to bring business operations back to normal from the primary site.

From the recovery plan, select "Recovery" to initiate a failover of all VMs being protected in the plan.

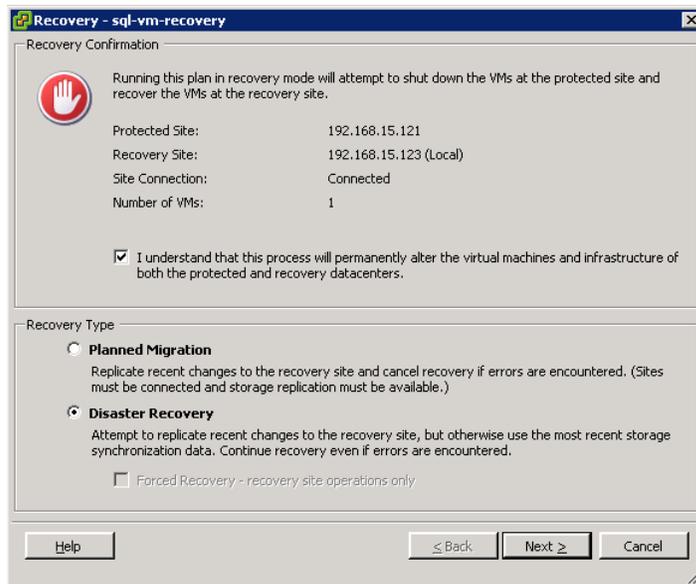


Figure 22 Running the SRM recovery plan



After it is failed over to the original primary site, check the SQL DB integrity. Now the original Primary becomes primary again, The next step is to re-protect to the recovery site to bring everything back to normal.

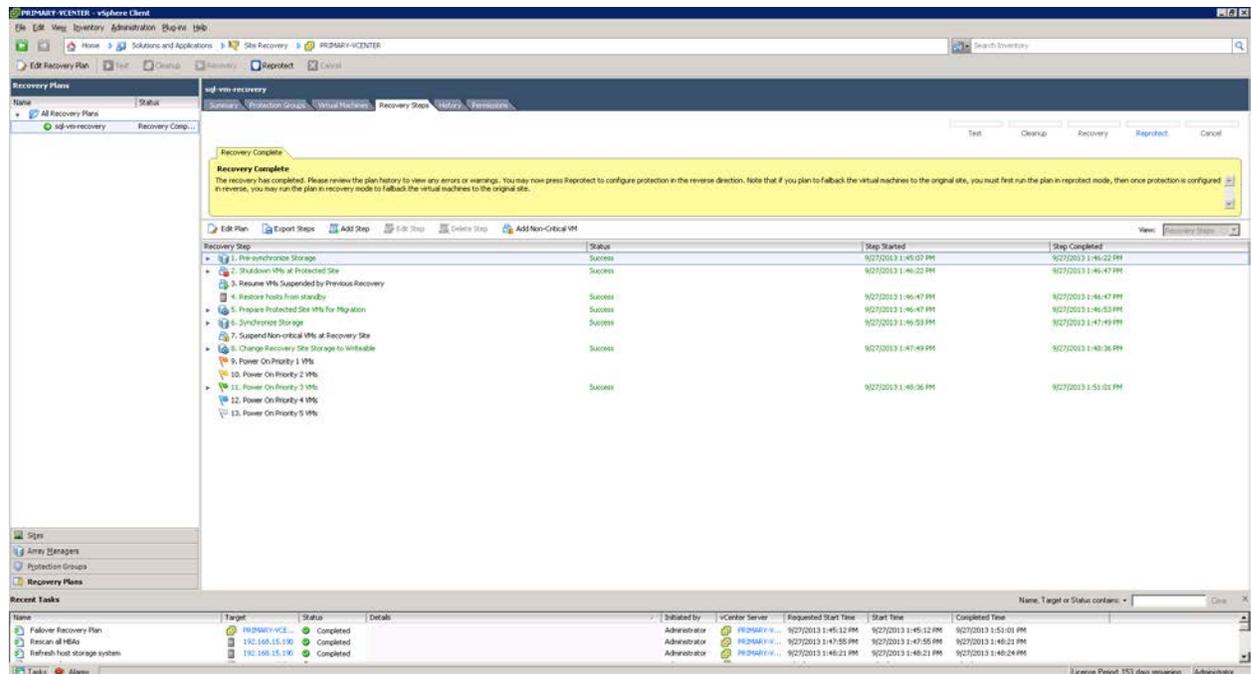


Figure 23 Completion of a successful site recovery

5.10 SQL test results

Using Quest Benchmark Factory for Databases, the "srm_test" database was populated with data using TPC-E like workload. In this test, the goal was to achieve an RPO of five minutes, while changes are updated to the remote EqualLogic array every five minutes through a scheduled task from the VSM. While running the TPC-E like load, the following SQL Query was executed to count the number of records in a table.

```
SELECT count(*) FROM [SRM_TEST].[dbo].[E_TRADE]
```

259200000

A disaster was simulated by physically shutting down the ESXi server on the primary site. Once the VM was unreachable, failover was initiated using SRM to the remote site. Upon failover to the remote site, it was verified that the guest OS was running on the VM and was accessible. Once the VM was accessible from the SQL enterprise manager the same query was executed to verify the number of records in the table.

```
SELECT count(*) FROM [SRM_TEST].[dbo].E_TRADE
```

259200000



All updates to the database within the last five minutes were lost because VSM scheduled replication happened every five minutes. This validated the last recovery point of the database upon a successful failover of VMs using SRM.

Note: RPO of five minutes is not always possible for all environments. This depends upon many factors including the size of the database, database change rate, WAN link and other WAN factors such as latency, packet loss, and others. For EqualLogic replication sizing and better understanding of the factors that impact RPO and RTO, refer to [Dell EqualLogic Auto-Replication Best Practices and Sizing Guide](#).

Once the SQL server was running from a VM on the remote ESXi server, SRM Re-Protect was enabled to reverse the replication direction from the remote to the primary site. To check the database integrity, the following tests were performed.

The records from the same table were deleted by using the following SQL query:

```
delete from SRM_TEST.dbo.E_TRADE where T_CHRG = 5 ;  
  
(61651155 row(s) affected)
```

The new records, before failing over to the original primary site, are:

```
select count(*) from SRM_TEST.dbo.E_TRADE ;  
  
197548845
```

SRM was used to fail back to the original primary site, and then same query was run on SQL server to verify the number of records in the table to be 197548845.

This also verified changes made to the SQL server after failover were replicated to the original primary site and that the database was active when it failed back.



6 Scenario 2 - SRM using EqualLogic replication with VSM and ASM/ME

ASM/ME enables you to create fast, space-efficient, point-in-time copies of EqualLogic volumes as part of a data protection strategy. You can quickly backup and restore EqualLogic volumes on multiple Windows machines and also manage multiple hosts from a single GUI.

ASM/ME is a component of the Host Integration Tools Microsoft Edition (HIT/ME), and it is installed with HIT/ME. ASM/ME has advanced operations and management capabilities that offer DBAs additional functionality other than day-to-day protection. These operations may require some manual operations that are not automated by ASM/ME.

SRM can be used in conjunction with ASM/ME to achieve application (SQL) consistent failover and failback of VMs to and from the DR site. This option provides flexible and granular database protection. You can use Smart Copies (snapshots), which provide application consistent snapshots from which the database can be recovered upon any corruption or human errors. ASM/ME provides a GUI to recover the database in place or revert the database to a point-in-time copy using the snapshots. Smart clones provide application consistent clones of database volumes, which allows duplicating database instances very quickly.

ASM/ME must be installed on the Windows VM running SQL server. Database volumes for SystemDB, TempDB, logs, and data all should be residing on separate EqualLogic volumes while the boot disks for VMs come from another EqualLogic volume via VMFS datastore. The datastore hosting boot disks for VMs is replicated using VSM and ASM/ME is used to replicate SQL specific volumes.

To provide flexibility and granularity of recovery in a virtualized infrastructure, it is critical that the backend storage is configured appropriately. Figure 24 shows the layout of VMs running Windows Server 2008 where ASM/ME is configured to protect the SQL database.



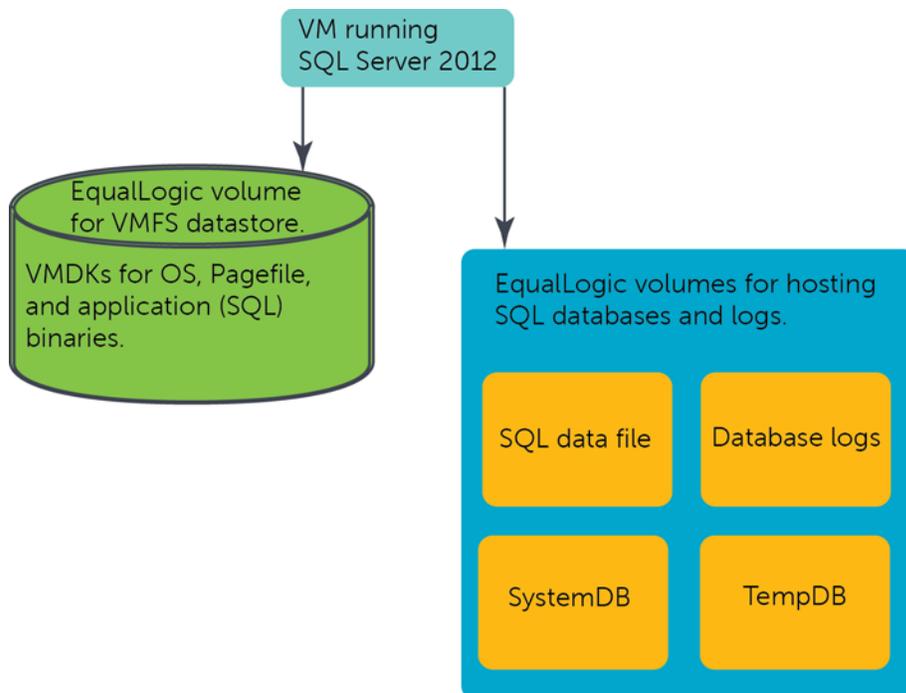


Figure 24 VM datastore layout for ASM/ME

6.1 SRM testing with ASM/ME

This solution is validated on four VMs running different instances of SQL Server 2012 each configured with the similar EqualLogic storage configuration. In this configuration, the boot disks for the VMs are created on a VMFS datastore on the ESXi server, while SQL specific volumes are created separately on the EqualLogic storage and presented to the VM guest OS and accessed directly via the Windows iSCSI initiator. One EqualLogic volume is presented to the ESXi host that is used for boot disks for all VMs. Four volumes are created on the EqualLogic storage for each VM to run a SQL database. Replication of these volumes is configured, but no replicas or replication schedules are created because they are managed by the EqualLogic tools such as VSM and ASM/ME. When replication is enabled for these volumes, enable the feature **Keep failback snapshot** because it provides a quick mechanism to resynchronize the changes back to the primary EqualLogic storage upon failure. This adds additional capacity requirements on the primary EqualLogic site to keep a base snapshot.

6.2 ASM/ME configuration for SQL Server

1. Install and configure 64 bit EqualLogic Microsoft Host Integration Tools (v4.6). This will install the EqualLogic MPIO driver for Windows and Auto Snapshot Manager for Windows.
2. Manage multiple Windows hosts running ASM/ME using a single interface from a single host. All SQL instances running on different VMs can be managed from a single VM.
3. Configure the MPIO driver to discover the newly added volumes and log on to iSCSI targets.
4. Format new disks from the Windows disk management tool with NTFS file system.
5. Install SQL server and select these volumes (as per table 4) to host different databases.

6. Once SQL is installed and configured, create a new test user database srm_sql1 and populate data using Quest Benchmark Factory for Databases tool with TPC-E like workload.
7. Configure ASM/ME on the Windows VM. ASM/ME needs to access the VSS control volume from the host. From the ASM, configure access to the VSS control volume.
 - a. Add PS Group from the settings.

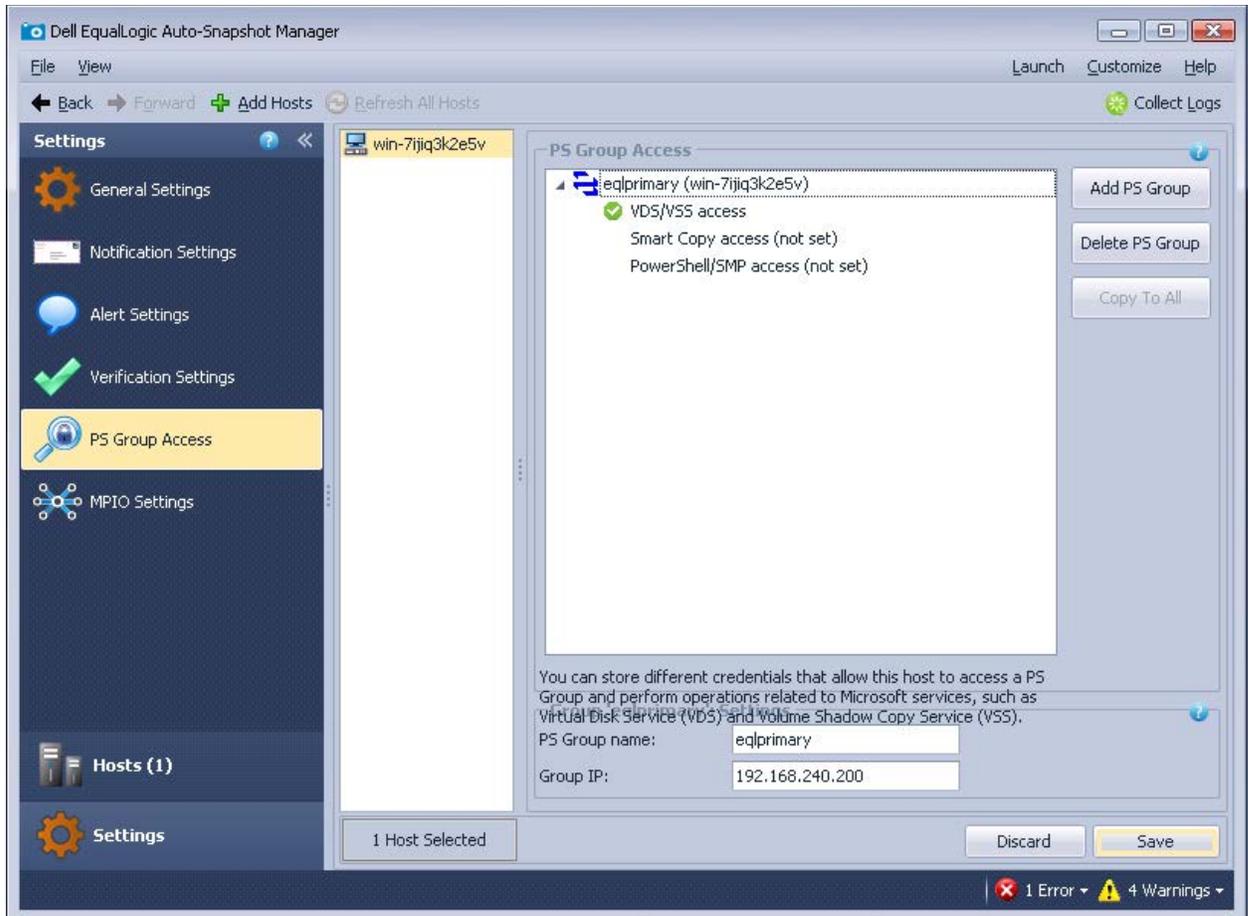


Figure 25 Adding EqualLogic PS Group to be managed by ASM/ME

- Configure VSS/VDS access. First create a local chap user account from the EqualLogic Group Manager and enable VSS/VDS management access to the group.

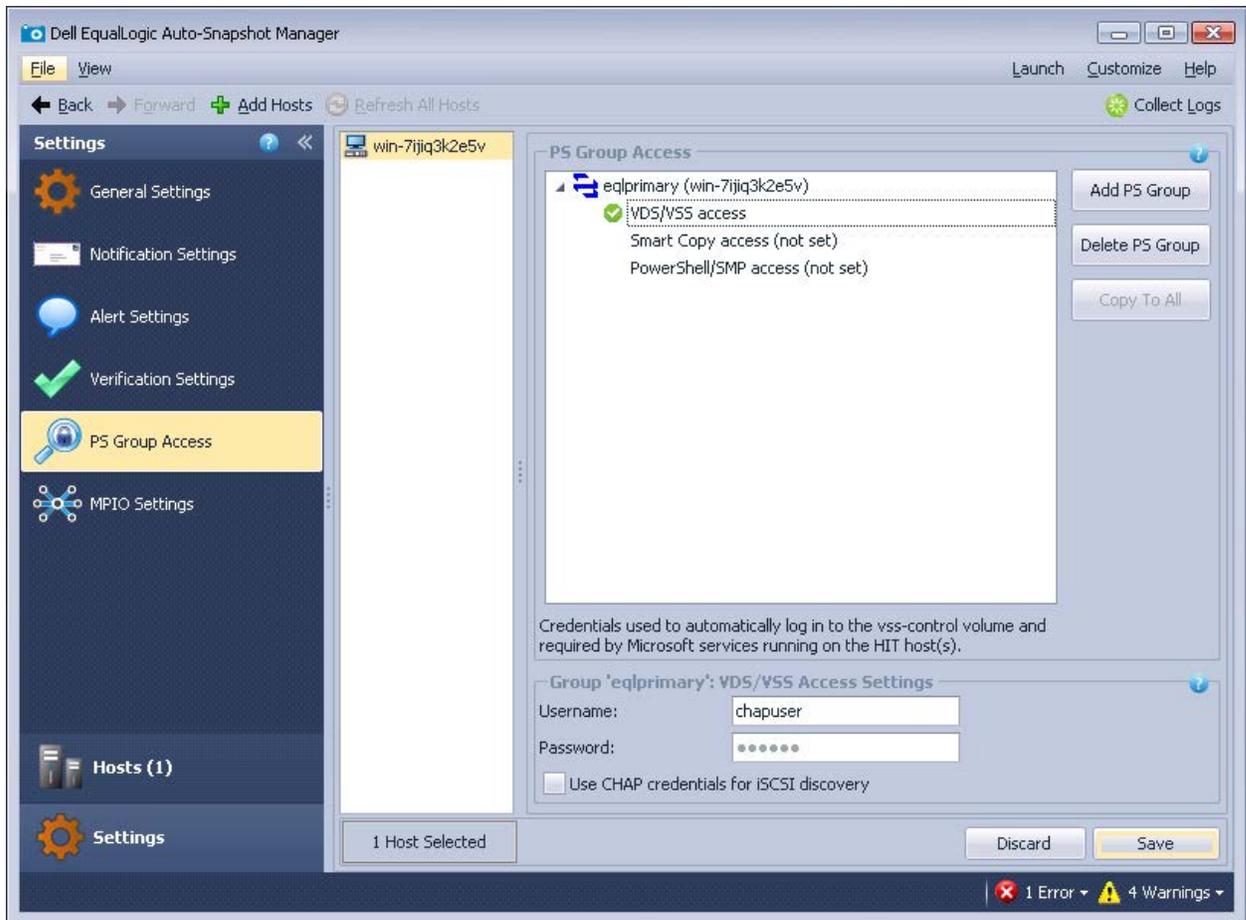


Figure 26 Configuring VSS/VDS access from the host using ASM/ME GUI

Once VSS/VDS access is configured, the host displays all iSCSI connections and SQL databases that are part of ASM/ME data protection.

6.2.1 ASM/ME with SQL Server overview

ASM/ME uses Microsoft's VSS architecture to provide application integration with EqualLogic SAN copy operations. During the VSS operation flow, Auto-Snapshot Manager initiates the process by requesting the SQL Server VSS Writer to prepare a database for a Smart Copy operation. The SQL Server VSS Writer component places the database in a consistent state and the PS Series VSS Provider service initiates the SAN copy using PS Series hardware snapshots, clones, or replication functions. The end result is a data-consistent point-in-time "Smart Copy" of the SQL Server database and volumes. Smart Copies can then be used to fully restore a database or simply recover object level data using various recovery options available to the Smart Copy set.

6.2.2 Snapshot Smart Copies

Snapshots created from ASM/ME provide a SQL transaction consistent point-in-time copy of SQL Server databases. Snapshot Smart Copies can be used to restore the database to a point-in-time, restore the



database as “new” side-by-side with the old database, or mount just the database volumes as read-only for recovering any necessary files. This also provides the capability to selectively restore a single database from a group or databases on the same host. These operations will revert back the database to a point-in-time taken by the snapshot. Smart Copies create array based snapshots and these smart copies require additional space allocation (snapshot reserve) in the EqualLogic volume to store any data changed after a snapshot was created. It is important to monitor reserve utilization and tune the reserve setting for optimal storage utilization. EqualLogic SAN HQ helps to monitor the snapshot reserve space and this should be adjusted to meet the snapshot retention requirements.

6.2.3 Clone Smart Copies

Clone Smart Copies create a complete new copy of the database. This allows you to quickly deploy databases in a testing environment. Clone Smart Copies are exact duplicates of the original volume or volumes including all the data and full size of each volume included in the Smart Copy. Clone Smart Copies are most useful when the original volume or database environment needs to be recreated such as test or development scenario.

With ASM/ME, Smart Copy clones can be used to Restore “All as New” or side-by-side restore operations to duplicate production environments or create exact copies of database environments for testing and development scenarios. In addition, these database copies can also be used to offload operations from the production database such as reporting or data mining.

6.2.4 Replica Smart Copies

The Replica Smart Copy option is available only if replication is configured in the PS Series group for volumes that make up a SQL server database. Replicas are snapshots that are sent from one PS Series group and stored on another PS Series group, and hold only the changed data from the last replica operation. Database volumes display the error “replication partner not available” under ASM/ME when the VM is not able to access the remote EqualLogic group. This is fine because the VM is not required to access the remote EqualLogic site in many cases when there is a failover using the SRM upon a disaster. If something happens to the local EqualLogic site, the VM will not be in operating condition because the VM is booted from a datastore that resides on the primary EqualLogic.

The masterDB and user databases are protected because they reside on unique EqualLogic volumes which are not part of the VMFS volume, so the first step from ASM/ME is to create a collection group for the SQL host to be protected.



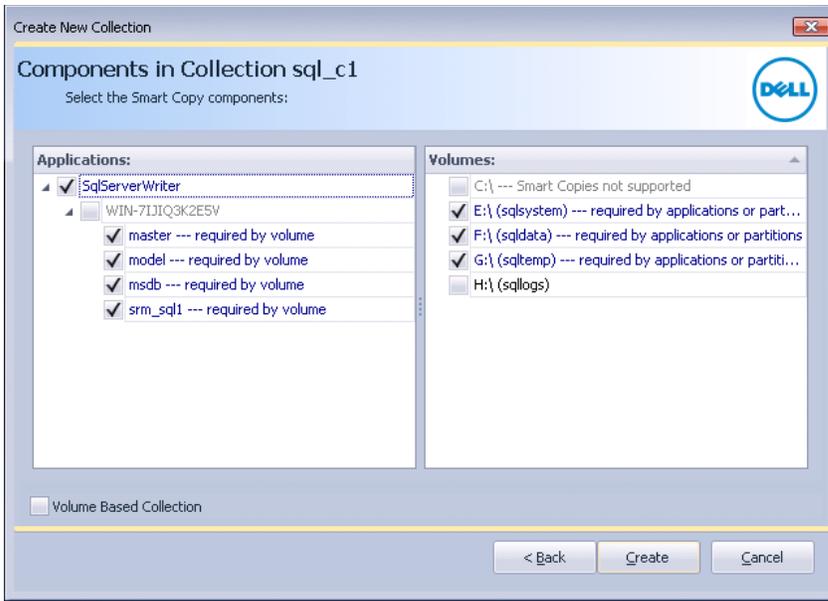


Figure 27 Creating a new collection group for all SQL components

Create a Smart Copy set (replica), with multiple components. This maintains consistency across all EqualLogic volumes which are part of the SQL instance.

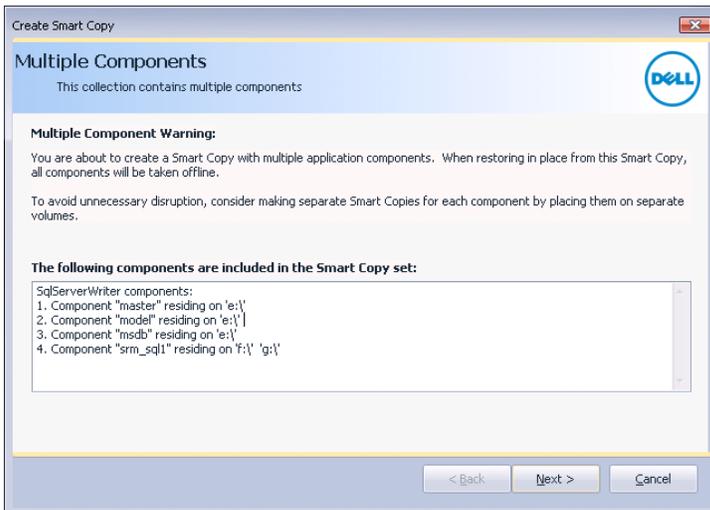


Figure 28 Selected components for the Smart Copy replica



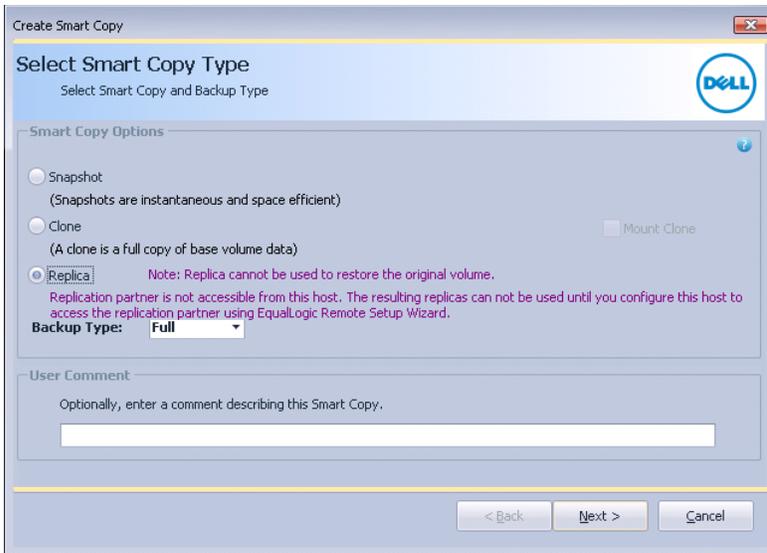


Figure 29 Selecting Smart Copy replica for SQL volumes using ASM/ME

Create a Smart Copy replica of the Smart Copy set that contains user and system databases.

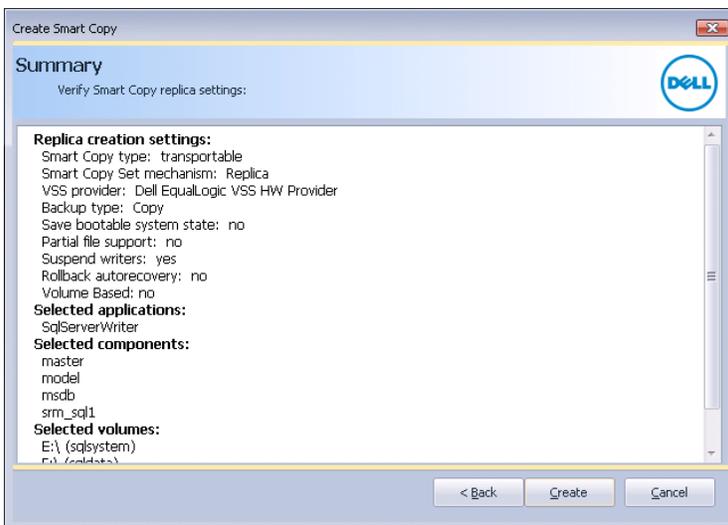


Figure 30 Creating a Smart Copy replica for the collection group



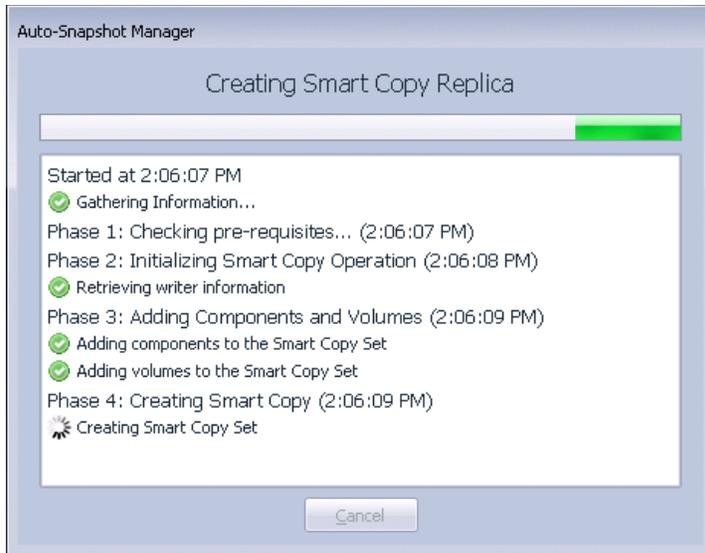


Figure 31 Sample output of Smart Copy replica creation

All VMs running their own SQL Server instances can be managed from a single host via ASM/ME. This creates application consistent Smart Copy replicas on the remote EqualLogic storage. This excludes TempDB, because TempDB is not required to be part of replication. When SQL Server is restarted it flushes everything on the TempDB and restarts cleanly. The smart replicas can be unreachable if the host does not have connectivity to the partner EqualLogic group (This is typical because VMs on the primary site do not require access to the EqualLogic storage network on the remote site).

Note: SQL Server database Smart Copies have two options for backup types that determine the behavior of the Smart Copy operation on a SQL Server database. The following backup types are available with ASM/ME:

Full – This backup type notifies the SQL Server that there was a backup operation. The operation applies a checkpoint and timestamp to the database log file. For SQL Server environments that use log file backups, this backup type allows log backup files to be applied to the restore process to add additional granularity during the restore process.

Copy – This backup type creates a copy of the volume or SQL Server database, and specifies an out-of-band backup operation that has no effect on application log files or backup operations. This backup type is supported only with SQL Server 2005 and later.

To configure a new schedule for creating smart replicas, From the collection group created, right click to create a new schedule for creating point-in-time replicas at the remote site.

Select **Collection** to create a new schedule.

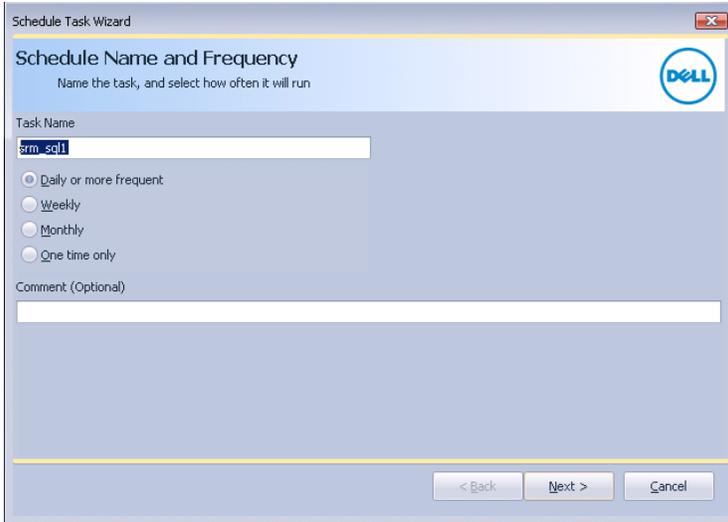


Figure 32 Create ASM/ME Smart Replica schedule

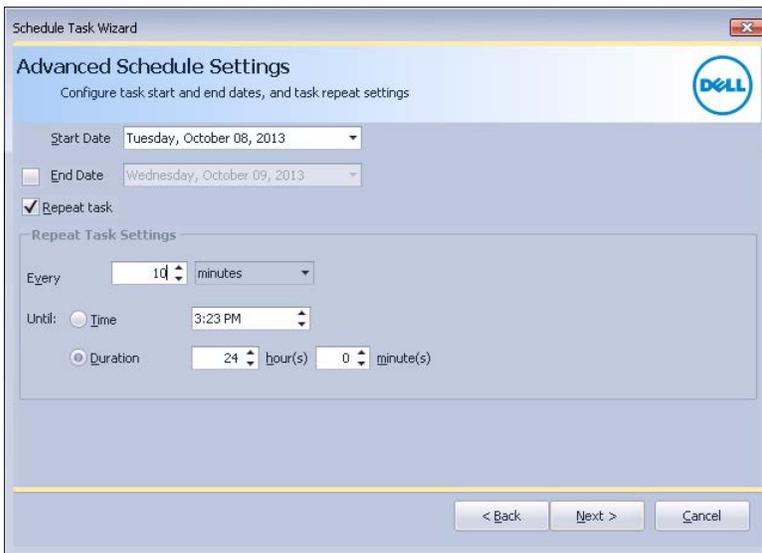


Figure 33 Configure ASM/ME Smart Replica schedule



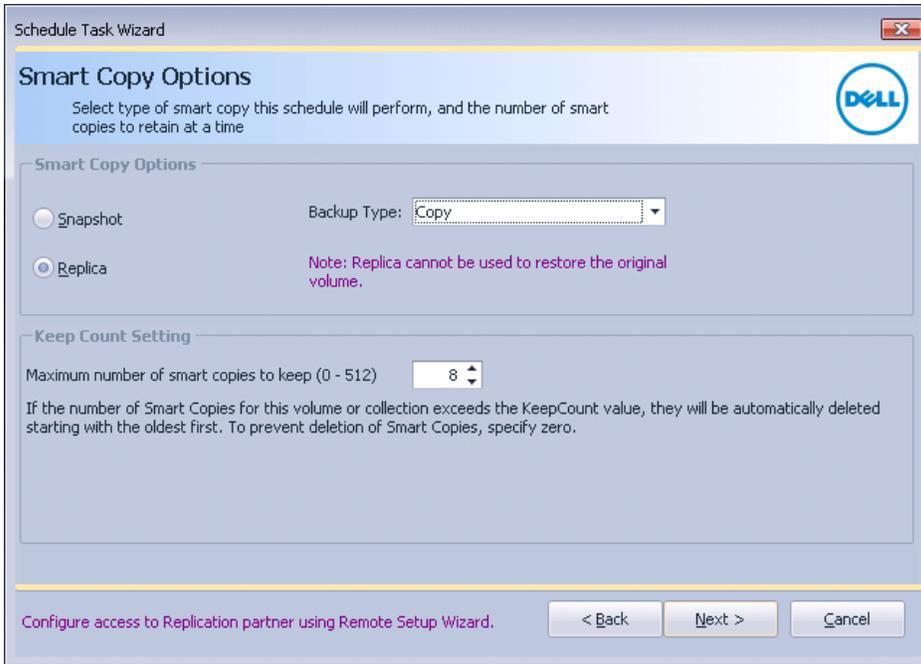


Figure 34 Selecting Smart Copy replicas in the schedule

Smart Copy replicas for the boot disk datastore volume must be configured from the VSM in addition to the ASM/ME replicas for the database. Do not include PS Series volumes accessed by iSCSI initiators because these are replicated using ASM/ME.



Figure 35 Smart Copy replicas of datastore(s) using VSM



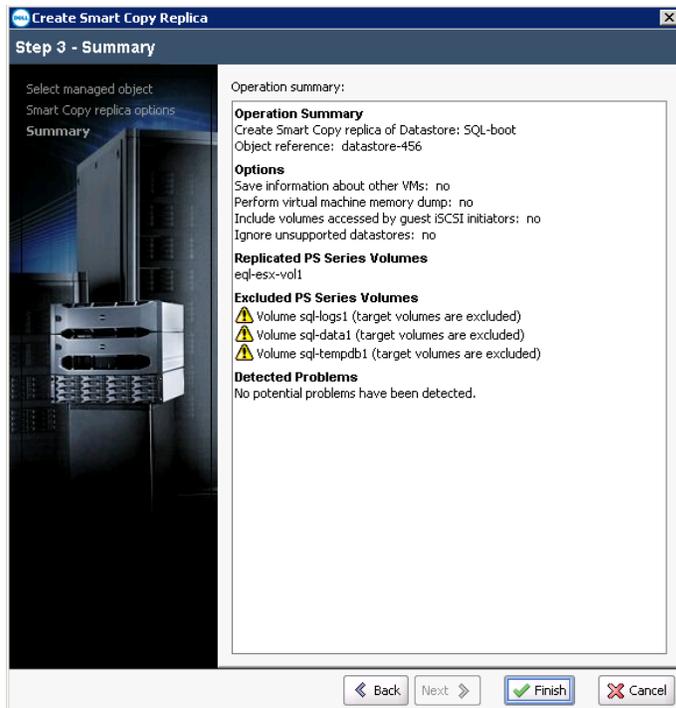


Figure 36 Creating Smart Replicas of datastore(s) using VSM

When using ASM/ME on the host, VSM cannot quiesce the VM residing on the datastore. For VSM to take a hypervisor consistent snapshot, additional steps are required as per the VMware article below.

The guest OS has reported an error during quiescing. The error code was: 5 The error message was: 'VssSyncStart' operation failed: IDispatch error #8449 (0x80042301)

It fails to create a snapshot and quiesces the VM when ASM/ME is configured on the VM. To create VM snapshot and quiesce the virtual machine, the following steps need to be done. Refer to the following VMware article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1028881

1. Open C:\ProgramData\VMware\VMware Tools\Tools.conf in a text editor. If the file does not exist, create it.
2. Add these lines to the file:
3. [vmbackup]
4. vss.disableAppQuiescing = true
5. Save the file and exit the editor.
6. Restart the VMware Tools Service for the changes to take effect:
7. Click Start > Run, type services.msc, and click OK.
8. Right-click the VMware Tools Service and click Restart.



6.3 Execute recovery plan in a true DR scenario

Using SRM, create a recovery plan to include all four VMs and perform a recovery of these VMs on the remote site.

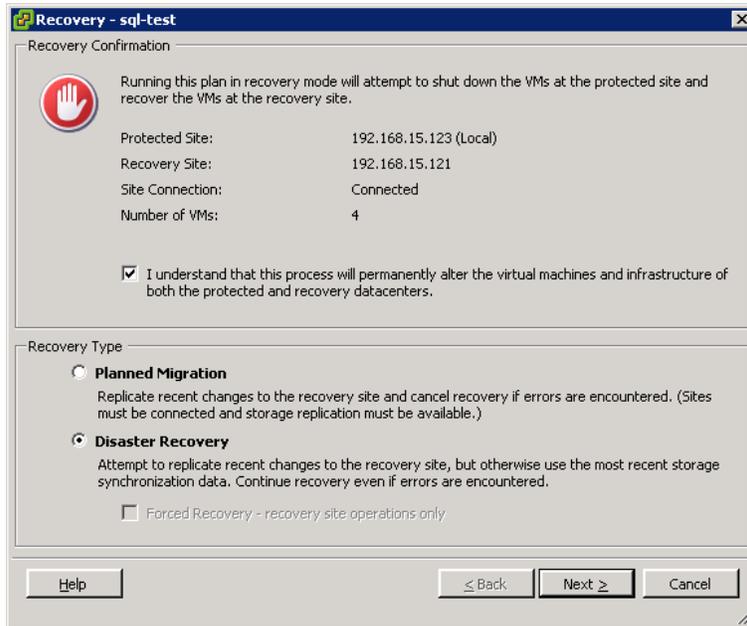


Figure 37 Initiating site recovery to the recovery site

Once VMs are failed over and restarted on the remote site, some manual steps outlined below must be performed to restart SQL instance as SQL database volumes are missing from the VM configuration.

6.3.1 Steps to recover SQL Database

1. Map virtual NICs on the primary ESXi server to the remote ESXi server for SAN access to the remote EqualLogic array. Complete this step while creating the protection group for VMs by mapping the network resources with the remote ESXi server.
2. From the EqualLogic Group Manager on the remote site, promote the replica set to a volume. Make the volume promotion permanent because ASM/ME cannot replicate on recovery volumes. If the primary EqualLogic site is not completely lost, these volumes should be demoted to replica sets.
3. From the EqualLogic Group Manager, assign SQL volumes to Windows VMs.



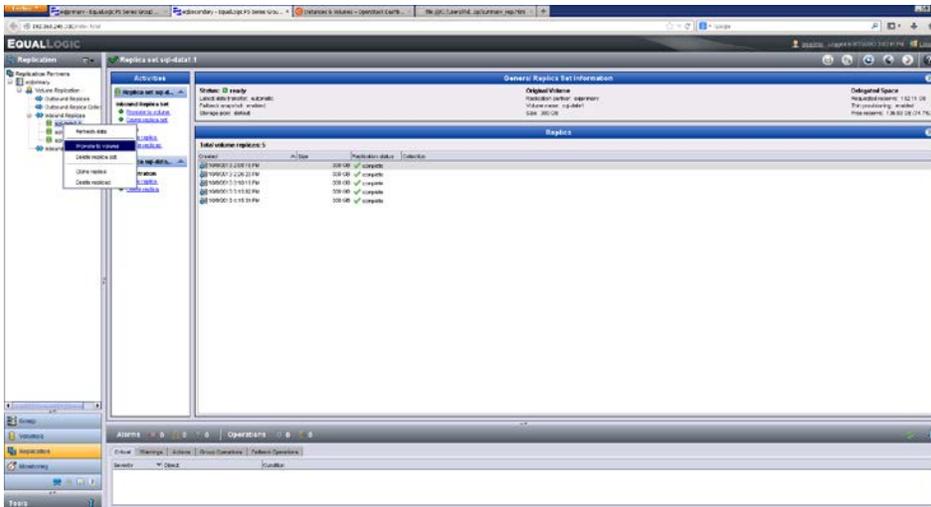


Figure 38 Promoting SQL replica volumes from EqualLogic Group Manager

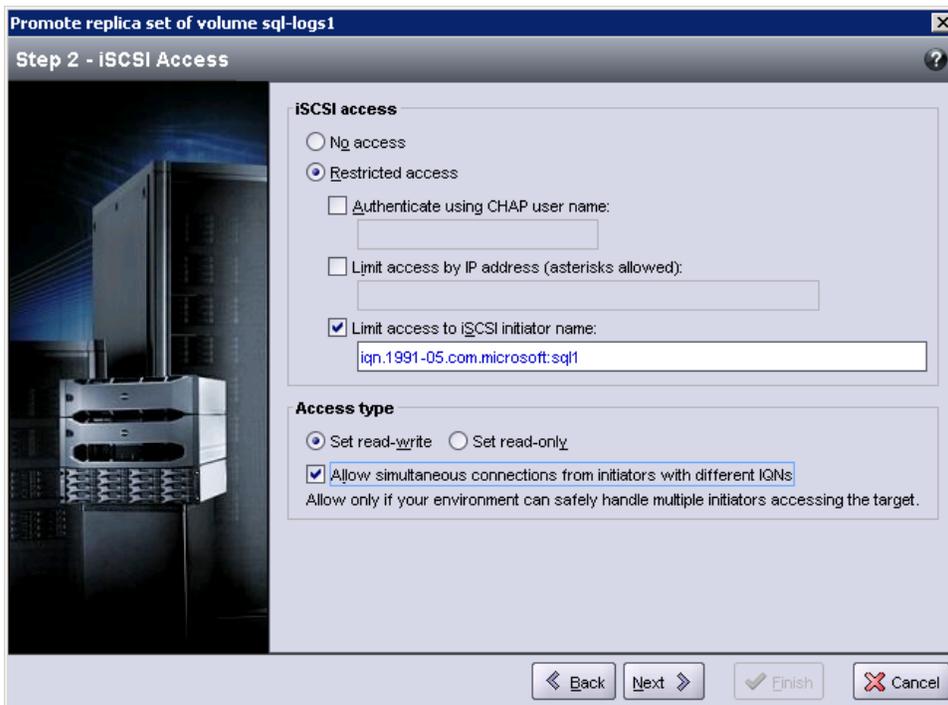


Figure 39 Mapping promoted volumes to VM

4. Make sure the windows VM networking is setup correctly to access the remote EqualLogic storage backend. Setup VM NICs to access the new EqualLogic group on the DR site.
5. Remove existing iSCSI connections from the initiator, and add the new EqualLogic Group Manager IP address. Make the volumes persistent.
6. From the Windows host iSCSI initiator, discover the new EqualLogic Group and logon to the EqualLogic volumes.



7. Since TempDB is not replicated, create TempDB on the DR EqualLogic. Assign the same drive letter as the primary site.
8. Once all volumes are discovered, restart the Windows VM to make sure SQL starts properly.
9. Once the database is restarted, enable replication using ASM/ME (schedule) in the reverse direction to the original primary site.

6.4 DR test results using ASM/ME

Using the Quest Benchmark Factory for Databases tool, "srm_sql1" database was populated with data using TPC-E like OLTP type database. In this test, the goal was to achieve a RPO of five minutes, because changes are updated to the remote EqualLogic array every five minutes via a scheduled task from the ASM/ME. The boot disk datastore is also replicated every five minutes using the VSM scheduler.

Before the disaster simulation, the following query was run to verify the number of records in a table.

```
SELECT count(*) FROM [SRM_TEST].[dbo].[E_TRADE]

323200000
```

At this point, the database was deleted to simulate a disaster. Upon on failover to the remote site, as shown in the last section, it was verified that the guest OS was running on the VM and was accessible. Once the VM was accessible, following the manual steps to restart the database, the following query was run to make sure the same number of records were in the table.

```
SELECT count(*) FROM [SRM_TEST].[dbo].E_TRADE

323200000
```

6.4.1 DR failback using ASM/ME

To bring business operations back to normal once the primary site is restored, reverse the SRM relationship. The first step is to re-protect the VMs from the DR site to the original primary site. Now the DR site becomes the new SRM protected site and the primary site becomes the new SRM recovery site. The next step is to configure ASM/ME to add the remote EqualLogic Group Manager and configure VSS control volume access to the remote EqualLogic array. Then create Smart Copy replicas from ASM/ME for the database instance. Also schedule VSM on the remote site vCenter to replicate datastore changes back to the primary site and ASM/ME Smart Copy replicas on the reverse direction from the DR site to the primary site. Once we have consistent replicas of both datastore volumes and database volumes on the primary site, schedule a downtime to failback to the primary site using SRM. Perform the same manual steps outlined before, like permanently promoting SQL replicas, recreating TempDB and restarting SQL Server.

In this SQL database recovery test, deleted the srm_sql1 test database was deleted on the DR site before failing back to the primary. Before the changes are replicated to the primary site, it failed back to the primary site using SRM and manual steps to recover the SQL instance. Then the sql_srm test database was attached from SQL enterprise manager, and SQL queries were run successfully.



7 Supported storage replication topologies with SRM

The testing used a single EqualLogic group that contained two members with a single pool configured at both sites.

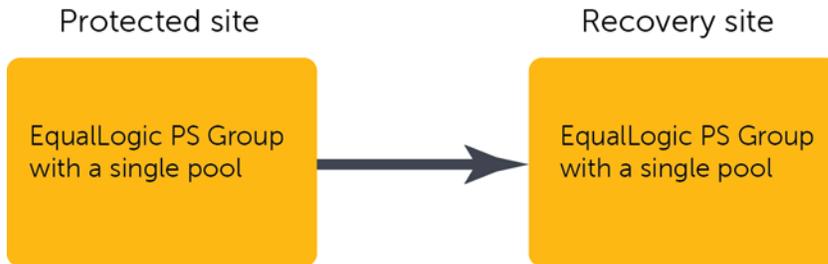


Figure 40 EqualLogic replication in a single direction: Supported EqualLogic replication layout for SRM

However, there may be scenarios where multiple EqualLogic groups or multiple pools are deployed within a same group. These topologies require additional considerations while using SRM. When setting up the EqualLogic replication between volumes, the delegated space must be configured from a single pool within an array for a single primary EqualLogic group. This can cause some issues upon failover because a single pool can only accommodate 1024 iSCSI connections and single pool can have different I/O characteristics than the primary site with multiple pools. Avoid running into the max connections limit on the failover of an EqualLogic array to avoid any potential performance issues. To avoid the limit, create an EqualLogic group for every pool to be replicated from the primary site. This may require some additional steps while configuring the SRM because more than one storage array will be configured. The solution can have different EqualLogic PS Series arrays at the primary and DR sites with different performance and cost characteristics.

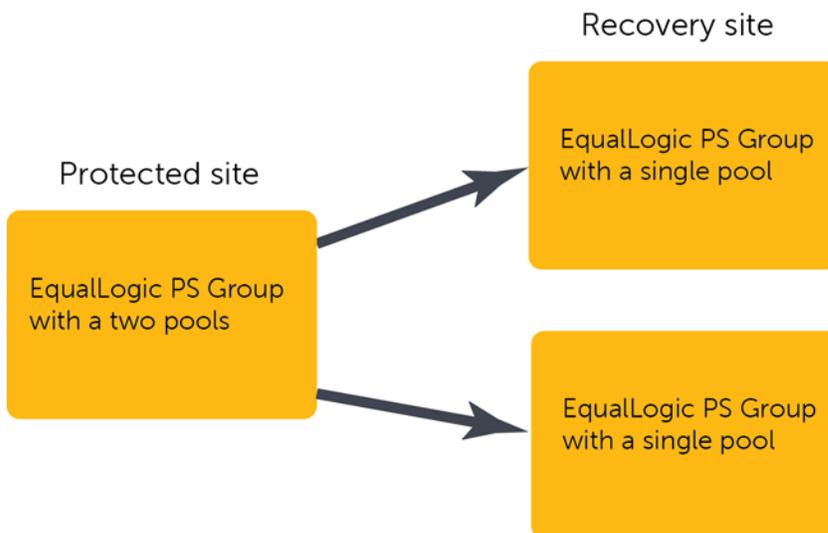


Figure 41 Supported EqualLogic replication layout with multiple pools for SRM



7.1 Additional best practices and considerations

This section describes optional steps involved in configuring SRM and SQL Server.

7.1.1 TempDB and OS page file considerations for replication

Replicating TempDB can flood the network traffic especially in certain database workflows. If you see network traffic across your replication network other than the database changes, consider options not to replicate the TempDB. The TempDB is a work area for the database to temporarily store data until a process is complete and the data can be written to the log files. Nothing is stored permanently on the TempDB drive. When you restart SQL Server it will also clean out the TempDB and start clean. Additional steps are required if you do not replicate the TempDB from SRM perspective because TempDB is not part of replication.

1. On the DR site, manually create a datastore for holding the TempDB for SQL instance.
2. From the DR Virtual Center, create a VM and name it the same as the production SQL server.
3. Create two virtual drives, one for the Windows OS and the other for TempDB. Install Windows on this new VM and format the second drive with NTFS file system with the same drive letter as the production SQL Server TempDB drive.
4. If any other folder structure is used on the production SQL server for the location of TempDB, recreate it here.
5. Now shut down the VM, open the settings and remove the second Hard Disk from the VM. Now that the drive is detached from the VM, right click the VM in inventory and click **Remove from disk**. This cleans up all the files in the folder on the Datastore with the exception of the detached drive that is needed.
6. When configuring the protection groups for VMs, on the Storage tab, you will now see a warning for the TempDB drive. This is expected because it was not replicated to the DR site.
7. Click the Hard Disk that has the warning on it, click **Browse**, and browse to DR Datastore where SQL TempDB is residing as created before.

This can have a replicated SQL server that is protected by SRM with minimized replication traffic for this VM by excluding the TempDB volume.

If you are using ASM/ME for replication, you can ignore replicating the TempDB volume for replication. Upon failover, a new volume must be created for TempDB on the remote EqualLogic array and formatted with the same drive letter as the primary SQL Server and attached to the VM using the manual steps outlined before in section 5.3.1.

7.1.2 Protect VMs with non-replicated Windows page files

While vCenter SRM allows you to replicate transient data, such as Windows paging files or VM swapfiles, such data need not be replicated. You may want to prevent replication of such data to avoid unnecessary consumption of network bandwidth. You need sufficient memory for SQL Server on the VMs to avoid excessive paging which results in unnecessary replication when the page file is also on the same location



as the OS boot disk datastore. If memory is a concern, avoid replicating the changes caused by paging by separating the page file on Windows to another datastore on EqualLogic that is not setup for replication.

By default, Windows VM hosts store the virtual memory pagefile on the system drive, but they can be configured to store the pagefile on a second virtual disk. This virtual disk can be stored in a non-replicated datastore on EqualLogic storage. However, to prevent Windows from placing the pagefile back in the default location when it is recovered at the DR site, additional steps must be taken in the SRM configuration.

SRM can be configured to connect a pre-existing VMDK to a recovered VM at the DR site. To make sure that this disk is recognized by the Windows VM as the same disk where the page file was stored at the primary site, the pre-created disk must have the same disk signature as the disk at the primary site. VMware disk cloning can be used in a procedure to accomplish this. Refer to the following paper for more information: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKCS&externalId=2009324

Every VM requires a swapfile, which is normally created in the same datastore as the other VM files. When you use SRM, this datastore is replicated. To prevent swapfiles from being replicated, create them on a non-replicated datastore.

7.1.3 Customize IP addresses at the DR site

In the solution testing, the same network IP addresses were used at both the primary and remote sites. It is not possible in all environments to have a different IP scheme for VMs upon failover. VMware provides a tool called the "dr-ip-customizer". This can assist with configuring SRM to automatically update VM IP address and network information when failover occurs. The dr-ip-customizer utility is a tool that generates a unique customization specification for each VM and applies that customization specification to the recovery plan for each VM. It makes it easy to change or update IP information for a VM. The tool is found in the Bin folder on the vCenter Server that is by default found in "C:\program files (x86)\vmware\vmware vcenter site recovery manager\bin". Directions for using the utility are provided in section "Customize IP Properties for a Group of VMs" in the Site Recovery Manager Administration Guide. The dr-ip-customizer utility is a tool that takes as input a file containing a comma-separated value (CSV) table of IP settings for multiple VMs, generates a unique customization specification for each VM, and then applies that customization specification to the recovery plan for each VM

7.1.4 Behavior of VMware datastore names during SRM failover

During failover and failback, by default the datastore volumes that are recovered by SRM during a DR failover will be prefixed with the snap-xxxx character string. To disable this, right-click the **Site Recovery** icon in the SRM GUI, select **Advanced Settings**, and select the checkbox for "SanProvider.fixRecoveredDatastoreNames" to remove the snap-xxx-prefix from the recovered data stores.



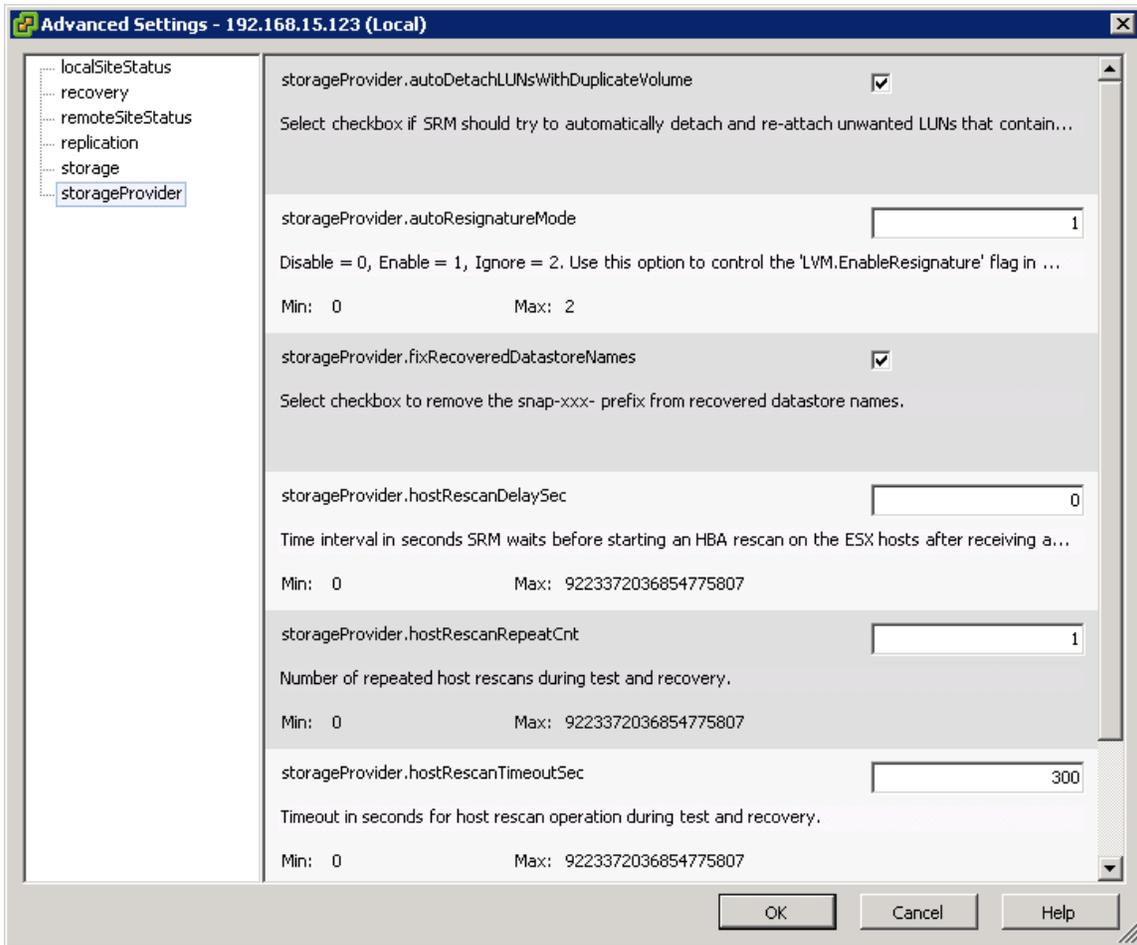


Figure 42 Disabling the snap-xxx prefix from recovered datastore names



8 Conclusions and findings

- vCenter SRM provides an automated DR of critical applications such as SQL Server running on VMs by simplifying and accelerating the recovery process.
- Group SQL VMs in fewer protection groups, because this enables fast recovery as supposed to creating one for each SQL VM. This reduces the overhead on SRM processes and can result in better RTO.
- RPO and RTO are two important metrics IT administrators keep in mind while designing or running the DR plan. In this solution, RPO is fulfilled by EqualLogic Auto-Replication where replication schedules can be managed by VSM or ASM/ME. SRM automates all tasks required for a successful failover to the recovery site which minimizes the RTO.
- RTO for the DR solution has many parameters such as network latency, number of VMs in the protection group, and number of protection groups in the recovery plan. IT administrators and DBAs should carefully consider these factors while designing a robust DR solution in order to comply with their business requirements.
- By using VSM, it provides a complete automatic failover, however the failover happens at a VM level and does not provide object level recovery of a database. This scenario provides better RTO compared to using ASM/ME.
- By using ASM/ME, it involves some manual steps during failover/failback, however this option provides granular database recovery.
- Multiple pools on the remote site require additional design considerations because the delegated space has to come from a single pool on the group at the remote site for each primary group. If the group has multiple pools on the primary site, each pool should be replicated to a different EqualLogic group on the remote site.
- Replicating TempDB can flood the network traffic in some environments, so it is recommended not to replicate the TempDB.
- Protect VMs with non-replicated Windows page files. By default, it replicates transient data, such as Windows paging files or VM swapfiles. Use more memory for VMs or move them to another datastore that is not replicated, but consider that this requires additional steps during failover.
- During failover and failback, by default the datastore volumes that are recovered by SRM during a DR failover will be prefixed with the snap-xxxx character string. It is recommended to disable this from vSphere client.
- VSM fails to create a snapshot and quiesce the VM when ASM/ME is configured on the VM. To create a VM snapshot and quiesce the VM, refer to the [VMware article](#) on how to disable application-consistent quiescing when using ASM/ME VSS requestor on the guest operating system.



A Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

- Dell EqualLogic Configuration Guide:
<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19852516/download.aspx>
- DR with EqualLogic and SRM 5 Technical Report
<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19968934/download.aspx>

Referenced or recommended publications:

- vCenter Site Recovery Manager
<http://www.vmware.com/products/site-recovery-manager/>

