



Dell EqualLogic PS Series Array: Understanding Synchronous Replication (SyncRep)

Abstract

This Technical Report describes how EqualLogic Synchronous Replication works to protect your data and offers configuration recommendations when planning to use EqualLogic SyncRep.

Copyright © 2012 Dell Inc. All Rights Reserved.

EqualLogic is a registered trademark of Dell Inc.

Dell is a trademark of Dell Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Information in this document is subject to change without notice.

Dell Inc. will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change.

Reproduction in any manner whatsoever without the written permission of Dell is strictly prohibited.

Authored By:

Rick Sussman

Chuck Farah

Special thanks to Carole Gelotti

March 2014

Preface

PS Series arrays optimize resources by automating performance and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software, host software, and free firmware updates.

Audience

The information in this guide is intended for administrators who wish to configure Synchronous Replication for real-time data protection.

Related Documentation

For detailed information about PS Series arrays, groups, volumes, array software, and host software, log in to the [Documentation page](#) at the customer support site.

Dell Online Services

You can learn about Dell products and services using this procedure:

1. Visit <http://www.dell.com> or the URL specified in any Dell product information.
2. Use the locale menu or click on the link that specifies your country or region.

Dell EqualLogic Storage Solutions

To learn more about Dell EqualLogic products and new releases being planned, visit the Dell EqualLogic TechCenter site: <http://delltechcenter.com/page/EqualLogic>. Here you can also find articles, demos, online discussions, technical documentation, and more details about the benefits of our product family.

For an updated Dell EqualLogic compatibility list please visit the following URL: <https://support.equallogic.com/compatibility>

Table of Contents

Revision Information.....	v
Introduction.....	1
How SyncRep Works	2
Initial Synchronization.....	3
SyncRep States.....	5
Synchronous Replication Configuration	8
Requirements and considerations for using SyncRep.....	8
Network considerations.....	9
Configuring SyncRep on a Volume	10
Volume with SyncRep configured.....	13
Configure Volume Collections with SyncRep.....	14
Changing Pool Relationship.....	17
Switch to SyncAlternate.....	17
Failover to SyncAlternate	19
Snapshots	25
APPENDIX A: Technical Best Practices	30
Business Purpose for SyncRep	31
Distance Considerations.....	34
Volume I/O and application profile Considerations	35
SAN Design for maximum availability and recovery	35
Snapshot space Considerations.....	36
APPENDIX B: SyncRep and Dell EqualLogic Host Integration Tools.....	37
Auto-Snapshot Manager/Microsoft Edition (ASM/ME).....	37
ASM/ME Impact during Switch.....	37
VMware integration with SyncRep	41
Auto-Snapshot Manager/Linux Edition (ASM/LE)	41
APPENDIX C: SAN HeadQuarters Management with SyncRep.....	42

Revision Information

The following table describes the release history of this Technical Report.

Report	Date	Document Revision
1.0	August 2012	Initial Release
1.1	April 2014	Updated with info on supported SyncRep configurations

The following table shows the software and firmware used for the preparation of this Technical Report.

Vendor	Model	Software Revision
Dell™	PS Series Array Firmware	Version 6.0.0 and above

The following table lists the documents referred to in this Technical Report. All PS Series Technical Reports are available on the Customer Support site at: support.dell.com

Vendor	Document Title
Dell™	Dell EqualLogic Group Manager Administrator's Manual, PS Series Firmware V6.0

Introduction

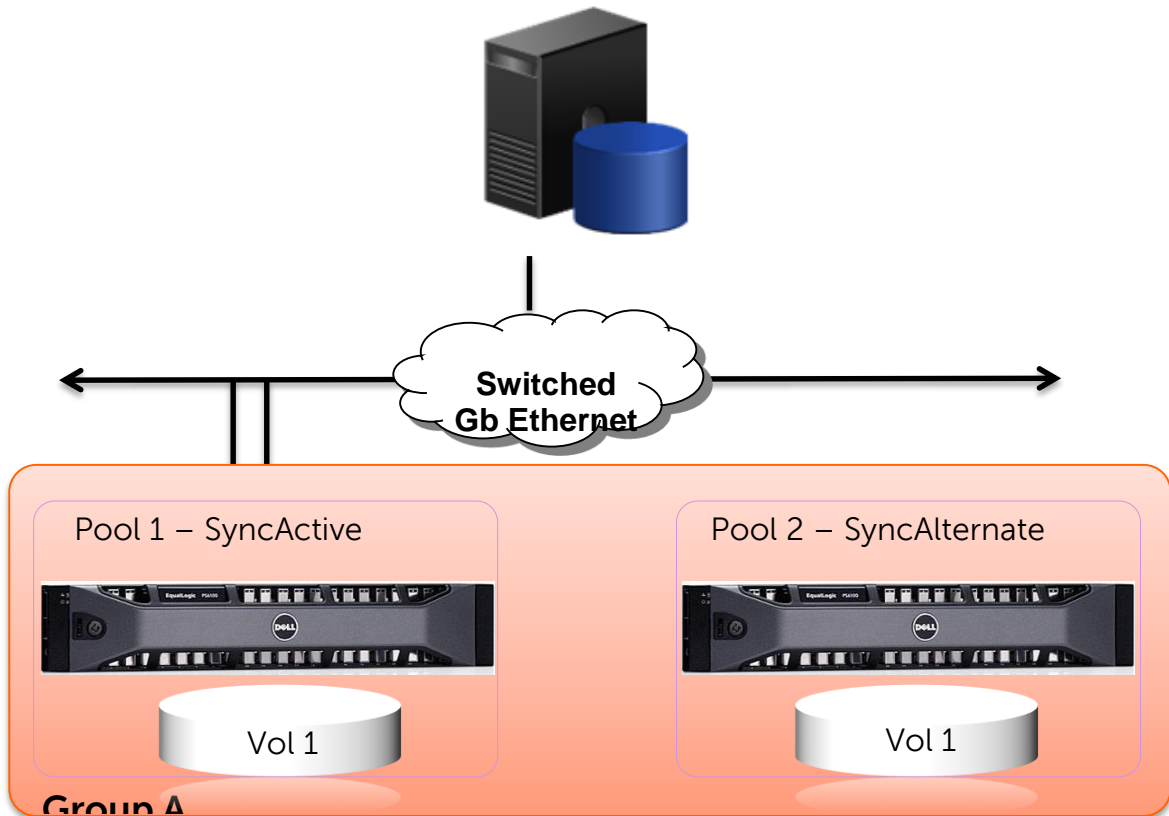
Synchronous Replication (SyncRep) is a new feature included in EqualLogic PS Series firmware v6.0. EqualLogic SyncRep is the simultaneous writing of volume data across two different storage pools in the same PS Series Group, resulting in two hardware-independent copies of volume data. SyncRep will only acknowledge write completion when both the source and target volume has completed. If one pool is not available due to a power failure or other resource outage, you can still obtain volume data from the other pool.

- SyncRep is the process of copying data over a storage area network so there are two copies of the data.
- SyncRep writes data to the active and alternate sites synchronously before acknowledgement is sent back to the host (initiator) so that the data remains current in both locations.
- SyncRep is highly desirable for applications with short recovery time objectives (RTOs) and a real-time recovery point objective (RPOs).
- Provides Quick Recovery if on copy of the data becomes unavailable.
- SyncRep is a firmware feature and is operating system agnostic.
- SyncRep is a real-time data protection solution and is not recommended for disaster recovery.

Note: You cannot perform traditional replication on a synchronous replication-enabled volume. Refer to the **Dell EqualLogic Group Manager Administrator's Manual, PS Series Firmware Version 6.0** for more information.

How SyncRep Works

SyncRep is enabled on a per-volume basis. In volumes for which synchronous replication is not enabled, volume data and snapshots are located only in the pool to which the volume is assigned. In synchronous replication-enabled volumes, volume data exists simultaneous in two pools.




- **SyncActive** – The pool to which the iSCSI initiators (hosts) are connected to when reading and writing volume data.
- **SyncAlternate** - When volume data is written to the SyncActive pool, the group simultaneously writes the same data to this pool. You can switch the roles of the SyncActive and SyncAlternate pools.

SyncActive and SyncAlternate pools can be switched. The former SyncActive pool then becomes the SyncAlternate pool, and vice-versa. No iSCSI target configuration changes are required. During the switch, host connections are logged out. Depending on the configuration, the iSCSI initiator may reconnect automatically and transparently when the switch has completed.

Initial Synchronization

The first time a volume is configured with SyncRep, a mirror volume is created in the SyncAlternate pool and the SyncActive volume will replicate all data to the SyncAlternate volume. This is a process that occurs in background and non-replicated changes will be tracked and sent to the SyncAlternate volumes. All new I/O will continue to both pools.



1. Create SyncAlternate Volume.
2. The SyncActive volume and SyncAlternate will receive all updates.
3. Changes that need to be sent to the SyncAlternate will be tracked.
4. The host receives acknowledgement of changes from the SyncActive volume.


For SyncRep initial synchronization consider these factors:

- The initial volume size (SyncActive)
 - Large volumes when full of data will take time to completely synchronize.
- The network bandwidth between pools.
 - Bandwidth constraints may affect the time to complete initial synchronization of large volumes. For example, a 500 GB volume may take at least an hour over a dedicated 1 GigE network to initially synchronize.
- Other activity occurring on these pools during initial synchronization.
 - Writes at the SyncAlternate Pool will favor user application traffic and may delay the overall time to synchronize.

SyncRep States

A volume that is configured for SyncRep can be in one of the following three states:

- **In sync** - Both the SyncActive and the SyncAlternate pools contain the same volume data. Switching roles is allowed.




In sync process steps:

1. iSCSI initiator (host) sends a write to the group.
2. All **writes** are done to BOTH volumes; all **reads** are done from the SyncActive volume ONLY.
3. SyncActive and SyncAlternate volumes acknowledge to the group that the **writes** are complete.
4. The **writes** are then confirmed to the iSCSI initiator (host).

Note: A pool switch for a synchronous replication volume swaps the roles of the SyncActive and SyncAlternate pools. As part of the operation, the SyncActive pool becomes the SyncAlternate pool and the SyncAlternate pool becomes the SyncActive pool. A **Switchover** is a planned maintenance on SyncActive where the volumes are **in sync**.

- **Paused** – SyncRep can be paused through the Group GUI. While SyncRep is paused, the volume is still online, and initiators can connect to and write to the SyncActive volume. SyncRep can be paused and can later be resumed. For example, this could happen in a maintenance window during which the SyncAlternate pool is taken offline. If data is written to the volume while synchronous replication is paused, it is written only to the SyncActive pool, and the two pools are **out of sync**. The group tracks all volume writes while synchronous replication is paused and when synchronous replication is resumed, the tracked changes are written to the SyncAlternate pool.
- **Out of sync** - SyncActive pool and SyncAlternate pool may not contain the same volume data; the SyncActive pool contains the most recent volume data. A SyncActive volume can go out of sync if synchronous replication is **paused**, the SyncAlternate pool is unavailable or has no free space, or there is insufficient bandwidth between the SyncActive and SyncAlternate pools to keep up with the write demand on the volume.



Out of sync process:

1. SyncAlternate is unavailable.
2. All **reads** and **writes** occur in SyncActive pool.
3. Changes are tracked and **failover** is allowed.

Note: If the SyncActive volume is out of sync, the option to switch the pools is replaced with the option to **failover to the SyncAlternate** pool. A **Failover to the SyncAlternate** option is available when a fault has already occurred. You can only failover to the SyncAlternate pool when the volume is already **out of sync** and the volume is taken offline.

Note: When Synchronous Replication is **paused** by an Administrator, the volumes will go **out of sync**.

Synchronous Replication Configuration

Requirements and considerations for using SyncRep

- PS Series Array firmware v6.0.0 or higher.
- Two pools, each containing at least one array member.
- Requires fast, reliable connectivity between pools.
 - The tolerance for latency that can be introduced to the network is determined by the application that is using the volume.
 - Latency of the link is the primary consideration and will be the gating factor for application write performance.
 - Bandwidth – at least 1 Gb (gigabit) or more is preferred.
 - Distance < 1Km.
 - Since Synchronous Replication implementation is two pools in the same PS Series group and all members must be in the same VLAN/subnet, and a “flat network” is required. Please see the [Network](#) section for more information.
- Volumes cannot be configured for both traditional replication and Synchronous Replication.
- Volume rethinning (unmap) is not recommended on volumes where Synchronous Replication is used. Please see the Dell EqualLogic PS Series Storage Release Notes Version 6.0 for more information.
- Supports up to 32 SyncRep enabled volumes in a PS Series group (PS4xx0 supports up to 4 Sync Rep volumes)*.
- Minimum of 5% snapshot reserve required.
- No technical limit to the amount of time or data.
- SyncRep is not available for volumes bound to members.

*** Note:** A PS Series group must contain only PS6xx0 arrays in order to support 32 volumes for SyncRep. If a PS4xx0 array is hosted anywhere in the group the max # of volumes allowed for SyncRep will be 4.

Network considerations

The SAN network configured in an EqualLogic SyncRep solution follows the same network best practices as for setting up a PS Series group. Below is a short list of possible considerations which are documented and located at the [Dell TechCenter](#) or Dell EqualLogic [support](#) site:

- SyncRep requires a “flat” network infrastructure. Both pools should be in the same subnet (not routed) without compression, de-duplication or encryption devices; all which adds latency.
- Recommend using network switch models which are capable of accommodating bursts of traffic without dropping packets. Buffers and Flow control need to be configured according to published reference architectures. For more information please see Dell TechCenter: [Reference Architectures](#)
- LAN speeds, 1Gige, 10Gige or mixed: See Dell TechCenter: [Deploying Mixed 1 Gb -10Gb Ethernet SANs](#)
- SyncActive and SyncAlternate pools should have the same models of switches configured if possible and not shared with Client LAN traffic.
- If SyncActive pool is on its own unique set of switches and SyncAlternate is also on a unique set of switches, ensure that at worse case there is only a single hop between the SAN switches connecting the two pools.
- Inter-switch connection sizing should follow the Dell EqualLogic Configuration Guide recommendations. See the chapter on “Sizing Inter-Switch Connections” in the [Dell EqualLogic Configuration Guide](#).

Note: Network infrastructure should follow the best practices for Dell PS Series EqualLogic arrays.

Configuring SyncRep on a Volume


Configure SyncRep on an existing volume.

The screenshot displays the EqualLogic PS Series Group Manager web interface. The left sidebar shows a tree view of volumes, with 'SyncRepVol1' selected. A red arrow labeled '1' points to the 'Volumes' button in the sidebar. Another red arrow labeled '2' points to 'SyncRepVol1' in the volume list. A third red arrow labeled '3' points to the 'Configure SyncRep' link in the 'SyncRep' section of the 'Activities' panel. The main content area shows the configuration details for 'Volume SyncRepVol1', including general information, storage pool, volume space, snapshot space, and volume distribution.

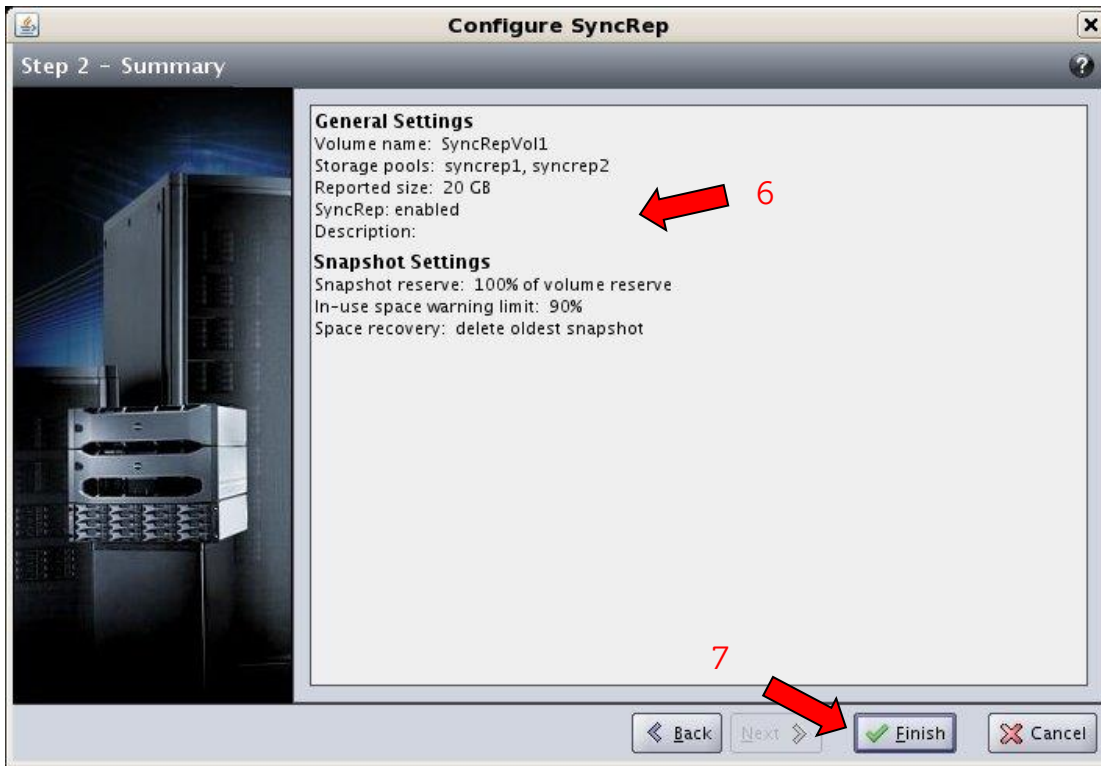
1. Select the **Volumes** button.

2. Select the volume for which you are configuring SyncRep (**SyncRepVol1**).

3. In the **Activities** panel, select **Configure SyncRep** to open SyncRep wizard.



4. Specify the SyncAlternate pool for the volume, (in this example there is only one pool to choose).
5. Click **Next** to go to the **Summary** step of the wizard.



6. Verify that the information contained in the summary is correct.
7. Click **Finish** to complete the wizard and enable synchronous replication.

Volume with SyncRep configured

The screenshot displays the EqualLogic PS Series Group Manager interface. The left sidebar shows a tree of volumes, with 'SyncRepVol1' selected. The main area shows the configuration for 'Volume SyncRepVol1'. The 'Replication' tab is active, showing the following information:

- General Volume Information:**
 - Status: online
 - General Settings: Volume name: SyncRepVol1, Reported size: 20 GB
 - SyncRep: Status: in sync, Unreplicated data: 0 MB
- Volume and Snapshot Space:**
 - Storage pool: syncrep1 (SyncActive), syncrep2 (SyncAlternate)
 - Volume Space: Reported size: 20.01 GB, Thin provisioning: disabled, Volume reserve: 20.01 GB (100%), Total volume reserve: 40.02 GB
 - Snapshot Space: Snapshot reserve: 20.01 GB (100%), Total snapshot reserve: 40.02 GB, Space recovery: delete oldest snapshot, Number of snapshots: 0
 - Volume Distribution: RAID preference: automatic, Load balancing: enabled

1. SyncRep status (Status: **in sync**).
2. SyncRep icon.
3. SyncRep operations (Switch to SyncAlternate, Pause SyncRep, Disable SyncRep, and Disconnect SyncActive).
4. SyncRep pool information.

Configure Volume Collections with SyncRep

Volume Collections allows multiple volumes to act as one object.


A SyncRep collection:

- All volumes in the collection are **paused** at the same time.
- All volumes in the collection are **switched over** or **failed over** at the same time.
- Provides the ability to create crash consistent snapshots for multiple volumes at a single point in time.

Note also:

- A volume must be **in sync** before it can be added to a SyncRep volume collection.
- All volumes in a SyncRep volume collection must have SyncRep enabled before being added to the volume collection (which will be enabled for SyncRep).
- SyncRep enabled volumes can only belong to one SyncRep enabled volume collection at a time.
- SyncRep relationship switch cannot be performed on individual volume in the collection.
- SyncRep collections are only considered to be **in sync** when all volumes in the collection are **in sync**.

SyncRep Collection steps:



1. **Create/Modify** a collection and include all volumes you wish to manage as one object with SyncRep.

Volume collection SyncRepCollection

Activities

- Collection SyncRepCo...
 - Volume Collection
 - Modify volume collection
 - Delete volume collection
 - Snapshots
 - Create snapshot collection
 - Schedules
 - Create schedule
 - Replication
 - Configure replication
 - SyncRep
 - Enable SyncRep for collection
 - Configure SyncRep**
- Volume SyncRepCoIV1
 - Volume
 - Modify settings
 - Clone
 - Set offline
 - Set access type
 - Delete volume
 - Convert to template
 - Move volume
 - Folder
 - Move to folder
 - Access
 - Create access policy
 - Snapshots
 - Modify snapshot settings
 - Create snapshot
 - Restore volume
 - Delete snapshots
 - Schedules
 - Create schedule
 - Replication
 - Configure replication
 - SyncRep
 - Configure SyncRep

Collection Status

Snapshots
Snapshot collections: 0
Last snapshot collection: none created

Replication: disabled
SyncRep: disabled

Collection Volumes

Total volumes in the collection: 2

Volume	Storage pool	Reported size	Volume reserve	Volume status	Replication partner	SyncRep status	iSCSI connections
SyncRepCoIV1	default	5.01 GB	525 MB	online			0
SyncRepCoIV2	default	5.01 GB	525 MB	online			0

2. For each Volume in the Collection, **Configure SyncRep**.

Volume collection SyncRepCollection

Activities

- Collection SyncRepCo...
 - Volume Collection
 - Modify volume collection
 - Delete volume collection
 - Snapshots
 - Create snapshot collection
 - Schedules
 - Create schedule
 - Replication
 - Configure replication
 - SyncRep
 - Enable SyncRep for collection
 - Configure SyncRep
- Volume SyncRepCoIV2
 - Volume
 - Modify settings
 - Clone
 - Set offline
 - Set access type

Collection Status

Snapshots
Snapshot collections: 0
Last snapshot collection: none created

Replication: disabled
SyncRep: disabled

Collection Volumes

Total volumes in the collection: 2

Volume	Storage pool	Reported size	Volume reserve	Volume status	Replication partner	SyncRep status	iSCSI connections
SyncRepCoIV1	default	5.01 GB	525 MB	online		in sync	0
SyncRepCoIV2	default	5.01 GB	525 MB	online		in sync	0

3. **Enable SyncRep for Collection** will enable SyncRep on the collection of volumes.

The screenshot shows the 'Volume collection SyncRepCollection' interface. On the left is an 'Activities' sidebar with options like 'Modify volume collection', 'Delete volume collection', 'Create snapshot collection', 'Create schedule', 'Switch collection to SyncAlternate', 'Disable SyncRep for collection', and 'Disconnect SyncActive volumes'. The main area has tabs for 'Status', 'Snapshots', 'Replica', and 'Schedules'. The 'Status' tab is active, showing 'Collection Status' and 'Collection Volumes' sections. Two red arrows point to the 'in sync' status in both sections.

Collection Status

Snapshots
 Snapshot collections: 0
 Last snapshot collection: none created

SyncRep
 Status: in sync
 Storage pools: default (Active), SSD-SAS

Collection Volumes

Total volumes in the collection: 2

Volume	Storage pool	Reported size	Volume reserve	Volume status	Replication partner	SyncRep status	iSCSI connections
SyncRepCoIV1	default	5.01 GB	525 MB	online		in sync	0
SyncRepCoIV2	default	5.01 GB	525 MB	online		in sync	0

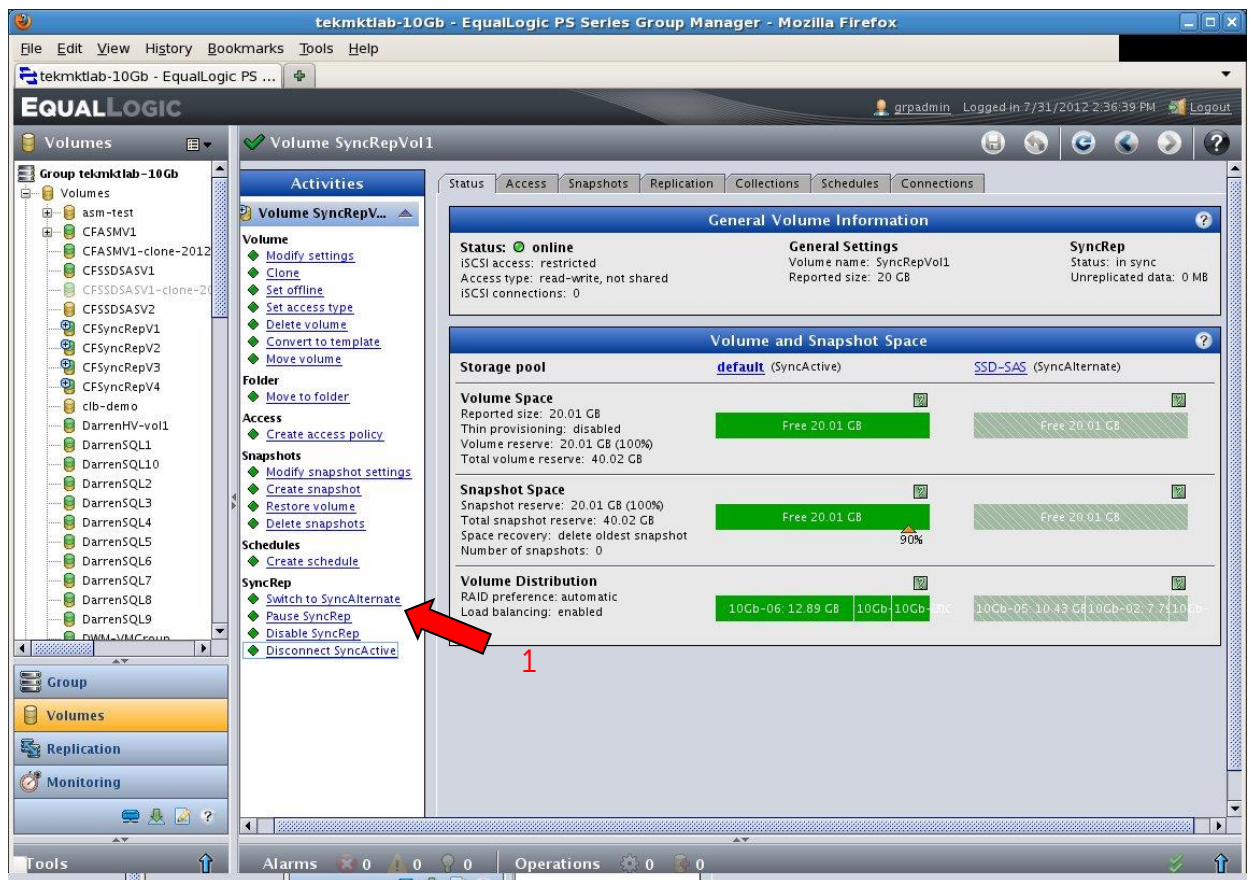
Note: Collection Status changes to **in sync** only after all volumes are **in sync**.

Changing Pool Relationship

Pool relationship change refers to the ability to change which pool is SyncActive and which is SyncAlternate. The reasons for a pool relationship change are SyncActive member failure, SyncActive pool failure or maintenance on SyncActive members. There are two ways pool relationships are changed: **Switch to SyncAlternate** and **Failover to SyncAlternate**.

Switch to SyncAlternate

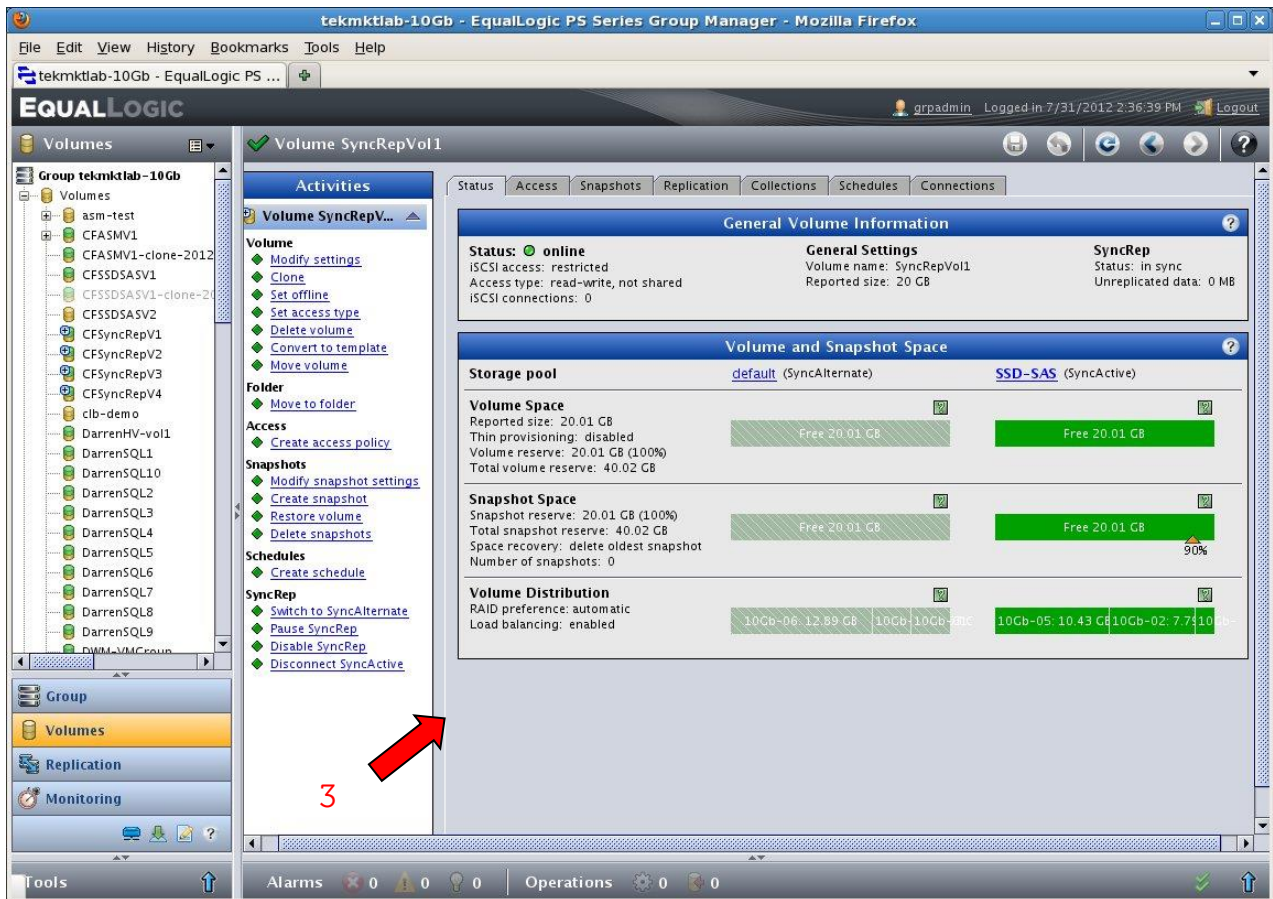
This pool relationship change is done manually and is only available when the volume is **in sync**. During the switch, host connections are logged out. The iSCSI initiators will automatically reconnect when the switch has completed. No iSCSI target configurations changes are required. No data will be lost.



1. Click **Switch to SyncAlternate**



2. Click **Yes**



3. SyncActive has switched pools

Failover to SyncAlternate

Failover is a pool relationship change that is required because of a fault occurring in the SyncActive pool. Failover to SyncAlternate is used to bring the SyncAlternate copy of data online and can only be run when the SyncRep status is **out of sync**. There are more steps to a manual failover than a manual switchover.

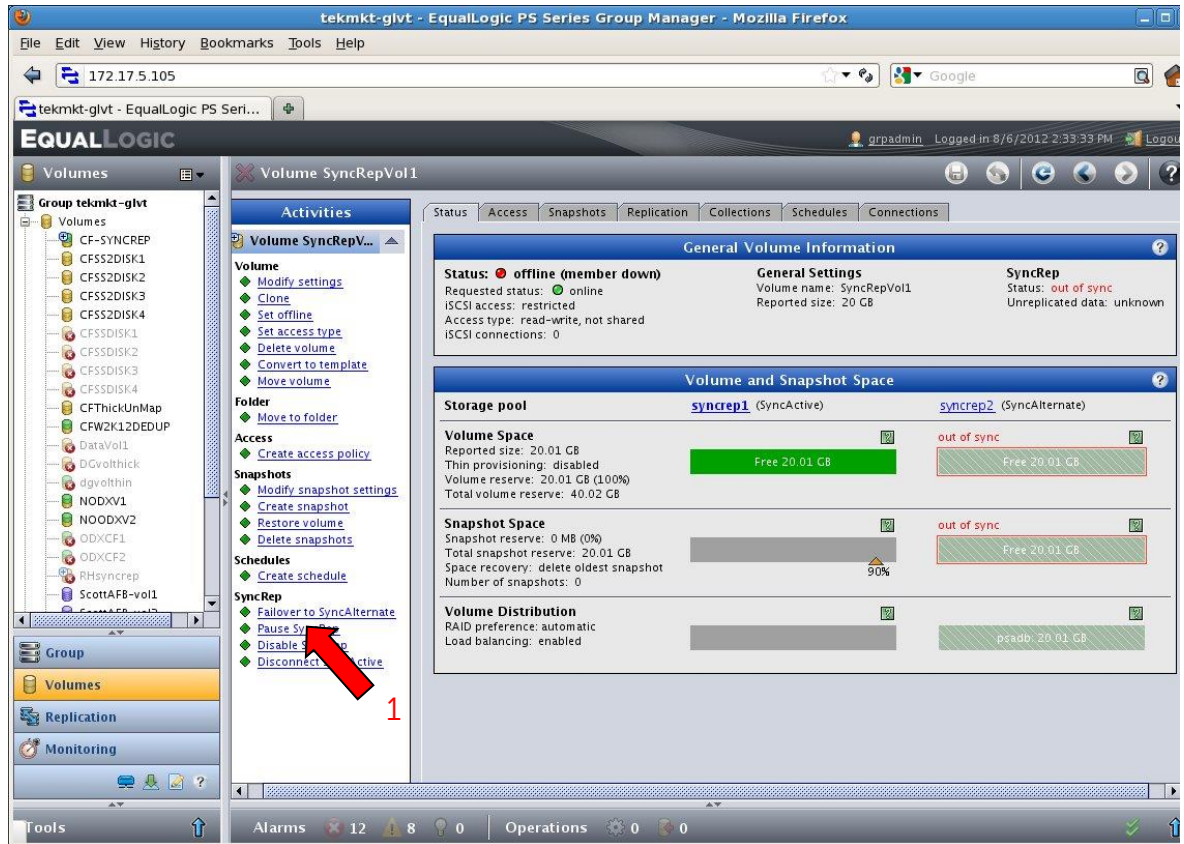
There are three possible Failover situations which may occur:

- **Out of sync** when both pools are online, (one member down in a multi-member pool for example).
- **Out of sync** due to an infrastructure problem or similar where both pools are online but cannot talk to each other, (a switch down for example).
- **Out of sync** because the SyncActive pool is offline, (see example below).

The new SyncActive volume (formerly SyncAlternate) accepts writes and tracks changes made while the original SyncActive copy is offline. Depending on the host configuration and application requirements, it may be necessary to log into the volume again from the host initiator and restart the applications.

Note: Dell recommends contacting support to be sure that all variables are considered before a Failover is initiated.

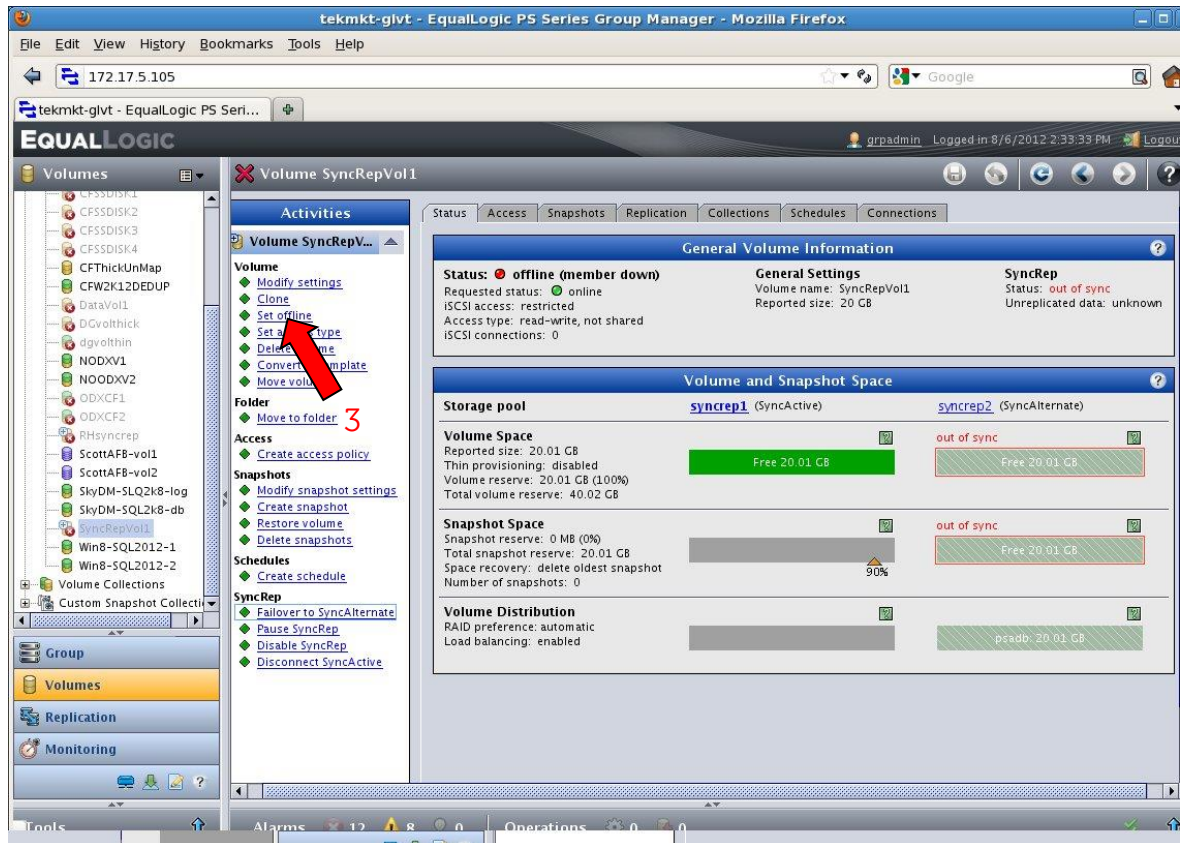
The following outlines the Failover situation when the SyncActive pool is offline causing the volumes to become **out of sync**.



1. Click **Failover to SyncAlternate**, The Failover to SyncAlternate option will appear in the Activities list when the volume goes **out of sync**.



2. Click **OK**



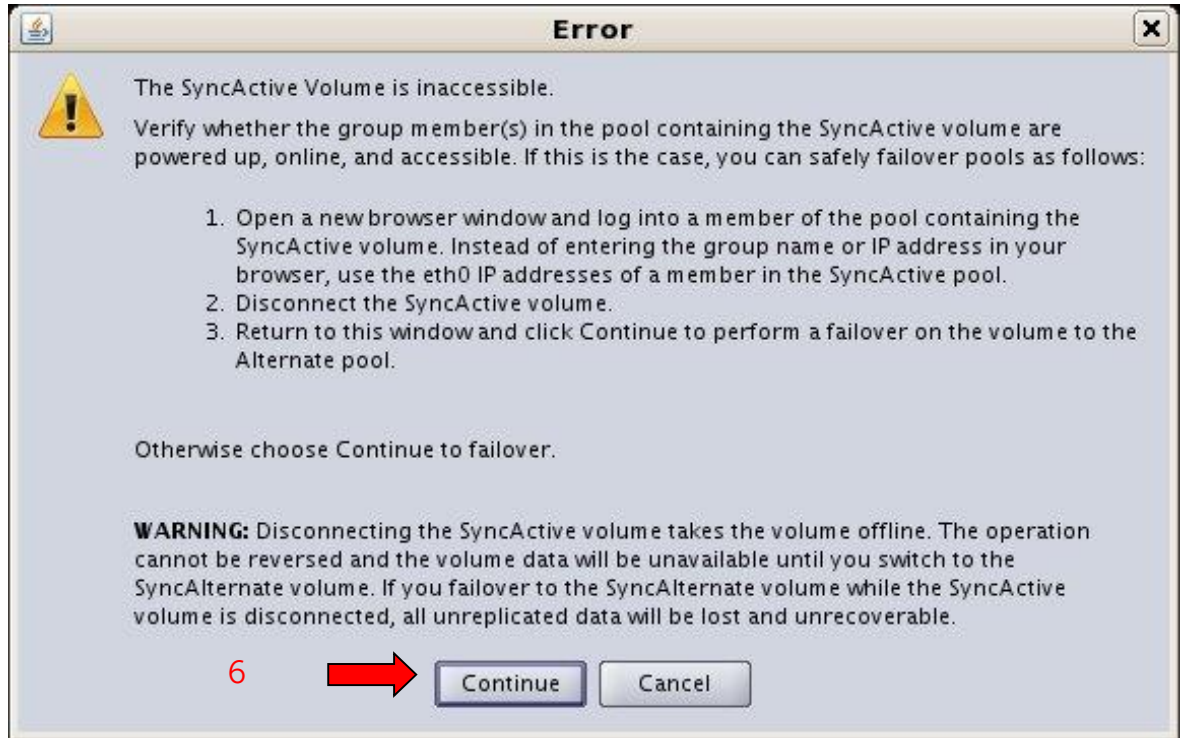
3. Set the volume offline.



4. Click Yes

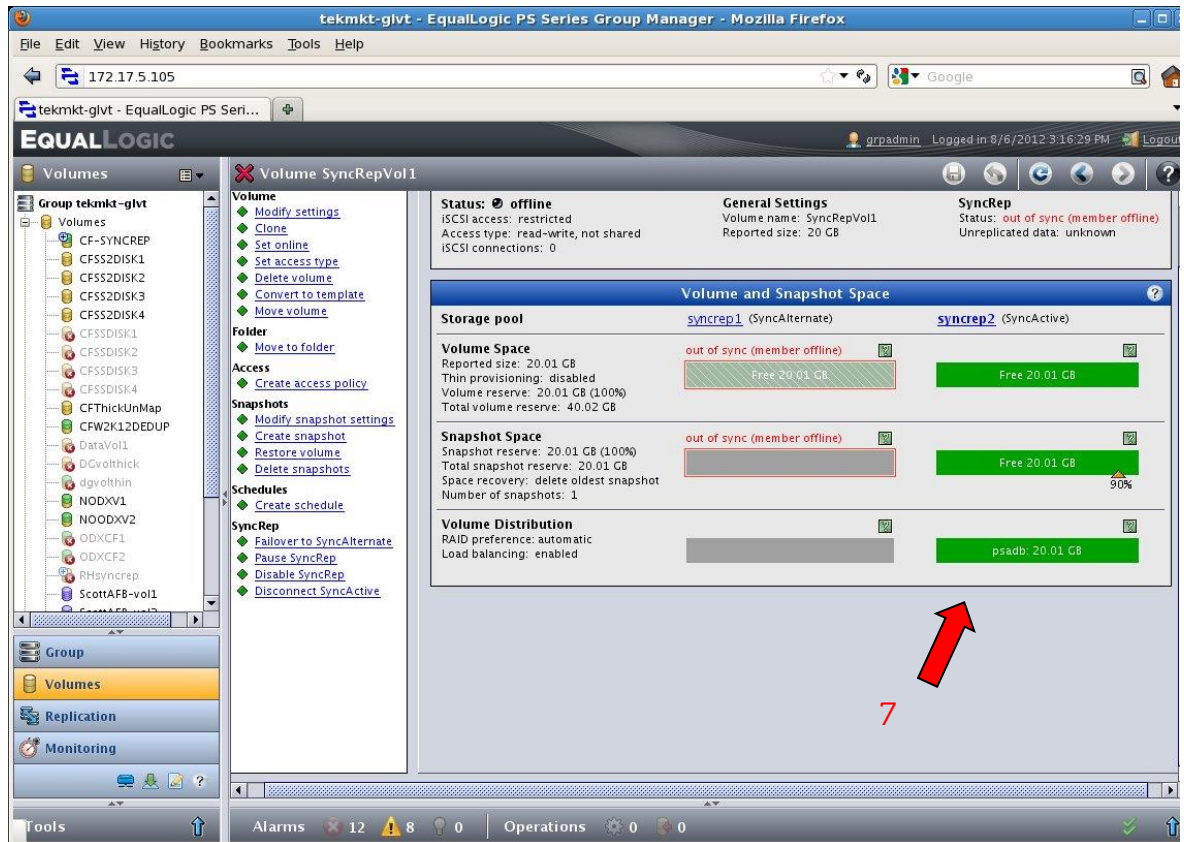


5. Click **Yes**



Note: If enabled, use the *management port* to log into a member of the pool containing the SyncActive volume.

6. Click **Continue** if the **SyncActive** pool has been verified offline or unreachable. If the SyncActive pool is still online, log directly into a member in the SyncActive pool and choose **Disconnect SyncActive** to properly failover to the SyncAlternate pool.



7. The **Failover to SyncAlternate** has completed. The original **SyncAlternate** has become the **SyncActive**.

Note: After access to the original volume has been restored the system will resume synchronous replication, all changes will be tracked. When the volume returns to **in sync** you can **switch** back to the original pool configuration or choose to continue with the new configuration.

Snapshots

Snapshots are valuable methods of data protection and are fully supported on SyncRep enabled volumes. SyncRep uses two types of snapshots:

- User snapshots – created by the user for data protection.
- System snapshots – created by the PS Series groups for state data of the volume.

The following snapshot criteria are true of SyncRep enabled volumes:

- Scheduling or manual creation of a snapshot will locate that snapshot in the SyncActive pool only.
- As with snapshots of volumes for which synchronous replication is not enabled, the snapshot data may be accessed by setting the snapshot online or cloning it. It does not matter whether the snapshot resides in the SyncActive pool or the SyncAlternate pool.
- Switchover/Failover moves the new snapshot or schedule to the new SyncActive pool, but existing snapshots remain in the pool they were created in.
- System snapshots are created and removed automatically to save the point in time state information of a volume when **out of sync**.

Note: System snapshots (**Synchronous replication system snapshots**) are deleted automatically when the volume is **in sync**.

Snapshot pool residence during **switch** operations:

General Volume Information

Status: ● online
iSCSI access: restricted
Access type: read-write, not shared
iSCSI connections: 4

General Settings
Volume name: CFSyncRepV1
Reported size: 20 GB

SyncRep
Status: in sync
Unreplicated data: 0 MB

Volume and Snapshot Space

Storage pool: [default](#) (SyncAlternate) | [SSD-SAS](#) (SyncActive)

Volume Space
Reported size: 20.01 GB
Thin provisioning: enabled
TP warning mode: None
Volume reserve: 2.01 GB (10%)
Total volume reserve: 4.01 GB

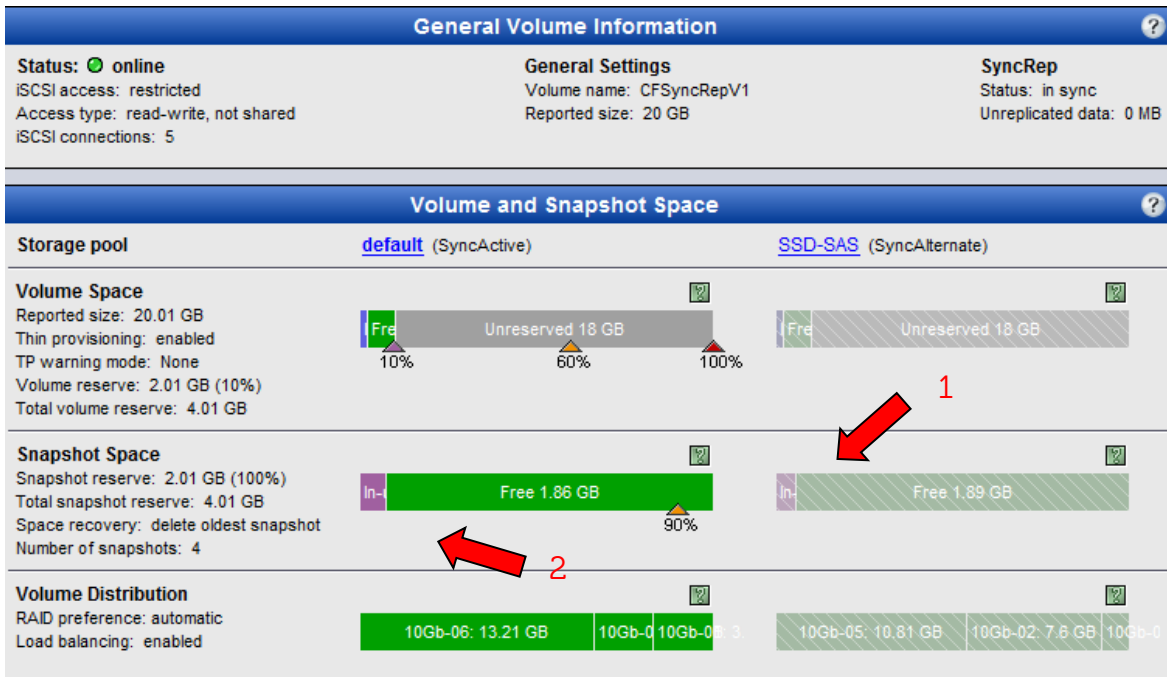
Snapshot Space
Snapshot reserve: 2.01 GB (100%)
Total snapshot reserve: 4.01 GB
Space recovery: delete oldest snapshot
Number of snapshots: 2

Volume Distribution
RAID preference: automatic
Load balancing: enabled

Original snapshots will remain in the **default** pool.

1. Originally, the SyncActive volume was in the **default** pool and snapshots were taken prior to switch.
2. Currently, the **SSD-SAS** pool is the SyncActive for this volume and any new snapshots will be contained here.

Switch back - New snapshots will be in the pool where the volume is SyncActive.



After switch back:

1. All snapshots created in the **SSD-SAS** pool will remain in the volumes SyncAlternate pool.
2. New snapshots for the SyncActive volume will now reside in the **default** pool along with the existing snapshots prior to the switch to the **SSD-SAS** pool.

Note: The snapshots will remain with the pool where the volume was SyncActive at the time of snapshot or clone creation.

Automatic System Snapshot creation:

The group creates system snapshots of SyncRep volumes for any of the following:

- Anytime an administrator performs a switch, failover or disconnect.
- Whenever the volume's state changes to **paused** or **out of sync**.
- Just before the group begins writing tracked changes to a SyncAlternate pool for an **out of sync** volume.

These snapshots reflect the volume's data at the point in time when the operation was performed or the volume's status has changed.

System snapshots– example if SyncActive volume is **paused**:

The screenshot displays the configuration for a storage volume named 'SyncRepColV2' with a reported size of 5 GB. It is divided into two sections: 'General Volume Information' and 'Volume and Snapshot Space'.

General Volume Information:

- Status:** online
- iSCSI access:** restricted
- Access type:** read-write, not shared
- iSCSI connections:** 0
- General Settings:** Volume name: SyncRepColV2, Reported size: 5 GB
- SyncRep:** Status: **paused**, Unreplicated data: 0 MB

Volume and Snapshot Space:

Storage pool: default (SyncActive) and SSD-SAS (SyncAlternate)

Volume Space:

- default (SyncActive):** Reported size: 5.01 GB, Thin provisioning: enabled, TP warning mode: None, Volume reserve: 525 MB (10.2%), Total volume reserve: 1.03 GB. The bar chart shows 10% free space, 60% unreserved (4.5 GB), and 100% total.
- SSD-SAS (SyncAlternate):** Status: **paused**. The bar chart shows 100% free space and 100% unreserved (4.5 GB).

Snapshot Space:

- default (SyncActive):** Snapshot reserve: 525 MB (100%), Total snapshot reserve: 1.03 GB, Space recovery: delete oldest snapshot, Number of snapshots: 1. The bar chart shows 90% free space (525 MB).
- SSD-SAS (SyncAlternate):** Status: **paused**. The bar chart shows 100% free space (525 MB).

Volume Distribution:

- default (SyncActive):** RAID preference: automatic, Load balancing: enabled. The bar chart shows 10Gb-06: 5.01 GB.
- SSD-SAS (SyncAlternate):** The bar chart shows 10Gb-05: 5.01 GB.

System snapshot process:

1. volume **paused** (or **out of sync** state)

View of **Synchronous replication system snapshot**.

The screenshot shows a web interface with two tabs: 'Status' and 'Access'. The main content is divided into two sections:

- General Snapshot Information**:
 - Status:** offline
 - iSCSI access:** restricted
 - Access type:** read-write, not shared
 - iSCSI connections:** 0
 - General Information:** Created: 8/2/2012 3:12:23 PM, Storage pool: default, Schedule: Size: 5 GB
 - Snapshot Collection:** Not part of collection
 - Description:** Synchronous replication system snapshot (indicated by a red arrow and the number 2)
- Snapshot iSCSI Settings**:
 - iSCSI target:** iqn.2001-05.com.equallogic:0-8a0906-117738306-7af1e5f9975501ad-syncrepcolv2-2012-08-02-15:12:23.2865.1
 - Public alias:** SyncRepColV2-2012-08-02-15:12:23.2865.1

2. **Synchronous replication system snapshot** created.

The screenshot shows a web interface with four tabs: 'Status', 'Snapshots', 'Replica', and 'Schedules'. The main content is divided into two sections:

- Snapshot Summary**:
 - Snapshots:** Snapshot collections: 0, Last snapshot collection: none created
 - Snapshot Schedules:** Snapshot schedules: 0, Running schedules: 0, Next snapshot: none scheduled
- Snapshots**:
 - No snapshot collections created**
 - | Created | Schedule | Status | Permission | Connections |
|---------|----------|--------|------------|-------------|
| None | | | | |

 (indicated by a red arrow and the number 3)

3. After volume is **in sync** the **Synchronous replication system snapshot** is deleted.

APPENDIX A: Technical Best Practices

Several technical considerations should be evaluated to choose the appropriate volumes for SyncRep.

SyncRep considerations include:

- Business Purpose
- Application performance requirements
 - Distance between pools
 - Bandwidth available
 - Change Rate (unique writes to pages)
 - Latency
- SAN Design
 - Network infrastructure
 - Array configuration
- Space available for additional copy of the data and snapshots

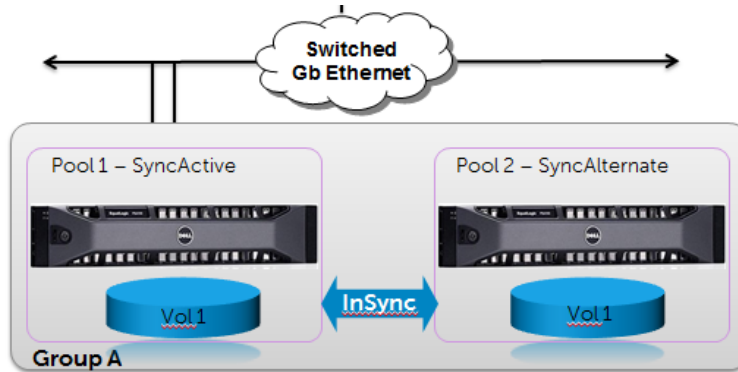
The following sections will try to outline some of these considerations in more detail.

Business Purpose for SyncRep


The design goal of EqualLogic SyncRep is to enable customers to elevate the availability of key volumes by protecting a pool within a PS Series Group. Below are guidelines to determining the appropriate business use cases for SyncRep.

Example Use Cases:

- Datacenter protection: single member pools same rack or floor.



Normal operations: Pools are in sync



Pools are out of sync

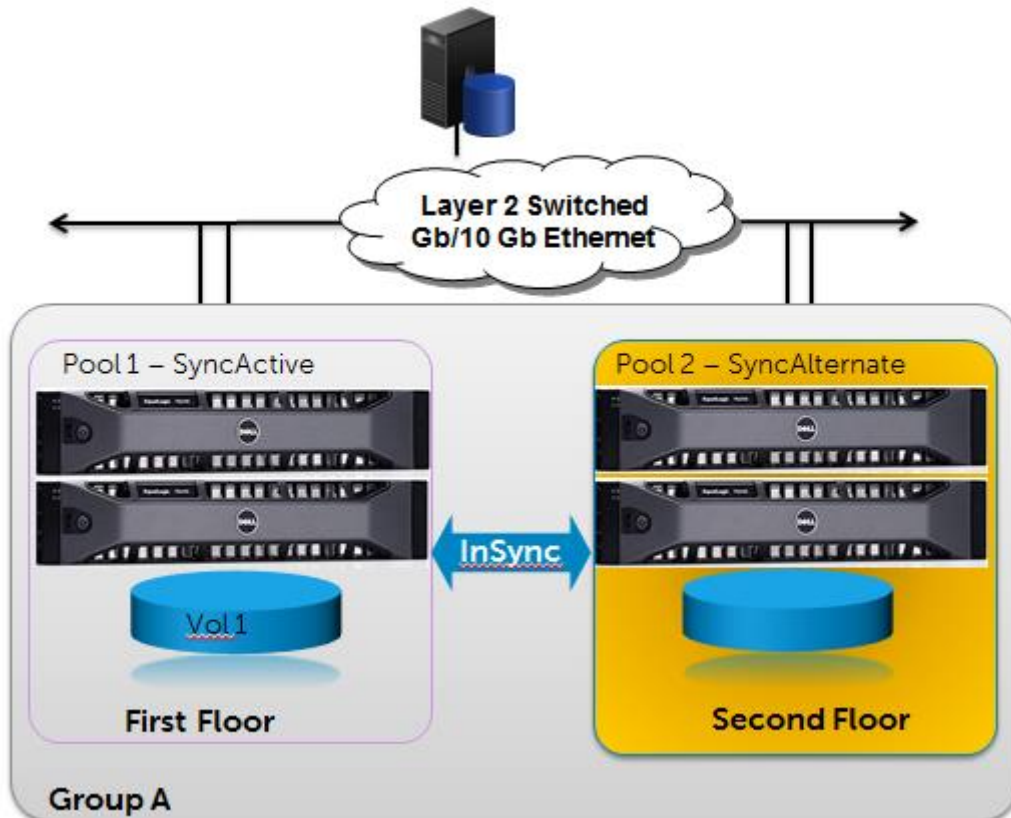
During failover to Pool1 changes are tracked to allow for re-sync when Pool2 is available.

Note: The SyncActive Volume is able to failover to the SyncAlternate pool in the event of I/O access problems.

- Extended Datacenter protection: pools located in different racks or floors. For instance a group that spans two buildings or floors.

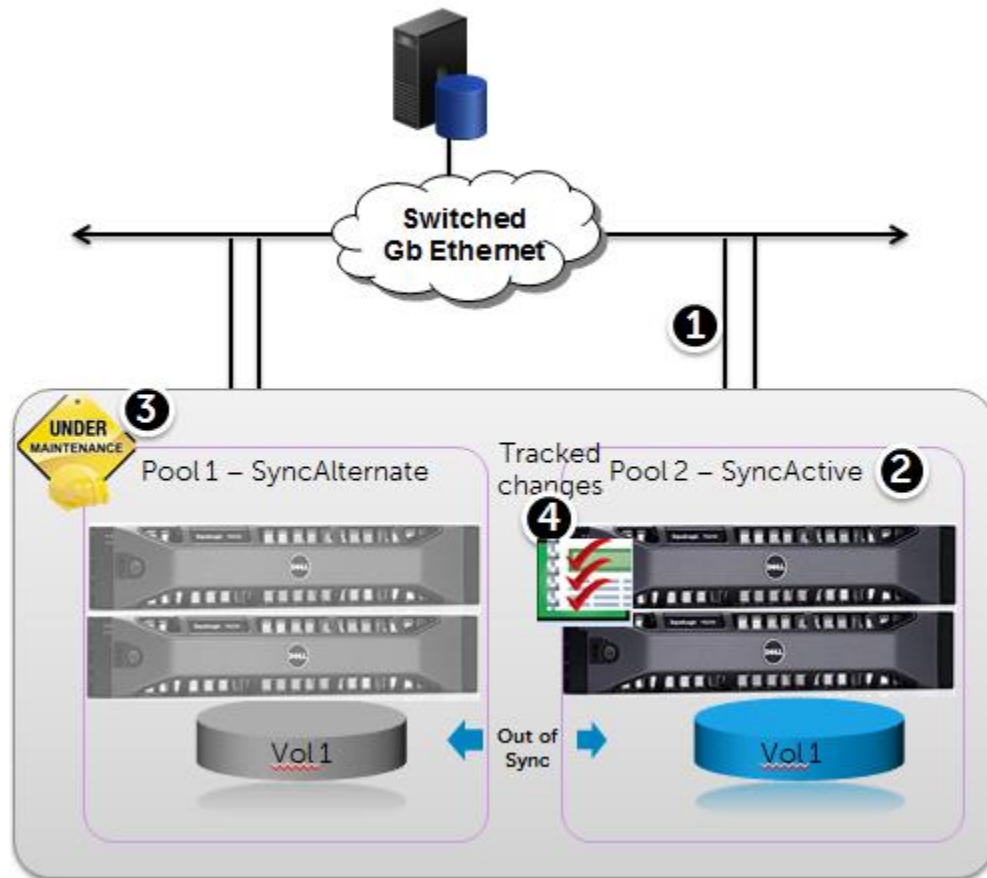
Note: Network Switches, NICs and other components should follow EqualLogic Network Reference Architecture recommendations for the Group. See network section for more information.

Multi Member Pools separated by floors within same Group.



Note: All members within a pool should be in close proximity to each other. For instance Pool1 members are all on the First Floor.

- Planned Maintenance. Switch to SyncAlternate for continued availability during scheduled outages.



Steps for maintenance:

1. **Switch** between SyncActive and SyncAlternate Pools.
2. Connect to the New SyncActive Pool – (no interruption to host I/O).
3. Perform the desired maintenance to the SyncAlternate (Pool1).
4. Changes are tracked and will update the SyncAlternate once returned to **in sync**.

Distance Considerations

Short distances are required to maintain an acceptable level of latency tolerance. Latency is dominated by two factors, the distance between the two pools that have SyncRep enabled and the algebra associated with the round trip time in TCP. Regardless of the data that needs to be transferred or the bandwidth that is available, it takes a finite amount of time for the data to travel from source to destination.

In addition to distance, many other factors such as packet loss, number of network hops; switch infrastructure and cable media type contribute to latency. To minimize the impact on the application performance, the source and destination PS Series arrays need to be located within one kilometer in distance.

Note: Recommend dedicated GigE or better, layer-2 network and switches within 1km apart. Please see the [Network](#) section in this document.

Volume I/O and application profile Considerations

- SyncRep will have the best performance with applications that have random read and write profiles. SAN HeadQuarters should be installed and used to monitor and analyze the workloads of the potential candidates for SyncRep.
- Some large sequential write profiles may degrade performance by 50% or greater on SyncRep volumes. Therefore, choosing the appropriate volumes for SyncRep is important. For instance a volume which is a backup to disk target would be a poor candidate for SyncRep.
- Consider performance guidelines for applications and databases such as SQL, Exchange and Oracle. For more information please refer to "Analyzing SAN Performance" in the **Dell EqualLogic Group Manager Administrator's Manual, PS Series Firmware Version 6.0** available on the Dell EqualLogic [support](#) site.

Note: Dell recommends using volumes with random read and light write profiles for SyncRep.

SAN Design for maximum availability and recovery

The SAN design should follow all best practices for PS Series group configuration.

- Switches, NICs and other Network infrastructure should conform to Dell PS Series EqualLogic iSCSI networking best practices and reference architectures located at the Dell EqualLogic [support](#) site or the [Dell TechCenter](#).
- For a highly available solution consider having servers on standby at the alternate location. This will increase the chances for applications to continue in the event of a localized disaster. For instance if the first floor floods with servers, switches and storage duplicated on the second floor the chances are greater for continued operations.
- Consider the PS Series generations, disk technologies and configurations when replicating between pools. For instance an active pool with PS6110XV (24 - 15K SAS Disk) Raid 10 and an alternate pool with PS5000E (16 - 7.2K Disk) configured Raid 50 may perform differently when switched depending on the I/O profiles.
- When the original SyncActive pool is repaired a re-synchronization may need to take place; consider the time and capacity needed for the changed data and bandwidth required to re synchronize the two pools.

Note: Avoid spreading members in the same pool over distance. Provide standby resources for more successful switch over to the SyncAlternate.

Snapshot space Considerations

SyncRep allocates the same percentage of snapshot reserve in both the SyncActive and the SyncAlternate pools.

For example, if you are creating a 20 GB volume and setting the snapshot reserve to 100%, you need to have a total of 40GB in each pool.

- 20 GB for the volume reserve; 20 GB of free space for snapshot reserve in the SyncActive pool
- 20 GB for the volume reserve; 20 GB of free space for snapshot reserve in the SyncAlternate pool

Note: Consider additional space for Snapshots in each pool. Remember Snapshots are located in the SyncActive pools at the time they were created.

APPENDIX B: SyncRep and Dell EqualLogic Host Integration Tools

Host Integration tools and applications such as SAN HeadQuarters, Auto-Snapshot Manager and VMware Integration tools are available to all Dell PS Series customers from the Dell EqualLogic [support](#) web site.

Auto-Snapshot Manager/Microsoft Edition (ASM/ME)

ASM/ME 4.5 provides application-consistent **Smart Copies** of SQL, SharePoint and Exchange databases.

ASM/ME 4.5 will create **Smart Copies** on the SyncActive pool. If a Smart Copy was created on the SyncAlternate pool they will not be available for mounts or restores. The ASM/ME GUI has several properties indicating SyncRep on the Volume and Collection views.

ASM/ME Impact during Switch

SyncActive-SyncAlternate volume switch (planned action) will occur for the ASM/ME protected volumes transparent to Windows. While the switch to the alternate pool is enabled any new **Smart Copies** (application consistent snaps or clones) will be located in the pool that has the SyncActive role. For a switch or failover a manual refresh in ASM/ME may be required to pick up the current status of the volumes.

Note that while **Smart Copies** (snaps or clones) will be visible through the GUI interface for both the SyncActive and SyncAlternate pools, only the current SyncActive pool **Smart Copies** will be available for mounts or restores.

SyncRep properties on the volume in ASM/ME

The screenshot displays the Oracle Enterprise Manager console for the host `w2k8r2sp1`. The left-hand navigation pane shows the tree structure under `Hosts (1)` and `Settings`, with `SyncRepASM` selected under `Applications`.

The main panel shows the properties for `G:\ (CFSyncRepV1) on w2k8r2sp1`. The **Smart Copy** section contains the following table:

Property	Value
Replication State	Not configured
Synchronous Replication	Enabled
Synchronous Replication State	In Sync
Smart Copy Count	0
Snapshot Reserve Available	98%
Snapshot Reserve Available (Alt.)	100%
Snapshot Reserve Policy	Delete oldest Smart Copy
Snapshot Space Borrowing	Enabled
Snapshot Borrow Space In Use	0 Bytes
Snapshot Borrow Space In Use (Alt.)	0 Bytes
Features	Supports Snapshot, Supports Clone, Supports Synchronous Replication

The **PS Details** section shows:

Property	Value
Type	Base volume
PS Series Group Name	tekmlab-10Gb
PS Volume name	CFSyncRepV1
PS Volume ID	68a09060-3830-6337-7e01-a580f9e5a124
PS Volume ID (reported)	68a09060-3830-6337-7e01-a580f9e5a124
Thin Provisioned Space	2.01 GB
VSS control volume status	OK

The **OS Details** section shows:

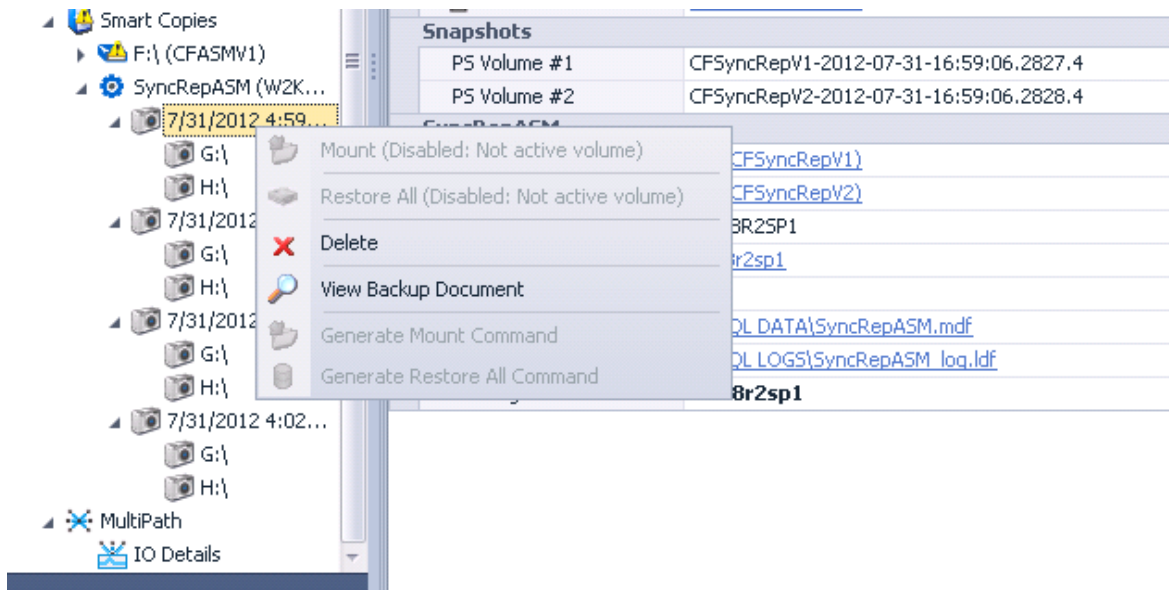
Property	Value
Volume ID	\\?\Volume{8ca67158-d11d-11e1-aca6-005056b27670}\

The **Volume G:\ (CFSyncRepV1)** summary section displays three pie charts:

- File System Allocation (20.01 GB):** 19.60 GB Free, 416.16 MB In-use.
- PS Utilization (20.01 GB):** 18.00 GB Unreserved, 1.55 GB Free, 465.00 MB In-use.
- Snapshot Reserve (2.01 GB):** 1.96 GB Free, 45.00 MB In-use.

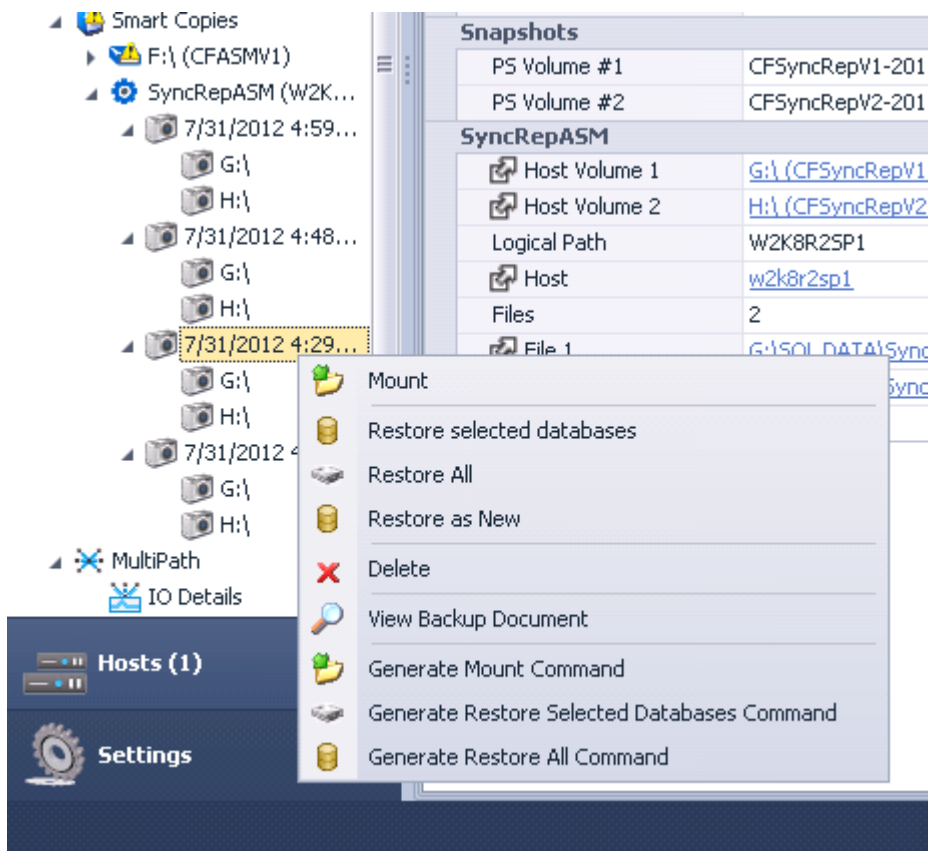
1. Synchronous Replication State
2. **Features** indicate the Volume is enabled for SyncRep.

Available ASM/ME Smart Copy operations on the SyncAlternate



Note: Smart Copies that reside in the SyncAlternate pool will only provide **Delete** or **View Backup Document** operations.

ASM/ME **Smart Copies** in the SyncActive Pool:



Note: The Current SyncActive Pool will provide all management capabilities.

VMware integration with SyncRep

Dell EqualLogic Multi-pathing Extension Module (MEM) Support:

Customers utilizing the Dell EqualLogic MEM module will need to update to version 1.1.1 for SyncRep support.

VAAI/UNMAP – Dell recommends that UnMap operations are not performed on SyncRep volumes. Please see the PS Series Storage Release Notes Version 6.0 for more information.

When the primary SyncRep volume goes offline, the naa value of the volume will change but the VM registration will still be valid. Additional steps must be taken to switch to the SyncAlternate pool, **Rescan for Datastores** and then power on the affected virtual machines.

The EqualLogic Virtual Storage Manager (VSM) v3.5 is a VMware vCenter plugin provided by Dell that is composed of tools to help administrators with storage provisioning tasks, data protection and disaster recovery and VMware View integration. VSM v3.5 has the following support for SyncRep in the following fashion:

Datastore Manager:

- Configure SyncRep on volumes from Datastores Option in VSM.
- Pause/Disable/Disconnect and Switch to SyncAlternate from Datastores Option in VSM.

Auto-Snapshot Manager/VMware:

- Create Hypervisor consistent snapshots of VMs residing on SyncRep datastores. These **Smart Copies** exist in the allocated snapshot space of the currently active SyncRep volume in the SyncActive pool. ASM/VE can determine **Smart Copies** that reside in the SyncActive pool and in the event of a switch to SyncAlternate can display any existing **Smart Copies** on that pool.

Auto-Snapshot Manager/Linux Edition (ASM/LE)

When ASM/LE creates snapshot type **Smart Copies**, they are created in the current SyncActive pool. If a SyncActive-SyncAlternate switch subsequently occurs, these **Smart Copies** will not be available for restore, (using the **asmcli restore** command), and ASM/LE will block attempts to restore from them. To proceed with the restore, the user can either switch the SyncActive-SyncAlternate back again, or choose to restore from a different **Smart Copy** which was created in the current SyncActive pool.

Note: It is possible to manually mount snapshots using the **asmcli mount** command and manually copying files to a different file system.

APPENDIX C: SAN HeadQuarters Management with SyncRep

SAN HeadQuarters is an advanced performance and health monitoring tool for PS Series Arrays. SAN HeadQuarters v2.5 and greater provides views into SyncRep activity.

SAN HeadQuarters v2.5 includes:

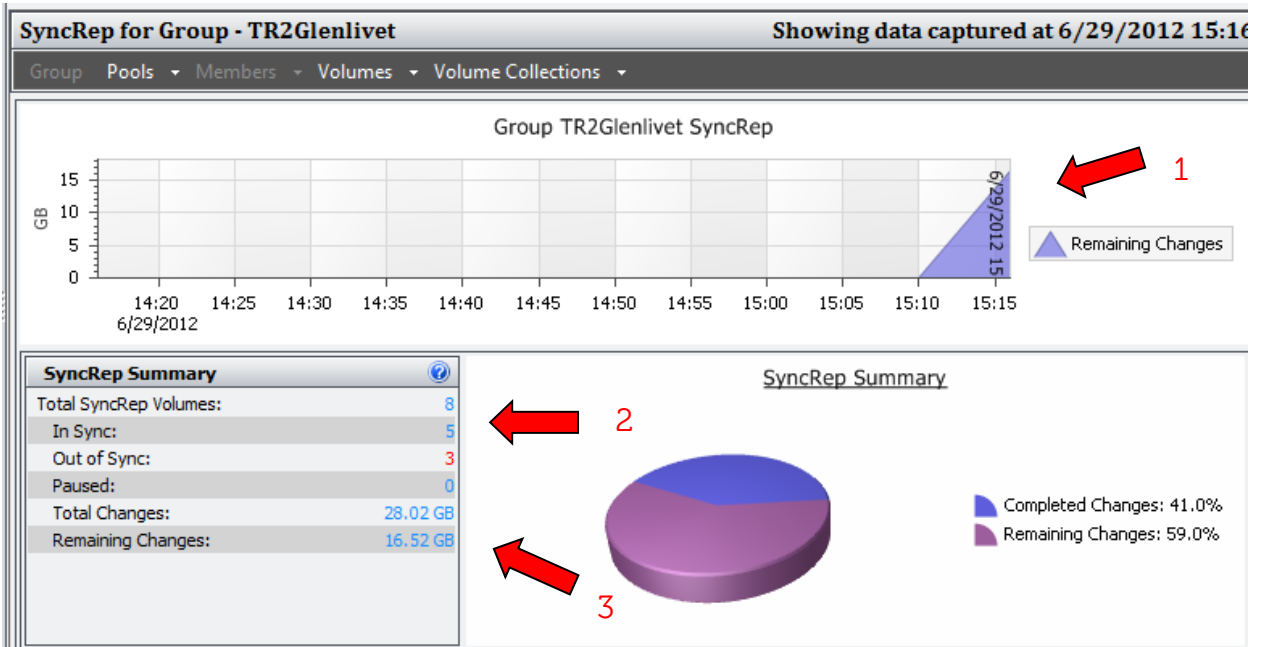
- Progress of initial or subsequent change synchronizations.
- Synchronization status over time.
- The ability to identify volumes participating in SyncRep.
- Details of SyncRep volumes at the Group, Pool or Volume level.
- Extensive tooltip explanation of Status and SyncRep definitions.
- A SyncRep report.

The SAN HeadQuarters view below shows groups eligible for SyncRep (Firmware version 6), volumes replicating and their current status.

Group / Vol...	Active P...	Alternate ...	Status	In Sy...	Out of Sync	Paused	Total Chan...	Remaining
HandsOnLab			Incompatible Firmware	0	0	0	0 KB	0 KB
NASGroup			Compatible Firmware ...	0	0	0	0 KB	0 KB
tekmkt-qlvt			Compatible Firmware ...	3	0	0	0 KB	0 KB
CF-SYNCREP	syncrep2	syncrep1	InSync	Yes	No	No	0 KB	0 KB
SkyDM-SLO2...	syncrep2	syncrep1	InSync	Yes	No	No	0 KB	0 KB
SkyDM-SQL2...	syncrep2	syncrep1	InSync	Yes	No	No	0 KB	0 KB
tekmktlab-10Gb			Incompatible Firmware	0	0	0	0 KB	0 KB
TR2Glenlivet			Compatible Firmware ...	8	0	0	0 KB	0 KB
CFSSpace 1D1	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
CFSSpace 1D2	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
CFSSpace 1D3	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
CFSSpace 1D4	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
CFSSpace 1D5	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
rsahgmtcf...	Pool1	Pool2	InSync	Yes	No	No	0 KB	0 KB
Win2012-DM1	Pool2	Pool1	InSync	Yes	No	No	0 KB	0 KB
Win2012-DM2	Pool1	Pool2	InSync	Yes	No	No	0 KB	0 KB

1. Indicates Groups that have SyncRep "Compatible Firmware" (v6.0).
2. Status of SyncRep volumes.
3. SyncRep Active and Alternate Pool identity.

The SAN HeadQuarters view below can show at the object level the activity over time for SyncRep. If the volume is undergoing replication during any point in the time selected “remaining changes” will be indicated. Also note that change tracking is provided using this view.



1. Remaining Changes for the interval – y axis will show amount of GB/MB of changes while the x axis will indicate the time.
2. Summary of SyncRep statistics for the selected time.
3. Indicates Total changes since last sync and any remaining changes at the selected time.

Technical Support and Customer Service

Dell support service is available to answer your questions about PS Series SAN arrays.

Contacting Dell

1. If you have an Express Service Code, have it ready.
The code helps the Dell automated support telephone system direct your call more efficiently.
2. If you are a customer in the United States or Canada in need of technical support, call 1-800-945-3355. If not, go to Step 3.
3. Visit support.equallogic.com.
4. Log in, or click "Create Account" to request a new support account.
5. At the top right, click "Contact Us," and call the phone number or select the link for the type of support you need.

Warranty Information

The MODEL array warranty is included in the shipping box. For information about registering a warranty, visit <http://support.dell.com/EqualLogic>.