

Application Note

January 2014

Configuring Remote Replication in an iSCSI-attached SAN Environment



Summary

This application note describes how to configure the remote replication in an iSCSI-attached SAN environment on a Windows Server 2008R2 using the SnapSAN Manager Server software.

Introduction

Remote replication solution includes methodologies and support for protecting production data while facilitating continuous business operations in the face of operational problems or large-scale catastrophic events. Remote Replication is a key storage feature which provides local copies of data to completely redundant environments at recovery sites located a great distance from the day-to-day operational environment. It provides the ability to create clones (point-in-time copies) of data with minimal impact to production.

The Master Volume (MV) or production data volume is what is being protected in case of a disaster. The copy or a clone is an identical copy of the Master Volume and can be referenced as a Replication Volume (RV). The Replication Volume is configured on another SnapSAN array located in a different building, city or state. Protection can be increased by configuring multiple RVs.

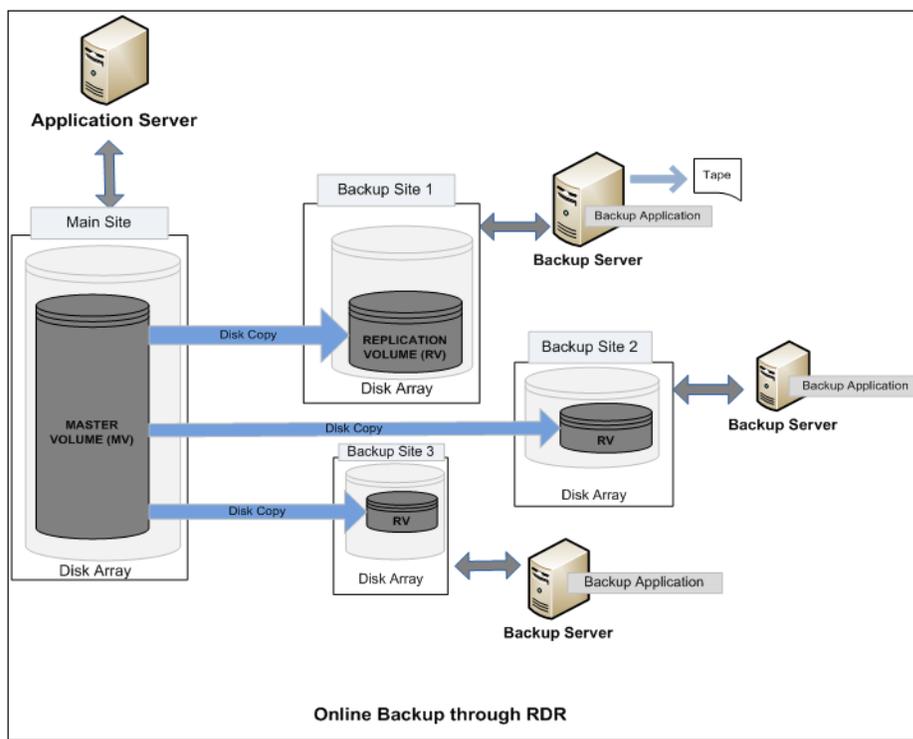
SnapSAN Storage arrays replication is Block-based (below the filesystem) using either Fibre Channel or iSCSI protocol. Data blocks changes are replicated from the MV to the RV using different replication methods. The replication methods are:

- Synchronous Mode
- Semi-synchronous Mode
- Background Copy Mode

This document provides the necessary steps in order to configure remote replication in a SnapSAN S5000 or S3000 environment. For additional information on the Remote Replication, please reference the *SnapSAN Replication/Mirroring User Guide*.

Here is an example of how Remote Replication is used in a backup server environment.

Online Backup:



Required Information, Tools, and Files

Before you begin these procedures, the following information, tools, and files are required.

Prerequisites

Prior to performing these procedures, ensure that you have the following:

- Overland Storage SnapSAN S3000/S5000 Disk Array must be installed and configured. You can get additional technical support from our website at <http://support.overlandstorage.com>, or by contacting Overland Storage using the information found on the [Contact Us](#) page on our web site.
- Verify that the following are installed on the Server:
 - Java Runtime Environment (JRE)
 - SnapSAN Manager Server Web Management Interface
- You must have two physical SnapSAN Disk Arrays that are attached together via iSCSI connectivity to each of the SnapSAN 1GbE or 10GbE controller ports, either direct connect or through an Ethernet switch.
- The examples we have provided in this document are of two separate servers. The primary application server which has the MV (Master Volume) and the backup server which has the RV (Replication Volume). This isn't a requirement, a single server can have both MV and RV assigned to it, but all precautions should be considered.
- This document assumes that the Storage Pool and Logical Disk have already been created. Additional information on Binding a Pool and Logical Disk can be found in the *SnapSAN S3000/S5000 Disk Array User Guide*:
<http://docs.overlandstorage.com/snapsan>

Versions

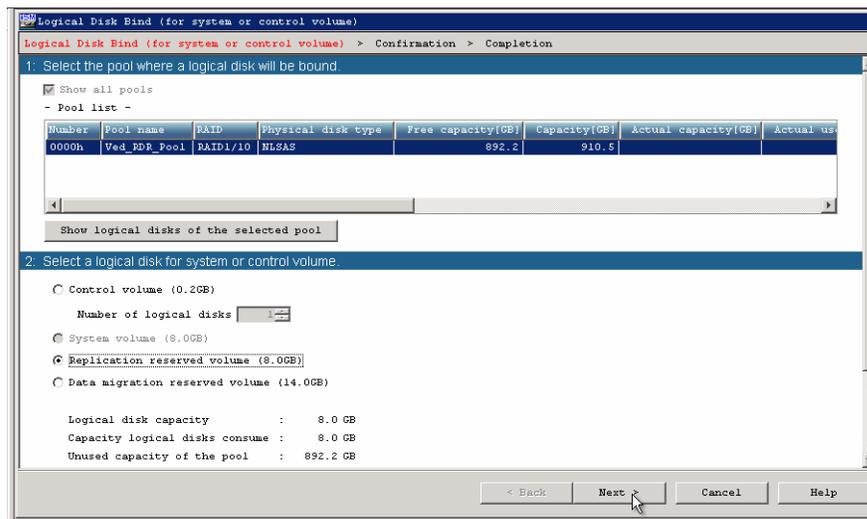
The test environment used for illustration in this document uses the following versions:

- Windows Server 2008R2
- Java JRE 7 update 40
- MS iSCSI Software Initiator version 6.1
- SnapSAN Manager Server version 8.2.060
- SnapSAN S3000 and S5000 software version 082R.007
- SnapSAN S3000 and S5000 firmware version U22R.007

Creating Replication Reserved Volume

The Replication Reserved Volume (logical disk) must be created on both the Primary and Remote SnapSAN arrays that are participating in the Remote Replication configuration.

1. From the SnapSAN Manager Monitor screen, select **Configuration > Logical Disk > Logical Disk Bind (for system or control volume)**.



2. Select the following **details**:
 - Pool name
 - Replication reserved volume
3. Click **Next**.
4. Click **Set**.
5. Click **Yes** to confirm.
6. Click **Finish**.

NOTE: Repeat these steps for any other SnapSAN array participating in the Remote Replication process.

Changing a Host Port to the Replication Port

Use the following steps to change the host connection port on the SnapSAN controller to a dedicated replication port. Changing the Host Connection Port to a Replication Port must be performed on both of the SnapSAN arrays that are participating in the Remote Replication configuration.

1. From the SnapSAN Manager Monitor screen, select **Configuration > Disk Array > Host Connection Port > Host Connection Port Settings (iSCSI)**.
2. Select the **port** to be used as a replication port and click **Edit**.
3. Change the **Port Type** to **Replication Port**, and click **OK**.
4. Click **Next**.
5. Click **Set**.
6. Click **Yes** to confirm.
7. Click **Finish**.
8. Click **Yes**.
9. Repeat [Steps 1–8](#) to create a replication port on the **remote SnapSAN array**.

Documenting the SAA ID

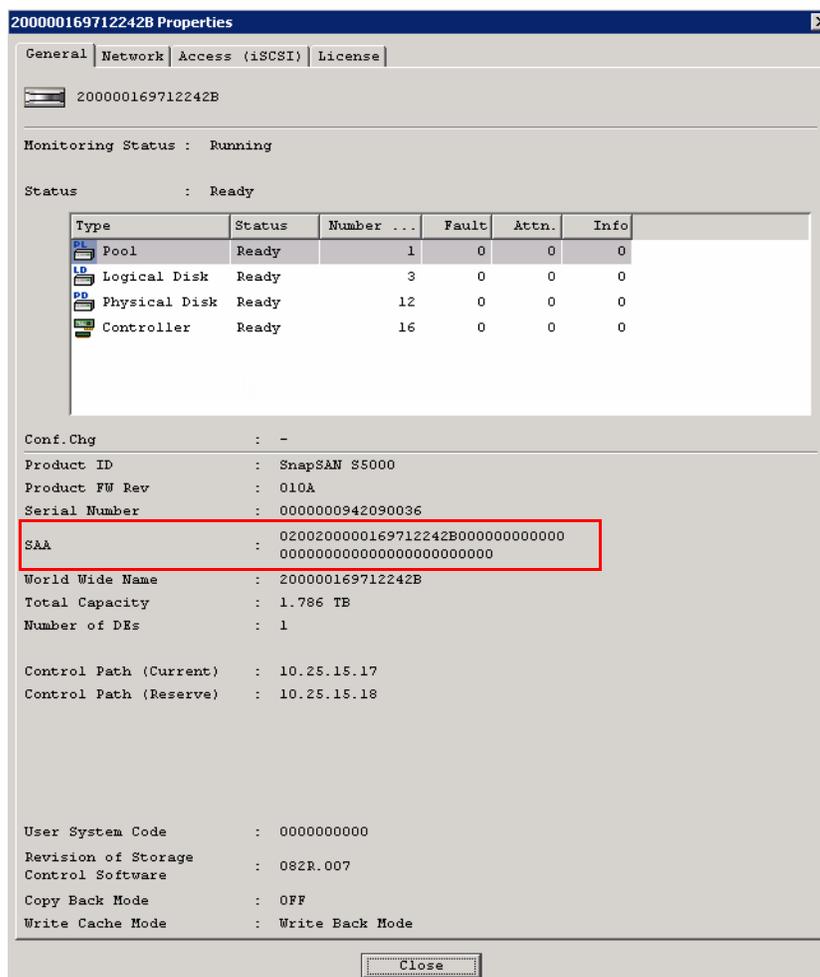
The SAA (Sub-system Absolute Address) ID is a unique identifier for each SnapSAN arrays. The SAA ID is used to configure remote replication and identifying and documenting the SAA is required on both SnapSAN arrays (Primary and Remote) that are participating in the Remote Replication configuration.

Use one of the following options in order to identify the SAA ID.

Option1: Using the SnapSAN Manager GUI

1. From the SnapSAN Manager Monitor screen, click **Screen**.
2. Right-click the disk array and select **Properties**.

Under the **General** tab, you can see the hardware information and SAA ID of the particular disk array.



Option 2: Using the command line interface (CLI)

Use the following CLI steps to gather the SAA ID:

1. Log in as **sysadmin/sys123** to the controller using SSH.

- Click **Authentication Settings** to start configuring the CHAP (Challenge Handshake Authentication Protocol) authentication.

Provide the following CHAP Authentication **details**:

- Select **Enable Initiator CHAP Authentication**.
- Provide the **CHAP Secret** password.
- Confirm the **password**.

- Click **OK**.
- Click **Set**.
- Click **Yes** to confirm.
- Click **Finish**.

You are taken back to the **LD Set Management** page from which you will see the available (and previously created) **LD set list**.

Platform	Name	Host Name	Configuration Lock	Action Mode	Target Mode
Windows (WN)	192-168-1-8			Normal	Normal
Windows (WN)	192-168-1-9			Normal	Normal
DEFAULT (DF)	rdrlink-----			Normal	Normal

- Repeat [Steps 1–11](#) for the **other** SnapSAN array.

Adding the Remote Disk Array to the Primary Disk Array

This section describes how to configure the RDR Setting. You must add both the Primary and Remote SnapSAN array in order to configure RDR settings. Repeat the steps outlined below on each of the SnapSAN arrays.

- Press and hold the **Shift** key while clicking **Configuration**.
- Navigate to **Maintenance Functions >RDR Settings**.

- Click **Add**.

Field	Recommended Input
Local Initiator alias (on demand)	Optional input.
Destination Target name	Enter the destination target IQN.
Local Initiator name	Enter the local disk array IQN.

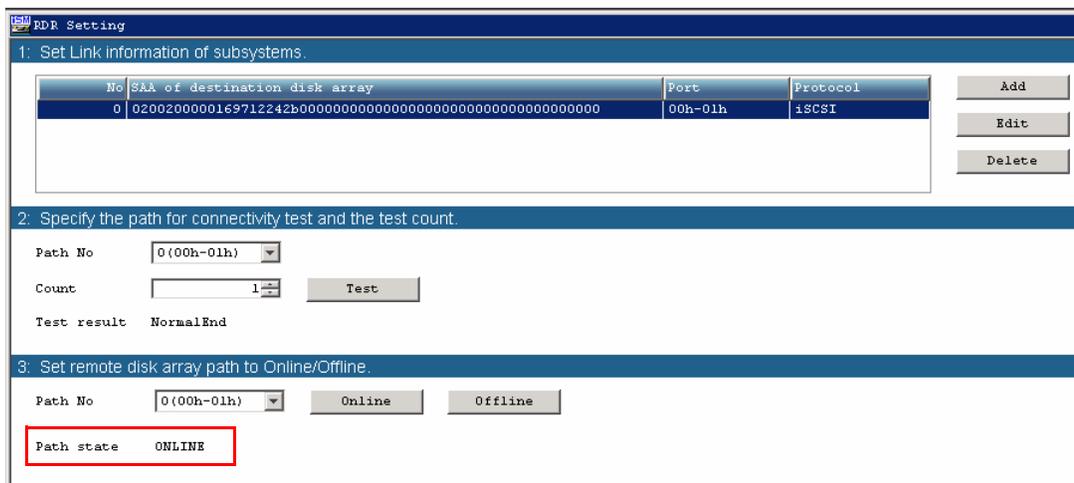
NOTE: To get the IQN information of the destination target and local initiator, login to the remote disk array through **ssh** or **telnet**, and run the following command:
`ismview -dpi`

You can see the Target Name (**Destination Target name**) and Initiator List (that is, the **Local Initiator name**) under the **iSCSI RDR Target Information** section, as shown in the following image:

```

--- iSCSI RDR Target Information ---
Platform          : DF
LD Set Name       : rdrlink-----
Target Name       : iqn.2001-03.jp.overland:storage01:ist-m000-sn-0000000942090036.df-rdrlink-----,target0002
Portal           :
Initiator List    : iqn.2001-03.jp.overland:storage01:saa-0200200000255c3a0edd0000000000000000000000000000000000000000000000
    
```

6. Click **OK**.
7. On the **RDR Settings** page, click **OK**.
8. Click **Yes**.
9. Select the added **SAA ID**.
10. Click **Online**.



The **Path state** types:

- **Offline** – Indicates an error in the configuration.
- **Freeze** – Indicates that the data replication function in the linked disk array is frozen. The Freeze path state automatically changes to **Online** within few seconds, and, if it does not, then click the **Online** button to manually change the path state to **Online**.

11. Repeat [Steps 1–10](#) to add the **local** SnapSAN array to the remote array.

Testing the Connectivity

Follow this procedure to test the connectivity between the Remote and Primary disk arrays.

1. From the **RDR Settings** page, click **Test**.
2. Click **Yes**.

A successfully established connection results in a **Normal End** status.

NOTE: Configuring the SAA on a single SnapSAN array will not establish a proper connection and the test results fails as **Abnormal End** status. Thus it is imperative to establish a proper connection, and configure both SAA on each SnapSAN array.

Configuring the Remote Replication Environment

Pairing MV and RV

1. From the SnapSAN Manager Monitor screen, select **Option > Replication**.
2. From the **Disk Array Link Information** tab, check if the disk arrays are properly linked for the remote replication.

Link ...	Link ...	Link Disk ...	Link Mode	Path Number	Path State	Dirac
00h	iSCSI	200000255C...	Normal	00h	Normal	01h

For example, if a SnapSAN S5000 is added to a SnapSANS3000 but the S3000 has not been added to S5000, then clicking the **Disk Array Link Information** tab displays a blue dot icon under the **Link** column and the **Path State** status appears as **Freeze**.

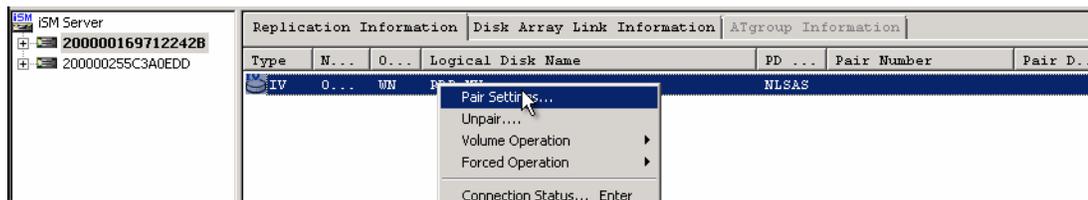
Link ...	Link ...	Link Disk Array Name	Link Mode	Path Number	Path State
00h	iSCSI	200000169712242B	Normal	00h	Freeze

3. After verifying the proper link connection, pair the MV and RV by using either Option 1 or Option 2:

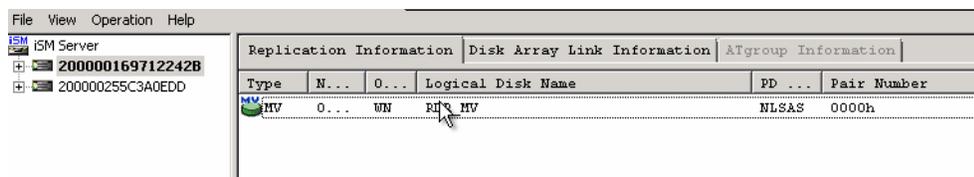
Option 1: Using the SnapSAN Manager GUI:

- a. In the left pane, click the **Primary SnapSAN array** (from which the data has to be replicated).

- b. In the **right** pane, right-click the **disk** to be replicated (known as the Master Volume or MV) and select **Pair Settings**.



- c. Provide the following **details**:
- Select **Remote Data Replication**.
 - In the field labeled **LD Set Name**, use default **ALL** for both the MV and RV sections.
 - In the field labeled **Logical Disk**, select the **master volume** under the MV section and the **replication volume** under the RV section.
- d. Click **Pair**.
- e. Click **Yes** to confirm.
- f. Click **OK**.
- g. Once the pairing has finished, close the **Pair Settings** dialog box.



Option 2: Using the command line interface:

- a. Login as **sysadmin/sys123** to the controller through SSH.
- b. Run the following command line to **pair** the MV and RV:

```
iSMrc_pair -pair -mv MV_rdr -mvflg ld -rv RV_rdr -rvflg ld
```

```
login as: sysadmin
Using keyboard-interactive authentication.
Password:
Last login: Wed Dec 11 21:54:48 2013 from cnt1
sysadmin@200000169712242B-0# iSMrc_pair -pair -mv RDR_MV -mvflg ld -rv VED_RDR_RV -rvflg ld |
iSMrc_pair: Info:      iSM13248: Command has completed successfully. (already) (code=0b0d-0d03-0003-0000)
ExitStatus:3
```

Testing Remote Replication

As a best practice, always test the replication environment after initial configuration. Use the following steps in order to copy sample data from the Primary MV to the Remote RV. The following steps describe how to prepare the MV and RV through Windows Disk Management.

1. From the Primary Windows host server that has the Master Volume or production disk, select **Computer Management > Disk Management**.
1. Change the status of the assigned disk to **Online**.
2. Initialize the **disk**.
3. Format the **assigned disk**.
4. Copy or create some **sample data** to the prepared Master Volume.

5. On the Remote backup host server, initialize the **RV** leaving it unformatted.

Performing the Remote Replication

1. To start replication, you must change the RV disk to an **Offline** status on the **Remote Server**.

On Windows Server 2008R2 and newer Operating Systems, the disk can be taken offline using Windows Disk Management. To do this on Windows 2003 and older Microsoft Operating Systems, the RV must be unassigned through the SnapSAN Manager from **Configuration > Host > Assignment of Logical Disk**.

2. From the left pane of the Repli interface in SnapSAN Manager, select the particular disk array from which data is being replicated, then right-click the **MV** and select the path **Volume Operation > Replicate**.
3. In the **Replicate** dialog box, under **Copy Conditions**, select the following **details**:
 - **Range** of replication as **All**.
 - **Copy Mode** as **Synchronous**.

Below is a sample table of copy control state:

Copy Control State		Copy State
Foreground Copy	Sync Mode	<ul style="list-style-type: none"> • Copy for reflecting difference is performed. • Updates to MV are reflected RV sequentially.
	Semi Sync Mode	<ul style="list-style-type: none"> • Copy for reflecting difference is performed. • I/O of updates to MV is completed when data is written to MV, and the data is copied to RV immediately after that. • Cannot be set to a pair (DDR) in the same disk array.
Background Copy	-	<ul style="list-style-type: none"> • Copy for reflecting difference is performed. • I/O of updates to MV is completed when data is written to MV, and the data is accumulated as difference information. For the accumulated difference, data is copied to RV asynchronously. The copy interval to RV (background copy level) can be changed in units of disk arrays. • When performing Separate, the state must be changed to Foreground Copy in advance.
Suspend	Suspend	<ul style="list-style-type: none"> • Copy for reflecting difference is not performed. • I/O of updates to MV is completed when data is written to MV, and the data is accumulated as difference information. Reflection to RV is not performed.
	Suspend due to a failure	<ul style="list-style-type: none"> • Forcefully suspended in the disk array due to a copy fault. • Copy for reflecting difference is not performed. • I/O of updates to MV is completed when data is written to MV, and the data is accumulated as difference information. Reflection to RV is not performed.

- **RV Status** as **Not Ready**.

4. Click **Replicate**.

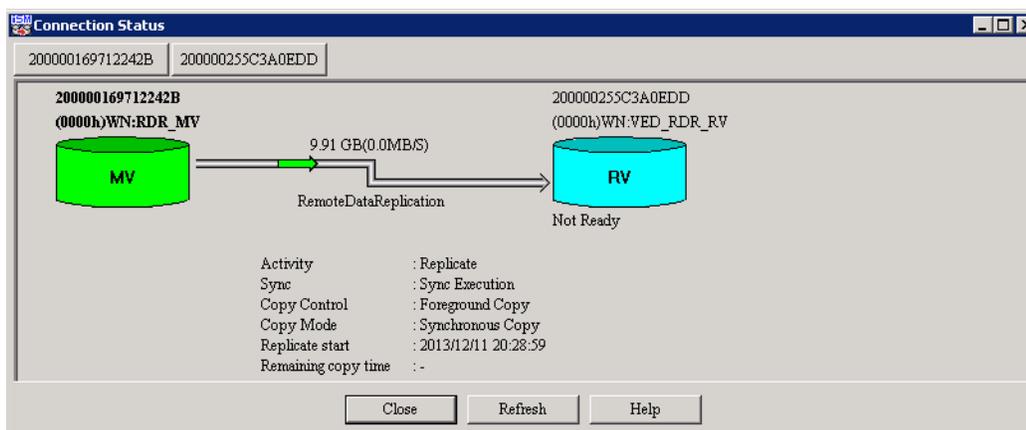
- If RV is dismounted, click **Yes**.



Once the replication starts, you can close the replication dialog box.

Checking the Replication Status

- From the Repli interface, select the Primary SnapSAN Array where the MV is located by double-clicking **MV** (or right-clicking the **MV** and select **Connection Status**.)
- Check the **details** of the data replication from MV to RV.



Verifying the Replicated Data

Separate the Connection between MV and RV

- From the Repli interface, select the **Primary SnapSAN array**.
- Right-click **MV** and select **Volume Operation > Separate**.
- Click **Separate**.
- Click **Yes** to confirm.

Assign RV to the Host

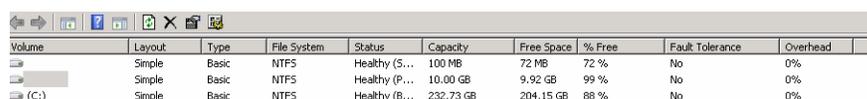
If the RV is not assigned to the host, use the following steps to assign it:

- From the SnapSAN Manager Monitor screen, select **Configuration > Host > Assignment of Logical Disk**.
- Select the **Host** and the **Logical Disk** for assignment.
- Click **Next**.
- Click **Set**.

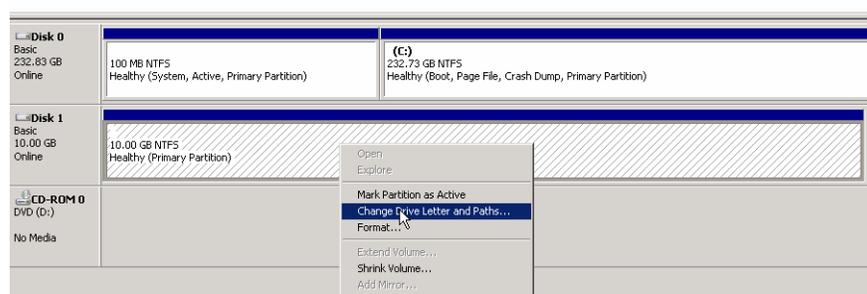
5. Click **Yes** to confirm.
6. Click **Finish**.
7. Click **Yes**.

Change the Drive Letter

1. From the remote backup host server, launch **Windows Disk Management**.
2. Under **Storage**, rescan the disks or refresh the details to rediscover the **RV disk**.
3. Find the new disk with no drive letter, right-click it, and select the option **Change Drive Letter and Paths**.

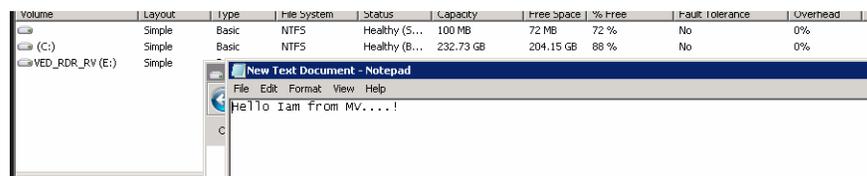


Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free	Fault Tolerance	Overhead
	Simple	Basic	NTFS	Healthy (S...)	100 MB	72 MB	72 %	No	0%
	Simple	Basic	NTFS	Healthy (P...)	10.00 GB	9.92 GB	99 %	No	0%
(C:)	Simple	Basic	NTFS	Healthy (B...)	232.73 GB	204.15 GB	88 %	No	0%



4. Click **Add**.
5. Select the following **details**:
 - **Assign the following Drive Letter**
 - **Drive**
6. Click **OK**.

The replicated data should now be the same as what was created on the Master Volume.



Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free	Fault Tolerance	Overhead
	Simple	Basic	NTFS	Healthy (S...)	100 MB	72 MB	72 %	No	0%
(C:)	Simple	Basic	NTFS	Healthy (B...)	232.73 GB	204.15 GB	88 %	No	0%
VED_RDR_RV (E:)	Simple	Basic	NTFS	Healthy (B...)	10.00 GB	9.92 GB	99 %	No	0%

Restoring Data from the RV to the MV

Fail-over to the RV is used primary in Backup or Disaster Recovery. The following steps describe how to use the Restore option or Fail-back from the RV to the MV (in reverse) in order to update the MV that had changes made while the MV was unavailable.

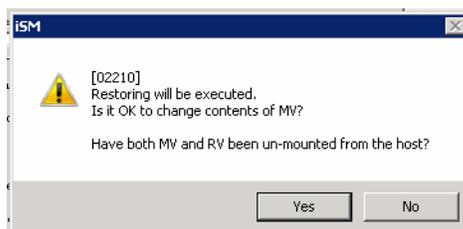
Prerequisites:

- The MV must not be formatted (delete the partition and make it unallocated).
- The volumes MV and RV must both be dismounted.

NOTE: When restoring the data without dismounting MV and RV from hosts, the restore task will be completed but with errors and after completing the restore process, no data will be restored.

To start data restorations from the RV to MV use the following steps:

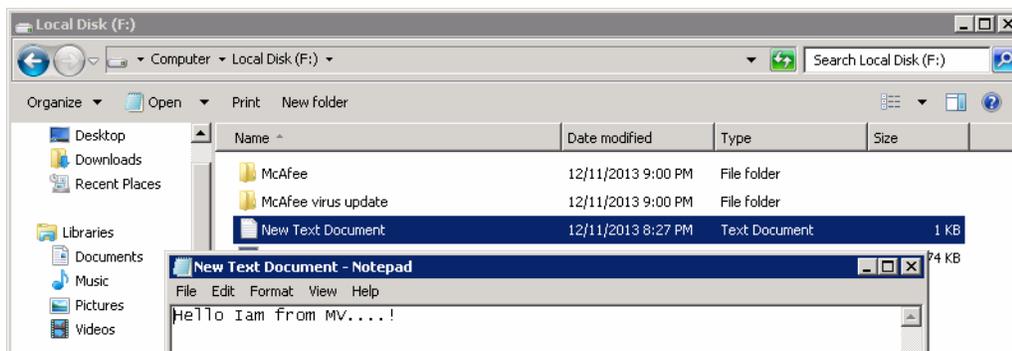
1. From the Repli interface, select the particular **disk array** in which data is replicated, right-click the RV, and select **Volume Operation > Restore**.
2. Select the details under the section **Restore Conditions** as per your requirements, and click **Restore**.
3. Do **one** of the following;
 - If MV and RV are dismounted, click **Yes**.
 - If they are not dismounted, first dismount them and then click **Yes**.



Verifying the Restored Data

To verify the restored data:

1. Separate the **connection** between MV and RV
2. Assign MV to the **host**.
3. If required (not found), change the **drive letter**.
Details for changing the drive letter can be found in **Change the Drive Letter**.
4. From the Primary host server, launch **Windows Disk Management**.
5. Under **Storage**, rescan the disks or refresh the details to rediscover the **MV disk**.
6. Find the new disk with no drive letter, right-click it, and select the option **Change Drive Letter and Paths**.



7. The replicated data should be the same as what was created on the Replication Volume.

Performance Analysis using Performance Monitor

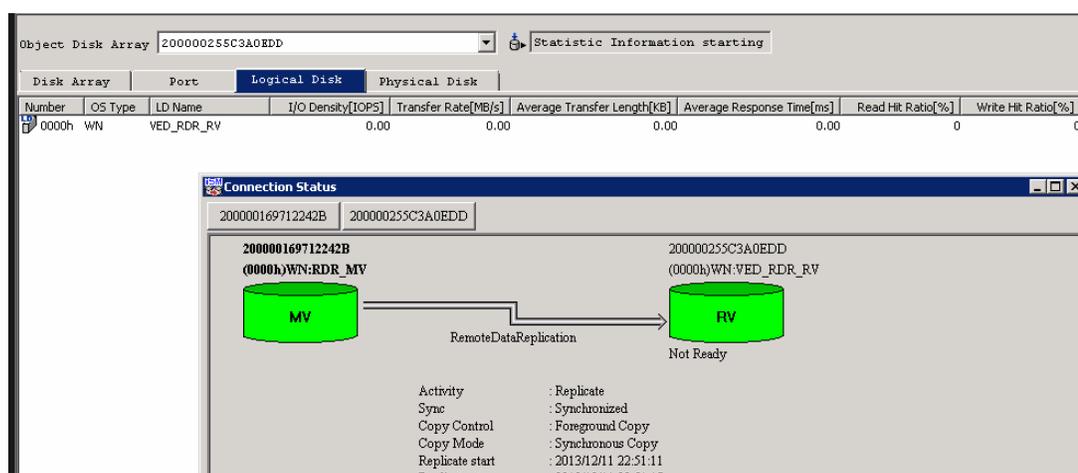
This section discusses the Synchronous and Semi-synchronous replication modes and their performances. For more details on how to configure the Performance Monitor, refer to the *Performance Monitor User Guide*.

Synchronous Replication

In Synchronous Replication mode, the data copy from MV to RV begins when the write request command is issued, and is done when the data copy has completed and the write request command completion notification is issued to the host:

1. The write request data from the host is stored in the cache memory area of the MV.
2. The write data from the cache memory is transferred to the remote SnapSAN array and stores the data in the cache memory of the RV.
3. Completion of the write request is sent to the host.
4. The data from cache memory is written to each physical disk of the MV and RV.

The example shows the performance details on an active Synchronous Replication:



Semi-synchronous Replication

The Semi-synchronous mode, the host is notified of the completion of the write request at the completion of writing to the MV, and then the data is copied from the MV to the RV. Even if the RV is located a great distance from the MV, this mode has less impact on I/O performance than Synchronous mode.

1. The write request data from the host is stored in the cache memory area of the MV.
2. Completion of the write request is sent to the host.
3. The write data from the cache memory is transferred to the remote SnapSAN array and stores the data in the cache memory of the RV.
4. The data from cache memory is written to each physical disk of the MV and RV.

