

Dell PowerVault MD Series Storage Arrays Storage Replication Adapter (SRA) Best Practices Guide (Client)



Notes, cautions, and warnings



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.



WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Copyright © 2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

2015 - 09

Rev. A00

Contents

1 Installing and configuring Storage Replication Adapter (SRA).....	4
Downloading SRA.....	4
Installation procedure.....	4
Password-protected storage arrays.....	5
NVSRAM settings.....	6
SRA device management service.....	9
Changing server settings in SraConfigurationData.xml.....	9
Changing virtual memory in NesSvc.ini.....	10
2 Asynchronous Remote Replication.....	11
iSCSI-Based Remote Replication.....	11
Support for four Remote Replication groups.....	11
Effects of 10-minute synchronization interval.....	11
General volume recommendations.....	12
Command line options.....	12
Advanced Settings in Site Recovery Manager and ESX/ESXi.....	12
ESX/ESXi host settings.....	13
3 Snapshot Repository sizing.....	14
Snapshot Group repository.....	14
Snapshot Volume repository.....	14
How SRA uses snapshots.....	15
4 Getting help.....	17
Documentation matrix.....	17
Dell documentation.....	17
VMware documentation.....	17
Contacting Dell.....	17
Locating your system Service Tag.....	18

Installing and configuring Storage Replication Adapter (SRA)


Downloading SRA

The Dell MD Series SRA is used with VMware Site Recovery Manager (SRM) to facilitate data center failover between separate vCenter Server environments. To use the SRA, download the latest version from the **Drivers and Download** page available at **Dell.com/support**.

- For the latest supported versions of SRA and VMware reference, see the Support Matrix available in the Manual section of your array at **Dell.com/support**.
- md5sum: You can find the md5sum value at the VMware Site Recovery Manager website for Dell MD Storage Replication Adapter.

You can calculate the Md5sums on any UNIX host with md5sum installed or by obtaining a Windows utility like md5sum.exe from etree.org/md5com.html and running the following command at command line interface (CLI).

```
md5sum <file_name>
```

 **NOTE:** Any references to filenames are for example purposes only.

If the file is downloaded with the installer, you can run the following command to verify the installer package.

```
SRAInstaller-xx.xx.xxx.xxx.xxx
```

```
[root@bvr-wb2 SRA]# ls -l
total 37980
-rw-r--r-- 1 root root      67 Oct  1 11:36      Installer.md5
-rw-r--r-- 1 root root 38838568 Oct  1 11:19 SRAInstaller-05.00.3050.0017.exe
[root@bvr-wb2 SRA]# md5sum -c      Installer.md5
SRAInstaller-05.00.3050.0017.exe: OK
[root@bvr-wb2 SRA]#
```

Figure 1. Example MD5 Evaluation

Installation procedure

After verifying that the downloaded file is complete and not corrupt, copy the installer to your intended SRM servers and run the SRA installer on those servers. You can view the latest information contained in the **readme.txt** file at the end of the installation by clicking **Yes**.

SRA is installed at the following location:

On x64 hosts C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\storage\sra

On x86 hosts C:\Program Files\VMware\VMware vCenter Site Recovery Manager\storage\sra

After SRA is installed, rescan for SRAs from the Site Recovery manager in the vSphere Client.

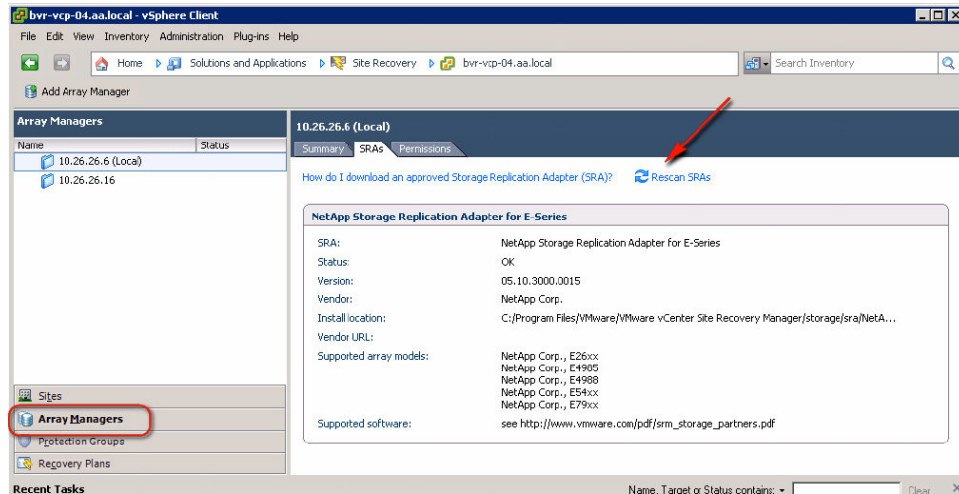



Figure 2. Site Recovery Manager (Rescan SRAs)

Password-protected storage arrays

If your environment implements password security on the storage arrays, modify the `SraConfigurationData.xml` file to prompt for the storage array password.

To modify the files:

1. Edit the `SraConfigurationData.xml` file available at `C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\storage\sra\Dell\config\SraConfigurationData.xml`
2. Locate the `<PasswordRequiredForArrayAccess>` tag.
3. Change the default value of `false` to `true`.
4. Save the file changes.

 **NOTE:** All storage arrays must utilize the same security measures. If one storage array has a password set, then the peer storage array must also have the password set. SRA does not support Mixed authentication mode.

```
configure how array access is performed.
when true, a password will be prompted for once and
then used for all array access
-->
<PasswordRequiredForArrayAccess>true</PasswordRequiredForArrayAccess>
```

Figure 3. SraConfigurationData.xml Password Value

NVSRAM settings

SRM test failover requires a change to the default NVSRAM settings on the MD storage array. During test failover, snapshots are created on the recovery site's storage array which is then mapped to the default host group and the ESX or ESXi host group.

Before completing the following tasks, the RAID controllers must be rebooted. Dell recommends you to complete the following tasks during maintenance or in a preproduction environment.

To change the NVSRAM settings:

- 1. Start **MD Storage Manager** on your storage array.
- 2. From the **Enterprise Management Window (EMW)**, click **Tools** → **Execute Script**.

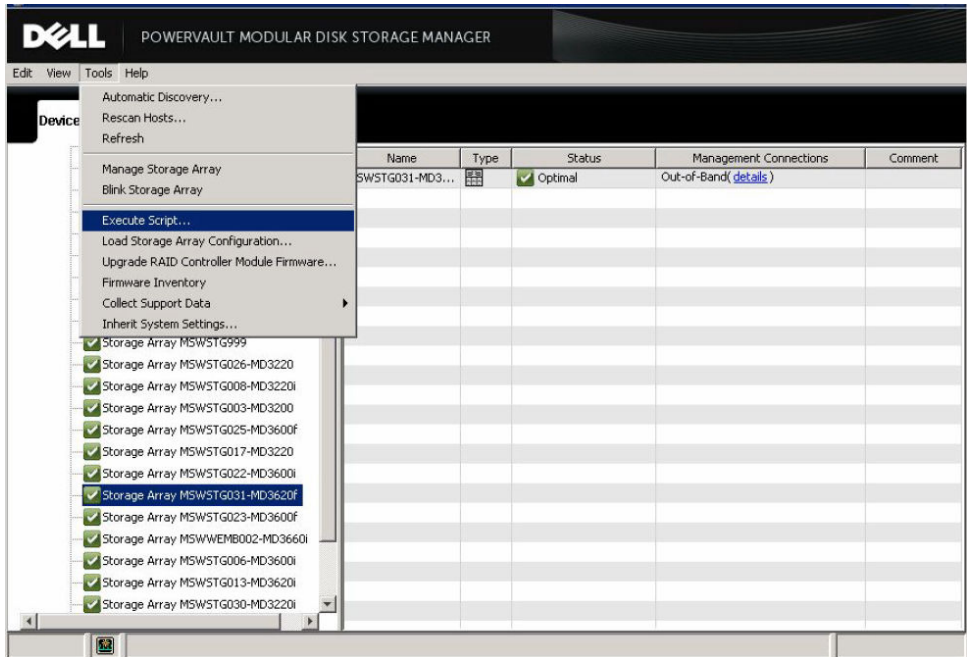


Figure 4. MD Storage Manager Execute Script

- 3. In the **Script Editor** window, to review your current NVSRAM settings, run the following command:
`show controller[0] NVSRAMByte[0x3b];`

NOTE: To test and run your command, on the **Script Editor** page, click **Tools** → **Verify and Execute**.

Your current NVSRAM setting for the specified RAID controller [0] is displayed.

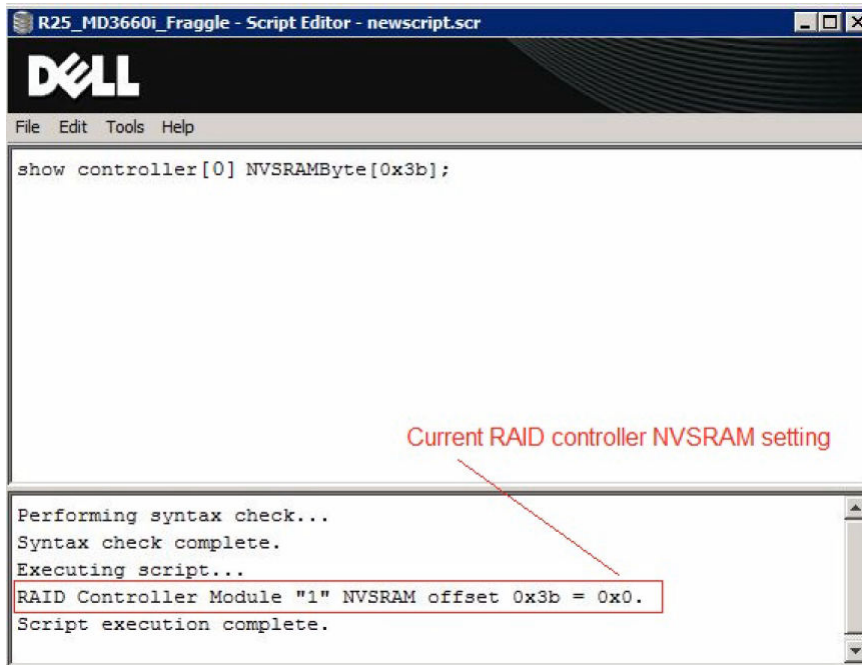


Figure 5. Review Current NVSRAM Settings

4. To change the NVSRAM setting of your primary RAID controller [0], run the following command.
`set controller[0] NVSRAMByte[0x3b]=2;HTTP/1`

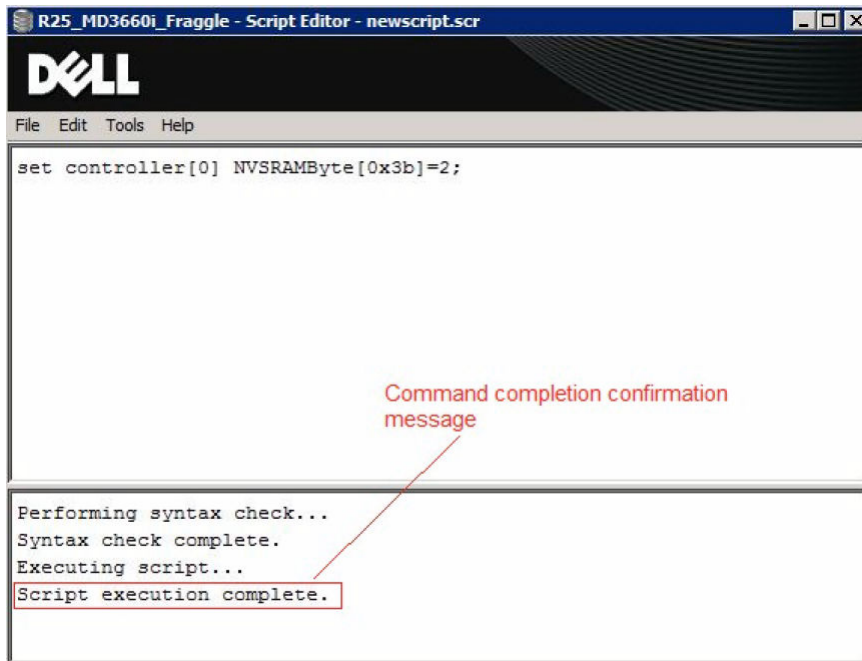


Figure 6. Change NVSRAM Setting on Primary RAID Controller

5. Ensure that your NVSRAM settings are changed by running the **show controller** command.
`show controller[0] NVSRAMByte[0x3b];`




Figure 7. Confirm New NVSRAM Setting

6. Reset the primary RAID controller module to load the new NVSRAM setting into controller memory:
`reset controller[0];`



NOTE: The reset controller command takes several minutes to run. Do not perform I/O operations involving the RAID controller until the command successfully completes.




Figure 8. Reset Primary RAID Controller to Load New NVSRAM Settings

7. Repeat task 3 through task 6 on the secondary RAID controller by substituting [1] for [0] in the command syntax examples shown.
8. Close the **Script Editor** dialog box.

SRA device management service

The process monitors and synchronizes communications between the MD SRA and the MD storage arrays. SRM workload is partitioned between a persistent server and multiple transient client programs which endure only during a single SRM command. Because many SRM commands are expected to be run during a single SRM workflow (for example, test failover), clients are frequently used. Servers, however, running as a Windows Service, persist and manage all communications with storage arrays. No special configuration other than default values is required for the server. However, two files can be edited to control host configuration, socket communication and virtual machine (VM) memory:

SraConfigurationData.xml	This file resides in the /config directory under the SRA installation directory. The specific location depends on your installation selections, the default location is C:\Program Files(x86)\VMware\VMware vCenter Site Recovery Manager\scripts\SAN\ Dell\SraConfigurationData.xml.
NesSvc.ini	This Win32 Service initialization file resides in the /win32svc directory under the same SRA default installation directory as the SraConfigurationData.xml.

Changing server settings in SraConfigurationData.xml

The portion of SraConfigurationData.xml file relevant to service configuration is:

```
<SraService>
  <SvcHost>localhost</SvcHost>
```

```
<ServicePort>1701</ServicePort>
<ListenBacklog>100</ListenBacklog>
<SraService>
```

- <SvcHost>** Specifies which host the Win32 Service initialization file (`NesSvc`) runs on. Currently, the only supported value is **localhost**. In future releases, it may be possible to share instances of `NesSvc` across multiple installations of SRA, which enhances performance and simplifies cooperation between multiple SRA instances.
- <ServicePort>** Determines which IP port the service uses for socket communications between the client and server. If different application on your system is already using port 1701 (default port), specify a different port.
- <ListenBacklog>** Configures a performance property of the port.
- >**



NOTE: Contact Dell technical support to modify the `<ListenBacklog>`.

Changing virtual memory in `NesSvc.ini`

No modifications are necessary to the Windows service initialization file (`NesSvc.ini`). If recommended by Dell technical support personnel, you can change the virtual memory settings to:

```
vmarg.1=-Xms256m
vmarg.2=-Xmx512m
```



NOTE: Any change to the `NesSvc.ini` file requires you to stop and restart the service in order for the changes to take effect.

Asynchronous Remote Replication

This feature performs remote replication using Point-in-Time (PiT) snapshots and is supported on both Fibre Channel and iSCSI MD storage arrays.

iSCSI-Based Remote Replication

With added support for Remote Replication over iSCSI protocol, greater replication distances are now possible. However, this improvement in replication distances impacts latency (compared to earlier Fibre Channel-only support). Therefore, take careful consideration during Datastore creation to ensure that only data that must be replicated is included on the virtual disks you replicate. After replication begins, closely observe the amount of data being replicated and the time it takes to synchronize data. If the amount of time required to synchronize data is more than the synchronization interval, a remote replication group becomes degraded and nonfunctional.

Proper sizing of your network infrastructure is a critical component of a successful disaster recovery solution. For more information, see the following documentation:

- *Dell PowerVault MD Series Storage Arrays Administrator's Guide* at Dell.com/support/manuals.
- VMware Site Recovery Manager Documentation Center at vmware.com/support/pubs.

Support for four Remote Replication groups

SRA supports up to four Remote Replication groups per storage array. All protected Datastores must reside in one of the four groups. Because a Remote Replication group is treated as a single entity, any action affecting the group (for example, swapping roles) affects all virtual disks in the group.

If cross-replication of Datastores is required (for example, replicating from the recovery site to protected site), the Datastore virtual disks for the recovery site must be contained in a separate Remote Replication group than the virtual disks on the protected site.

Effects of 10-minute synchronization interval

Remote replication groups require a 10-minute interval between automatic and manual synchronizations. If you request a manual synchronization when a Remote Replication group is inside a 10-minute synchronization interval, the manual synchronization does not begin until the previous synchronization interval has passed. This interval may cause a delay in the SRM workflow process which requires several synchronization operations to occur during test failover and/or failover workflows. SRA is optimized to avoid requesting a manual synchronization if no changes are detected within the Remote Replication group. However, if changes are detected, a synchronization is requested. The effect may be observed as a lack of progress or slow progress through SRM workflows.

General volume recommendations

When designing a disaster recovery strategy using VMware Site Recovery Manager and MD Series storage arrays, keep in mind these considerations:

- Protection works on a datastore level (storage array virtual disk). All VMs on the same datastore that requires protection are also protected and replicated.
- Multiple small-sized datastores and virtual disks must be used to limit the amount of data replicated to the recovery site.
- Locate (migrate) protected VMs to the same datastore and migrate any VMs that do not require protection to other locations.
- SRM does not provide for application-consistent failover, but does provide VM-consistent failover. Therefore, even with a successful failover of a VM, the applications that are running on the VM may not be in a consistent state. More recovery methods may be required to return to normal operation.
- Only MD Series storage arrays running firmware versions 07.84.XX.XX or later support asynchronous Remote Replication groups. The groups are treated as consistency groups, so all virtual disks in the group is treated as a single entity.
- Synchronization priority dramatically affects the speed of replication, assuming that sufficient bandwidth is available between the protected and recovery storage arrays.

Command line options

SRA installs a CLI utility that provides the following functions:

- To stop the SRA service, run the following command.
`<SRA_Path>\svrCmd.cmd`
Restarting the SRA service requires a start command from Windows services.
- To track the SRA service and capture detailed logging for each command issued and received, run the following command.
`<SRA_InstallationPath>\svrCmd track on (enable)`
`<SRA_InstallationPath>\svrCmd track off (disable)`
Log files are written to `<SRA_InstallationPath>\track` directory.

Advanced Settings in Site Recovery Manager and ESX/ESXi

The following advanced settings in Site Recovery Manager (SRM) are recommended for best performance when using the MD Series Storage Replication Adapter (SRA).

To establish these settings:

1. Click **Sites** in the SRM left pane, then right-click your site name and click **Advanced Settings**.

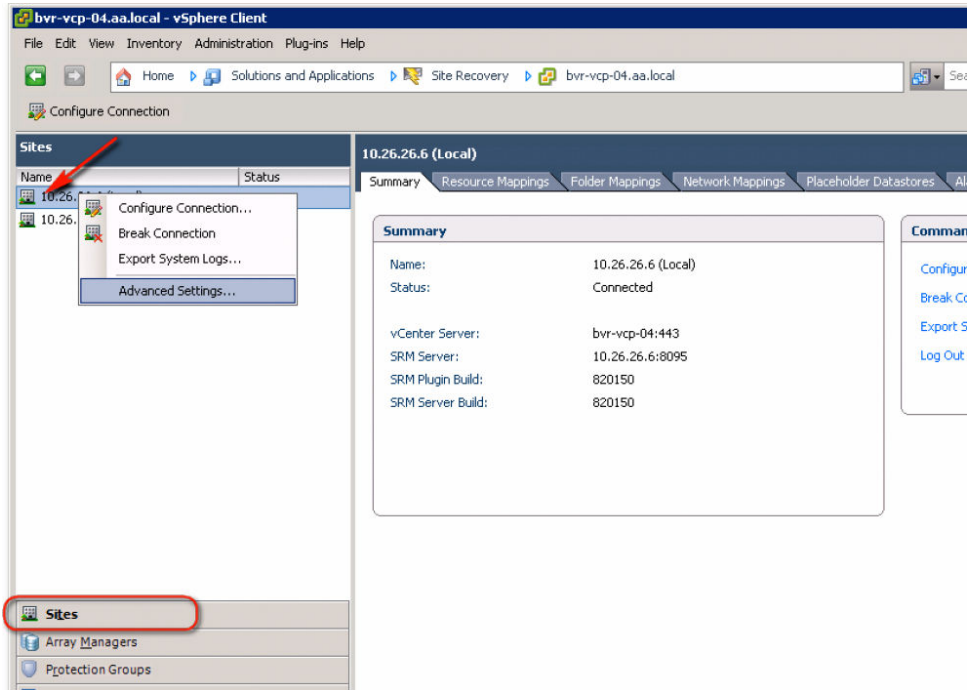


Figure 9. Site Recovery Manager Advanced Settings

2. Click storage and set the following values:
 - `storage.commandTimeout = 900`
 - `storageProvider.hostRescanRepeatCnt = 2`
 - `storageProvider.hostRescanTimeoutSec = 900`
3. Click **OK** to save your changes.
4. Click **storageProvider** and set the following values:
 - `storageProvider.fixRecoveredDatastoreNames = enabled`

ESX/ESXi host settings

The following changes are recommended for ESX or ESXi host settings:

- `Disk.MaxLUN` — Set this value slightly larger than the number of LUNs mapped to the ESX or ESXi host. This provides faster rescan operations by not scanning all 256 LUN possibilities.
- `Disk.UseDeviceReset = 1` and `Disk.UseLunReset = 0`, `Disk.UseDeviceReset` and `Disk.UseLunReset` are used together to indicate how device resets are issued.

Snapshot Repository sizing

A new feature of firmware 07.83.xx.xx are Point-in-Time Snapshots. These provide the ability to roll back snapshots to earlier point-in-time saves and optimize the data changes between snapshot images. This feature utilizes two separate repositories to facilitate tracking of changes to the base volume. They are the Snapshot Group repository and Snapshot Volume repository.

Snapshot Group repository

The Snapshot Group repository is used to track data changes to the base volume (volume that the Snapshot Image was created from). The Snapshot Group repository may contain multiple Snapshot Images (point-in-time records of base volume). A Snapshot Volume is created from these images and can then be mapped to a host for access.

Snapshot groups

Name	Status	Total Repository Capacity	Available Repository Capacity	Snapshot Image Limit	Snapshot Images	Scheduled
SRA_Primary_1_SG_01	✓ Optimal	10,000 GB	9,999 GB (100%)	32	0	No

Snapshot groups: Associated repositories

Associated Snapshot Group	Name	Status	Capacity
SRA_Primary_1_SG_01	repos_0003	✓ Optimal	10,000 GB

Figure 10. Modular Disk Storage Manager (MDSM) Snapshot Group View

Snapshot Volume repository

The Snapshot Volume repository is used to track data changes to the Snapshot Volume, if read/write access is allowed. After a Snapshot Volume is mapped to a host for access, any changes to the volume are tracked within this repository.

Associated replicated pairs:

Primary Virtual Disk	Secondary Virtual Disk	Status	Total Repository Capacity	Available Repository Capacity	Repository Status
SRA_Primary	SRA_Backup_1	✓ Optimal	42,792 GB	42,790 GB (100%)	✓ Optimal

Member replication repository members:

Member Name	Name	Status	Capacity
SRA_Primary	repos_0000	✓ Optimal	1,000 GB
SRA_Primary	repos_0001	✓ Optimal	38,792 GB

Figure 11. MDSM Snapshot Volume View

How SRA uses snapshots

The MD Series SRA utilizes Point-in-Time Snapshots if the feature is enabled on the storage array during test failover, the SRA creates a Snapshot Group, Snapshot Image, and Snapshot Volume on the recovery site's storage array for all volumes contained in the protection groups being tested. This requires the creation of the two snapshot repositories listed above and the default size for these repositories is 10 percent of the base volumes for each repository, for a total of 20 percent of the base volume size. This means that the amount of free capacity on the recovery site storage array must be 20 percent of the base volumes participating in the test failover. This value is controlled by the `SraConfigurationData.xml` file located in the config directory under the installation directory, typically:

`C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\scripts\SAN\Dell\SraConfigurationData.xml`

The value is set with the XML tag `<SnapshotBasePercentage>`.

```
<!--
SnapshotBasePercentage represents the initial size, expressed as a percentage
of volume size, of a snapshot which is formed for test failover.
-->
<SnapshotBasePercentage>10</SnapshotBasePercentage>
```

Based on how the VMs are used during test failover, it is possible to fine-tune this value for your environment to reduce the amount of free capacity needed for test failover. If the test VM residing on the recovery site ESX host does not write extensive data to the Datastores and no synchronization (or minimal changes) occur between the protected site volumes and the recovery site volumes during test failover, this value may be decrease to 2–5 to require even less free capacity during test failover. The Snapshot Volume and Snapshot Group repositories are deleted during the cleanup phase of test failover along with the Snapshot Image.

Because the Snapshot Volume is typically not used during the test failover process (little write activity), the size of these repositories may be decreased to small sizes in order to preserve free capacity on the recovery site's storage array. If a repository runs out of space during the test failover phase, the VMs on the recovery site lose access to the Datastore and underlying volume affected by the out-of-space condition of the repository, but the protected site VM functions as normal. These values are not recommended for Snapshot repositories used for other purposes. To monitor the sizes and status of the

repositories, in the left pane, click **All Logical Objects**. In the working pane, from the **Object Type** drop-down menu, select **Repositories**.

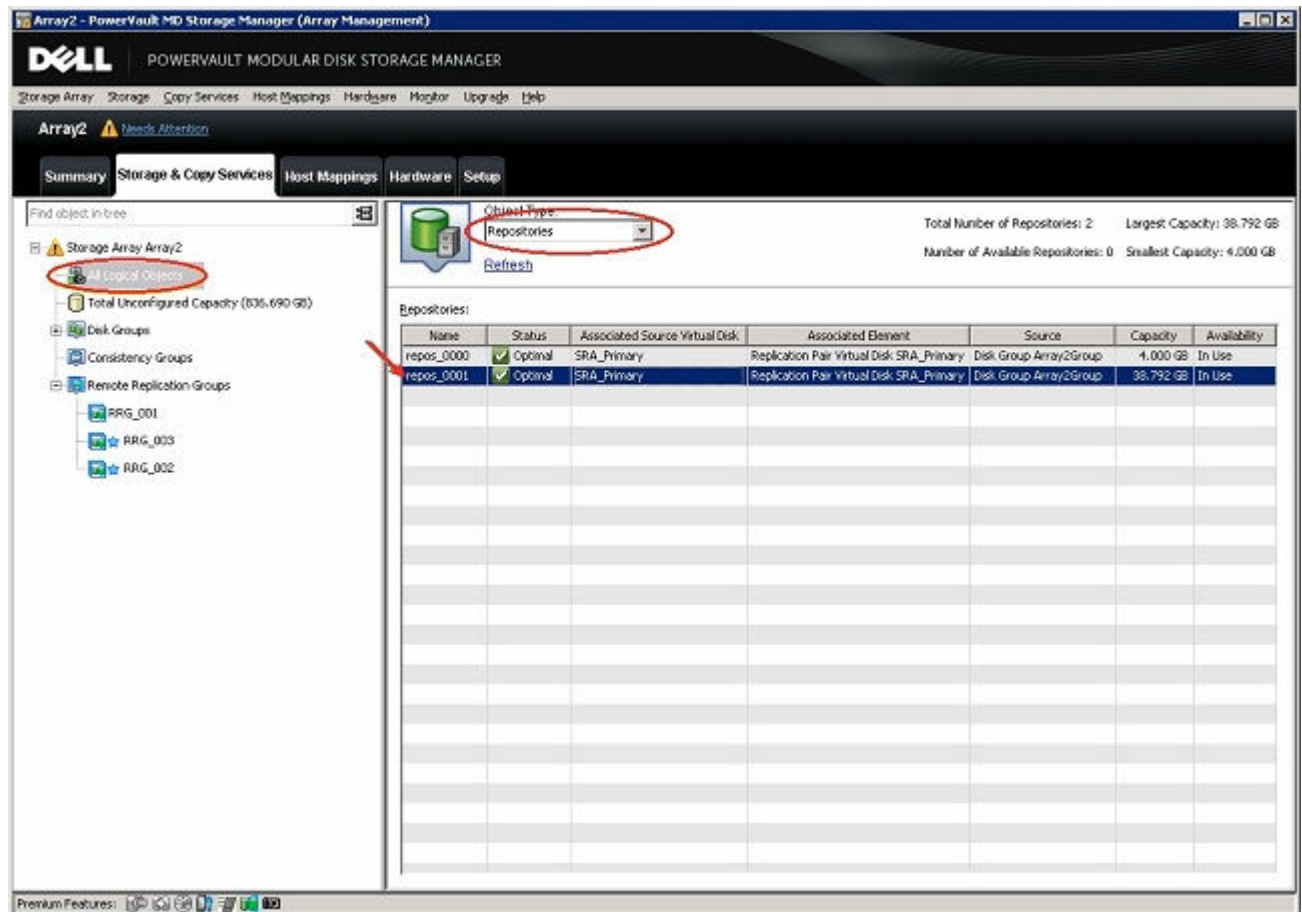


Figure 12. MDSM Snapshot Volume View

To view the repository space, mode, or timestamps, from the drop-down menu, select **Volumes** or **Snapshot Groups**.

Getting help

Documentation matrix

The documentation matrix provides information about documents that you can refer to for setting up and managing your system.

Dell documentation

- For all PowerEdge and PowerVault documentation, go to **Dell.com/support** and enter the system Service Tag to get your system documentation.
- For all Virtualization documents, go to **Dell.com/virtualizationsolutions**.
- For all operating system documents, go to **Dell.com/operatingsystemmanuals**.
- For all storage controllers and PCIe SSD documents, go to **Dell.com/storagecontrollermanuals**.
- For Dell Support Forums, go to **en.community.dell.com/support-forums/default.aspx**.
- For Dell Advanced Search, go to **search.dell.com/index.aspx**.

VMware documentation

- For vCenter SRM 6.0 documentation, go to
https://www.vmware.com/support/pubs/srm_pubs.html
- For vSphere 6.0 Documentation (ESXi, ESX, and vCenter Server), go to
<https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>
- For information about VMware Knowledge Base (Searchable Support Issues), go to
<http://kb.vmware.com/selfservice/microsites/microsite.do>
- For information about VMware Communities (Help Forums), go to
<https://communities.vmware.com/welcome>
- For VMware Compatibility Guide, go to
<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=io>

Contacting Dell



NOTE: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues:

1. Go to **dell.com/support**.
2. Select your support category.
3. Verify your country or region in the **Choose a Country/Region** drop-down list at the bottom of the page.
4. Select the appropriate service or support link based on your need.

Locating your system Service Tag

Your system is identified by a unique Express Service Code and Service Tag number. The Express Service Code and Service Tag are found on the front of a physical DR Series system by pulling out the information tag. This can also be found on the support tab in the GUI. This information is used by Dell to route support calls to the appropriate personnel.