



Confidence in a connected world.

# Enabling DR in SmartTier configurations using Veritas™ Volume Replicator

*August 2010  
Yatin Nayak  
VVR Engineering*

# Enabling DR in SmartTier configurations using Veritas™ Volume Replicator

I.	Introduction .....	3
II.	Configuration requirements.....	3
III.	Setting up storage tiers at the VVR secondary sites .....	4
IV.	Admin Models .....	4
a)	Admin Model for the VVR user – Setting up SmartTier in an existing VVR configuration .....	4
b)	Admin Model for the SmartTier user – Replicate an existing SmartTier configuration using VVR.....	6
V.	Benefits of enabling SmartTier in VVR configurations: .....	8
VI.	Benefits of enabling DR in SmartTier configurations: .....	9
VII.	Glossary of terms.....	10
VIII.	References .....	10

## I. Introduction

Veritas™ Volume Replicator (VVR) is a data-replication software solution designed to contribute to an effective disaster recovery plan. VVR enables you to maintain a consistent copy of application data at one or more remote locations. VVR is a separately licensed option of Veritas™ Volume Manager (VxVM) that works as a fully integrated component of VxVM. A key feature that VVR provides is “write-order fidelity” - this ensures data consistency by applying writes on the VVR secondary host in the same order that they were issued by the application on the VVR primary host.

SmartTier<sup>1</sup> is a feature of the Veritas™ File System (VxFS) which enables the effective exploitation of multi-tier storage. The SmartTier user has the ability to create multiple storage tiers and associate the relative importance of data to a specific storage tier. The user is also provided with control on the placement of the data on each of these tiers, based on specific file placement policies as chosen or created by the user. SmartTier works on a Multi-Volume File System (MVFS), which is a file system that resides on two or more storage volumes.

A Volume Set (VSET) is a VxVM object, which internally comprises of a collection of VxVM volumes. The VSET object enables support for MVFS. Starting with SF 4.1 release, VVR supports replicating a VSET as part of a Replicated Data Set (RDS), to a DR site.

SmartTier is built on top of a MVFS, and since VVR supports replicating a VSET, VVR is able to seamlessly integrate with SmartTier.

Please refer to the **Glossary of terms** section towards the end of this document for a brief description of the product terminology.

## II. Configuration requirements

1. Identical configuration of the VSET is required on all VVR secondary sites

A SmartTier-enabled file system is built on top of a MVFS, which in turn is built upon a VSET. A key requirement for replicating SmartTier-enabled file systems is that the VSET configuration should be identical on all the VVR primary sites. In other words, the names, indices and sizes of all the VSET component volumes as well as the VSET names themselves must remain the same across all the VVR sites.

2. The storage tiers need to be manually setup on all VVR secondary sites.

The volume tags assigned by the SmartTier user, which constitute the storage tiers, are not replicated by VVR to the secondary sites. Hence, the user needs to manually associate the same SmartTier tag to the data volumes on the secondary data volumes as on the primary data volumes. It is not required that the characteristics and cost of storage of these storage tiers on the VVR secondary sites be identical to that of the storage tiers on the VVR primary site. With the volume tags properly assigned on the VVR secondary sites, Storage Foundation can promote one of the VVR secondary sites to a VVR primary site in the event of a disaster. The existence of the same storage tier hierarchy on the new VVR primary site is required for proper enforcement of a previously active SmartTier file placement policy.

---

<sup>1</sup> SmartTier includes and extends functionality previously provided by Dynamic Storage Tiering (DST)

### III. Setting up storage tiers at the VVR secondary sites

Assuming the identical VSETs are already created on all the VVR sites, storage tiers can be setup on the VVR secondary sites by performing the following steps:

1. Query the SmartTier tags for each of the VSET component volumes on the VVR primary site using the vxassist/vxvoladm CLI.

For an example, please refer to step 2 below of the **Admin Models** section.

2. Assign the same SmartTier tag manually to each of the corresponding VSET component volumes on the VVR secondary sites again using the vxassist/vxvoladm CLI, one volume at a time.

For an example, please refer to step 3 below of the **Admin Models** section.

### IV. Admin Models

#### a) Admin Model for the VVR user – Setting up SmartTier in an existing VVR configuration

##### Assumptions:

The following assumptions have been taken into account for the implementation below:

1. The RDS name is 'rvg', replication may or may not have been started yet.
2. A Multi-Volume file system already mounted by the user and referenced by its mount point '/apmnt' but doesn't have SmartTier enabled yet

##### Implementation:

Here are the steps a VVR user needs to do in order to enable SmartTier:

1. On the VVR secondary site(s), create a volume set with an identical configuration as on the primary site i.e. the names, indices and sizes of all the component volumes must remain the same:

```
# /usr/sbin/vxmake -g <dg> vset vs appvols=<volumelist>
```

2. On the VVR primary site, query the volume tags associated to the component data volumes of the volume set on the primary. The following command lists all the volume tags for all the volumes in a specified disk group:

If using SF 5.1 and later releases:

```
# /usr/sbin/vxassist -g <dg> listtag
```

If using SF 5.0 and previous releases:

```
# /usr/sbin/vxvoladm -g <dg> listtag
```

3. On the VVR secondary site(s), associate the same tags to the component data volumes of the volume set as on the primary i.e. for each volume vol in the volume set, do this:

If using SF 5.1 and later releases:

```
# /usr/sbin/vxassist -g <dg> settag <vol> vxfs.placement_class.<tag>
```

If using SF 5.0 and previous releases:

```
# /usr/sbin/vxvoladm -g <dg> settag <vol> vxfs.placement_class.<tag>
```

**Note:** SmartTier recognizes volume tags of the form vxfs.placement\_class.class\_name. The prefix vxfs identifies the tag as being used by the VxFS file system. The placement\_class identifies the tag as a file placement class recognized by SmartTier. The class\_name represents the user-defined name of the file placement class to which the tagged volume belongs. For example, a volume having the tag vxfs.placement\_class.gold is said to belong to placement class gold. Only the gold part of the volume tag appears in file placement policies.

4. Create a SmartTier file placement policy using multiple invocations of fsppmk. Each invocation of fsppmk command generates an XML construct for a portion of a placement policy. The fsppmk command outputs the generated XML to stdout. A series of outputs from this command can be concatenated to generate a complete placement policy.

```
# /opt/VRTS/bin/fsppmk policy -n <policy_name> <options>
# /opt/VRTS/bin/fsppmk rule <options>
# /opt/VRTS/bin/fsppmk select <options>
# /opt/VRTS/bin/fsppmk create <options>
# /opt/VRTS/bin/fsppmk delete <options>
# /opt/VRTS/bin/fsppmk relocate <options>
...
```

Refer to the fsppmk(1M) man page for more details as well as an example on how to use fsppmk.

Alternatively, one of the pre-packaged (with VRTSvxfs package) file placement policies in /opt/VRTSvxfs/etc may be used as-is or after appropriate edits.

5. Validate the syntax and grammar of the SmartTier file placement policy (if not using one of the pre-packaged file placement policies) using the following command:

```
# /opt/VRTS/bin/fsppadm validate /appmnt <path_of_placement_policy_file>
```

6. Assign the placement policy to the VxFS file system

```
# /opt/VRTS/bin/fsppadm assign /appmnt <path_of_placement_policy_file>
```

7. Associate the volume set to the RDS :

```
# /usr/sbin/vradmin -g <dg> addvol <rvg> <vs>
```

8. If replication is not started yet, start replication using automatic synchronization (or one of the various initial synchronization methods) to the secondary site(s). This enables replication of data residing in the storage tiers to the secondary site(s):

```
# /usr/sbin/vradmin -g <dg> -a startrep <rvg> <sechost>
```

9. The user may now enforce a relocation of the existing files on the primary to the corresponding storage tiers, as defined by the active placement policy assigned to that file system, by running:

```
# /opt/VRTS/bin/fsppadm enforce -r /tmp/report /appmnt
```

The relocation of the files will happen on the secondary in the same order as done on the primary.

10. If the user desires to add a new volume to the volume set (at a specific index) as well as the RDS on all VVR sites, this can be done without taking the VxFS file system offline, i.e. with the file system still mounted. The following steps need to be done:

10.1 On the primary site, query the volume tag associated with that volume using the command in step 2 above.

10.2 On the secondary site(s), associate the same volume tag to the volume using the command in step 3 above.

10.3 Run the following vradmin command to add the volume to the volume set and to the RDS on all VVR sites:

```
# /usr/sbin/vradmin -g <dg> -tovset vs addvol <rvg> \  
    <newvol>[:index]
```

Note that vradmin command adds the volume to the volume set and to the RDS on all VVR sites; do not use the VxVM vxvol and vxvset commands to associate the volume to the RDS and volume set separately because these commands work on only one site at a time.

10.4 Finally on the primary site, run the following command to add the volume to the MVFS:

```
# /opt/VRTS/bin/fsvoladm add /appmnt <newvol> <size>
```

11. Optionally, if the user desires to exclude a specific volume belonging to a specific storage tier from replication, perform the following step:

```
# /usr/sbin/vradmin -g <dg> delvol <rvg> <vset_vol>
```

If the above command is run for all volumes of a specific storage tier, it will effectively exclude the entire storage tier from replication.

## b) Admin Model for the SmartTier user – Replicate an existing SmartTier configuration using VVR

### Assumptions:

The following assumptions have been taken into account for the below implementation:

1. User has a volume set named 'vs' on the primary site, on which the user has created a VxFS file system.
2. SmartTier is enabled on the VxFS file system, as per the user's specified file placement policy.

### Implementation:

Here are the steps needed to be performed by SmartTier users in order to enable replication to a remote site :

1. On the VVR secondary site(s), create a volume set with an identical configuration as on the VVR primary site i.e. the names, indices and sizes of all the component volumes must remain the same:

```
# /usr/sbin/vxmake -g <dg> vset vs appvols=<volumelist>
```

2. On the VVR primary site, query the volume tags associated to the component data volumes of the volume set on the primary. The following command lists all the volume tags for all the volumes in a specified disk group:

If using SF 5.1 and later releases:

```
# /usr/sbin/vxassist -g <dg> listtag
```

If using SF 5.0 and previous releases:

```
# /usr/sbin/vxvoladm -g <dg> listtag
```

3. On the VVR secondary site(s), associate the same tags to the component data volumes of the volume set on the secondary as on the primary i.e. for each component volume vol in the volume set, do this:

If using SF 5.1 and later releases:

```
# /usr/sbin/vxassist -g <dg> settag <vol> vxfs.placement_class.<tag>
```

If using SF 5.0 and previous releases:

```
# /usr/sbin/vxvoladm -g <dg> settag <vol> vxfs.placement_class.<tag>
```

**Note:** SmartTier recognizes volume tags of the form vxfs.placement\_class.class\_name. The prefix vxfs identifies the tag as being used by the VxFS file system. The placement\_class identifies the tag as a file placement class recognized by SmartTier. The class\_name represents the user-defined name of the file placement class to which the tagged volume belongs. For example, a volume having the tag vxfs.placement\_class.gold is said to belong to placement class gold. Only the gold part of the volume tag appears in file placement policies.

4. Create an RDS using the following steps:

#### 4.1 Create the SRL volume on the primary and secondary sites

```
# /usr/sbin/vxassist -g <dg> make srl <size>
```

#### 4.2 Create a primary RVG consisting of the volume set and the SRL

```
# /usr/sbin/vradmin -g <dg> createpri <rvg> vs srl
```

#### 4.3 Add the VVR secondary site to the RDS.

```
# /usr/sbin/vradmin -g <dg> addsec <rvg> <prihost> <sechost>
```

After successful completion of this command, the RDS now consists of a VVR primary and a VVR secondary.

5. Perform the following step to associate the volume set to the RDS:

```
# /usr/sbin/vradmin -g <dg> addvol <rvg> <vs>
```

6. If replication is not started yet, start replication using automatic synchronization (or one of the various initial synchronization methods) to the secondary site(s). This enables replication of data residing in the storage tiers to the secondary site(s):

```
# /usr/sbin/vradmin -g <dg> -a startrep <rvg> <sechost>
```

7. If the user desires to add a new volume to the volume set (at a specific index) as well as the RDS on all VVR sites, the following steps need to be done:

- 7.1 On the primary site, query the volume tag associated with that volume using the command in step 2 above.

- 7.2 On the secondary site(s), associate the same volume tag to the volume using the command in step 3 above.

- 7.3 Run the following vradmin command to add the volume to the volume set and to the RDS on all VVR sites:

```
# /usr/sbin/vradmin -g <dg> -tovset vs addvol <rvg> \  
    <newvol>[:index]
```

Note that vradmin command adds the volume to the volume set and to the RDS on all VVR sites; do not use the VxVM vxvol and vxvset commands to associate the volume to the RDS and volume set separately because these commands work on only one site at a time.

- 7.4 On the VVR primary site, run the following command to add the volume to the MVFS:

```
# /opt/VRTS/bin/fsvoladm add </mntpt> <newvol> <size>
```

8. Optionally, if the user desires to exclude a specific volume belonging to a specific storage tier from replication, perform the following step:

```
# /usr/sbin/vradmin delvol -g <dg> <rvg> <vset_vol>
```

If the above command is run for all volumes of a specific storage tier, it will effectively exclude the entire storage tier from replication.

## V. Benefits of enabling SmartTier in VVR configurations:

### a) Network bandwidth savings provided by the ability to do selective replication of the volume set

Though the VSET configuration should be identical on all VVR sites, VVR provides a facility to exclude unimportant VSET component volumes from replication on the primary and secondary sites (for example: a volume containing Oracle temporary logs which has scratch data). This can reduce the network bandwidth costs to a minimum.

### b) Data Availability even in case of VSET component volume failures

A MVFS has the unique feature of providing data availability in most conditions of VSET component volume failures. VVR honors this and does special handling to not disable the replicated volume group (RVG) for any component hardware data volume failures in the volume set, unless the metadata volume

(usually at index 0) itself goes bad. If the metadata volume suffers a hardware disk failure or goes bad for some reason, replication is stopped and manual intervention is needed by the storage administrator to recover the file system data from backup.

## **VI. Benefits of enabling DR in SmartTier configurations:**

### **a) Disaster recovery (DR), enabled by VVR**

The disaster recovery (DR) needs of SmartTier customers are fully met by VVR. In case of a disaster at the primary VVR site, the SmartTier configuration at one of the secondary VVR sites can be enabled.

### **b) Storage savings**

SmartTier brings with itself a tremendous value proposition by reducing the online storage hardware costs. Less important data can reside on lower-cost storage tiers. This can result in good cost savings, while still providing the DR benefits provided by VVR.

### **c) Ability to replicate the data movement enforced by SmartTier file placement policies to a remote VVR site(s)**

This is an automatic benefit of combining SmartTier with VVR. The data movement enforced by the user-defined SmartTier file placement policies on the primary site is automatically replicated to the secondary sites by VVR. Since VVR supports replication across heterogeneous UNIX/Linux platforms, this provides an additional benefit for seamlessly enabling the SmartTier file data movement across a variety of UNIX/Linux platforms and flavors.

## **Notes:**

Starting with SF 5.1 and later releases, fsppadm is part of the VRTSvxfs package and resides at /opt/VRTS/bin/fsppadm

In previous releases (SF 5.0MP3 and earlier), fsppadm was part of the VRTSfspro package, residing at /opt/VRTSfspro/bin/fsppadm

## VII. Glossary of terms

<b>VxFS</b>	Veritas™ File System
<b>VxVM</b>	Veritas™ Volume Manager
<b>VVR</b>	Veritas™ Volume Replicator
<b>DR</b>	Disaster Recovery
<b>SmartTier</b>	Also known as <b>Dynamic Storage Tiering (DST)</b> , a feature of VxFS which enables effective exploitation of multi-tiered storage architecture
<b>VSET</b>	Volume Set, a VxVM container object containing a collection of data volumes
<b>MDFS/MVFS</b>	Multi-device File System / Multi-volume File System, file systems that occupy two or more virtual volumes
<b>VVR Primary</b>	The node where the application is running, and from which data is being replicated to the VVR Secondary.
<b>VVR Secondary</b>	The node to which VVR is replicating data from the VVR Primary.
<b>RVG</b>	Replicated Volume Group, a component of VVR that is made up of a group of data volumes within a given VxVM disk group along with a few more VVR objects. VVR replicates from the Primary RVG, on the node where the application is running, to one or more Secondary RVGs.
<b>SRL</b>	Storage Replicator Log, a circular buffer of writes for an RVG. Writes stored in the SRL are waiting to be shipped to the Secondary from the Primary, or waiting to be written to the data volumes in the RVG.
<b>RDS</b>	Replicated Data Set, a logical grouping of the RVG on a VVR Primary and the corresponding RVGs on one or more VVR Secondary hosts

## VIII. References

1. Symantec Yellow Book: [Using Dynamic Storage Tiering](#)
2. Veritas™ Volume Replicator Administrator's Guide
3. Veritas™ File System Administrator's Guide
4. [Design doc for VVR-VSET Integration](#)
5. Man pages for reference:

vradmin(1M)

vxvset(1M)

vxassist(1M)

vxvoladm(1M) (only with SF 5.0 and previous releases, EOL'ed in SF 5.1)

fsppmk(1M)

fsppadm (1M)

fsvoladm(1M)

## About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at [www.symantec.com](http://www.symantec.com).

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation  
World Headquarters  
20330 Stevens Creek Boulevard  
Cupertino, CA 95014 USA  
+1 (408) 517 8000  
1 (800) 721 3934  
[www.symantec.com](http://www.symantec.com)

Copyright © 2010 Symantec Corporation. All rights reserved.  
Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

9/2010