

Veritas Storage Foundation™ and High Availability Solutions Quick Recovery Solutions Guide for Microsoft Exchange 2010

Windows Server 2008

5.1 Service Pack 1 Application Pack 1



Veritas Storage Foundation and HA Solutions Quick Recovery Solutions Guide for Microsoft Exchange 2010

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.1. Service Pack 1 Application Pack 1

Document version: 5.1.SP1AP1.0

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas and Veritas Storage Foundation 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice file accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrades protection
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our web site at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information

- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively. Enterprise services that are available include the following:

- Managed Services These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
- Consulting Services Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
- Education Services Education Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about enterprise services, please visit our web site at the following URL:

www.symantec.com/business/services/

Select your country or language from the site index.

Contents

Chapter 1	Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange 2010	
	About the solutions guides	12
	Quick Recovery solutions	12
Chapter 2	Using the Solutions Configuration Center	
	About the Solutions Configuration Center	13
	Starting the Configuration Center	14
	Available options from the Configuration Center	15
	About running the Configuration Center wizards	21
	Following the workflow in the Configuration Center	22
	Solutions wizard logs	25
Chapter 3	Quick Recovery concepts	
	About Quick Recovery snapshot solutions	27
	Advantages of Quick Recovery snapshots	28
	Quick Recovery process	29
	About the components used in Quick Recovery	29
	FlashSnap and FastResync	29
	Integration with Microsoft Volume Shadow Copy Service	30
	The VSS framework	30
Chapter 4	Planning a Quick Recovery snapshot solution	
	System requirements	35
	Supported software	36
	Storage requirements	36
	Methods of implementing Quick Recovery snapshots	38
	About the Quick Recovery Configuration Wizard	38
	About the VSS Exchange Snapshot Scheduler Wizard	39
	About the VSS Snapshot and Snapback wizards and the vxsnap utility	40
	Planning your Quick Recovery solution	41
	Backup types for snapshot sets	42
	About logs	42

Recommendations and best practices	42
Recommendations for Exchange storage configuration	43
Recommendations for maintenance and backups	44
VCS cluster considerations	44
Exchange 2010 Database Availability Group (DAG) considerations ...	47
Using snapshot schedules in a DAG	47
Guidelines for creating passive copy snapshots in a DAG	48
Recovery in a DAG	48
VVR considerations	48

Chapter 5 Configuring Exchange for Quick Recovery snapshots

Tasks for configuring Exchange storage	52
Reviewing the prerequisites	52
Reviewing the configuration	53
Creating dynamic disk groups	54
Creating dynamic volumes	55
Pointing the database and log paths to the SFW volumes	59

Chapter 6 Implementing snapshot sets with the configuration wizard

About the Quick Recovery Configuration Wizard	61
About snapshot sets	62
About snapshot templates	63
Templates and multiple components	64
Templates and schedule start dates	64
Tasks for implementing snapshot sets with the configuration wizard	65
Reviewing the prerequisites	65
Scheduling and creating snapshot sets	66
Viewing the status of scheduled snapshots	75
Troubleshooting scheduled snapshots	76
Deleting or modifying schedules	78
Synchronizing schedules after adding a cluster node	79

Chapter 7 Scheduling or creating an individual snapshot set

About scheduling or creating an individual snapshot set	81
Tasks to schedule a new snapshot	82
Tasks to create a one-time snapshot set	83
Reviewing the prerequisites	83
Preparing the snapshot mirrors	84
Scheduling a snapshot set	85
Viewing the status of scheduled snapshots	93

- Troubleshooting scheduled snapshots 94
 - Deleting schedules 96
- Creating a one-time snapshot set 97
- Refreshing a snapshot set manually 101
 - Reattaching the split-mirror snapshots 101
 - Creating the refreshed snapshot set 103

Chapter 8 Recovering Exchange mailbox databases

- About recovery using Quick Recovery snapshots 105
- Tasks for recovery using Quick Recovery snapshots 106
- Prerequisites for recovery 107
- Recovery using an Exchange 2010 passive copy snapshot in a Database Availability Group (DAG) 107
- Recovering using the VSS Restore Wizard 108
 - Recovering using the vxsnap utility 112
- Post-recovery steps 116

Chapter 9 Recovering after hardware failure

- About recovery after hardware failure 117
 - Recovery scenarios 118
 - Recovery tasks 118
- Tasks for recovering after hardware failure 119
- Reviewing the prerequisites 120
- Reviewing the sample configuration 120
- Scenario I: Database and transaction logs volumes are missing 121
 - Identifying the missing volumes 122
 - Deleting missing volumes from Storage Foundation for Windows ... 123
 - Replacing hardware and adding disks to the dynamic disk group 123
 - Changing the drive letter or mount points of the snapshot volumes 124
 - Restoring the mailbox database to the point in time 126
 - Refreshing the snapshot set 127
- Scenario II: Database volumes missing, transaction logs are available 128
 - Identifying the missing volumes 129
 - Deleting missing volumes from Storage Foundation for Windows ... 129
 - Replacing hardware and adding disks to the dynamic disk group 130
 - Changing the drive letter or mount points of the snapshot volumes 131
 - Restoring the mailbox database to the Point in Time (PIT) 133
 - Recovering the mailbox database to the Point of Failure (POF) 135
 - Refreshing the snapshot set 136
- Refreshing the snapshot set 137

- Refreshing the snapshot set on the current disks 137
 - Reattaching healthy snapshot volumes 138
 - Clearing the snapshot association for volumes whose
 - drive letters or mount points were reassigned 139
 - Creating snapshot mirrors of volumes whose
 - drive letters or mount points were reassigned 139
 - Creating the new snapshot set 140
- Moving the production volumes to different disks
 - and refreshing the snapshot set 144
 - Reattaching healthy snapshot volumes 145
 - Clearing the snapshot association for volumes whose
 - drive letters or mount points were reassigned 146
 - Adding mirrors to volumes whose drive letters or
 - mount points were reassigned 146
 - Creating snapshot mirrors of volumes whose drive letters
 - or mount points were reassigned 147
 - Creating the new snapshot set 148

Chapter 10 Vxsnap utility command line reference

- About the vxsnap utility 153
- Vxsnap keywords 154
 - vxsnap prepare 156
 - Syntax 156
 - Attributes 156
 - Example 157
 - vxsnap create 157
 - Syntax 158
 - Attributes 158
 - Examples 161
 - vxsnap reattach 161
 - Syntax 161
 - Attributes 162
 - Example 162
 - vxsnap restore 162
 - Syntax 163
 - Attributes 164
 - Examples 166
 - Recovery after hardware failure 166

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange 2010

This chapter includes the following topics

- [About the solutions guides](#)
- [Quick Recovery solutions](#)

About the solutions guides

[Table 1-1](#) shows the solutions guides provided with Veritas Storage Foundation HA for Windows for Microsoft Exchange Server 2010. Separate guides are available for earlier versions of Exchange, for Microsoft SQL, for Enterprise Vault, for Microsoft SharePoint Server, and for additional application solutions.

Table 1-1 SFW HA solutions guides for Exchange Server 2010

Title	Description
<i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange 2010</i>	Solutions for Exchange Server 2010 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows <ul style="list-style-type: none"> ■ High availability (HA) ■ Campus clusters ■ Replicated data clusters ■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication
<i>Veritas Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft Exchange 2010</i>	Quick Recovery solutions for Exchange Server 2010 using either Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.

Quick Recovery solutions

Quick Recovery is the process of creating and maintaining on-host point-in-time images of production Exchange databases that can be used to quickly recover missing or corrupt Exchange databases.

These images, called snapshot sets, contain snapshots of the volumes associated with the Exchange databases and transaction logs.

Quick Recovery is designed to augment your traditional backup methodology.

Using the Solutions Configuration Center

This chapter covers the following topics:

- [About the Solutions Configuration Center](#)
- [Starting the Configuration Center](#)
- [Available options from the Configuration Center](#)
- [About running the Configuration Center wizards](#)
- [Following the workflow in the Configuration Center](#)
- [Solutions wizard logs](#)

About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your Veritas Storage Foundation for Windows (SFW) or SFW High Availability (HA) environment. The Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2003, 2007, and 2010
- Microsoft SQL Server 2000, 2005, and 2008
- Enterprise Vault Server (high availability new server and disaster recovery solutions)
- Additional applications

You can use the Configuration Center and its wizards to set up your environment for any combination of the following solutions:

- High availability at a single site for a new installation
- High availability at a single site for an existing server

- Campus cluster disaster recovery, including the following:
 - Campus cluster using Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003, 2007, 2010, and for Microsoft SQL Server 2005 and 2008)
- Fire drill to test the fault readiness of a disaster recovery environment

The Solutions Configuration Center provides two ways to access Solutions wizards:

- The Applications tab lists solutions by application. It provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.
- The Solutions tab, for advanced users, lists wizards by solution without additional instructions.

Starting the Configuration Center

You can start the Configuration Center in the following ways:

- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the Solution Guides listed in the right pane match the selected application.

In addition, some choices can vary depending on the operating system of the node on which you launch the wizard. For example, since Microsoft Exchange 2003 runs only on 32-bit operating systems, on a 64-bit system only the Exchange 2007 and Exchange 2010 configuration wizards are shown.

Figure 2-1 shows the choices available on a 64-bit system when you click Solutions for Microsoft Exchange.

Figure 2-1 Solutions Configuration Center for Microsoft Exchange

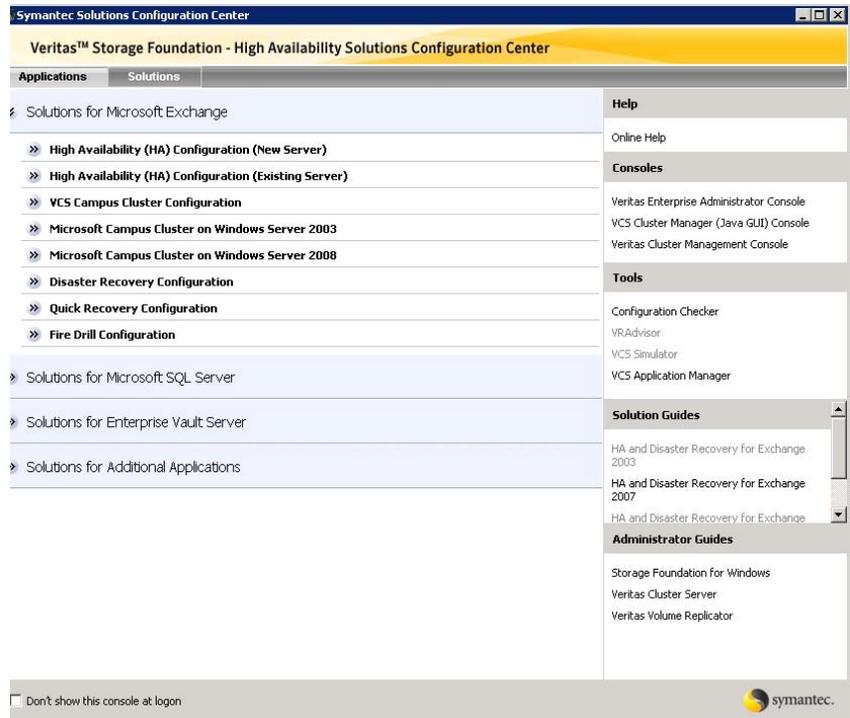


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

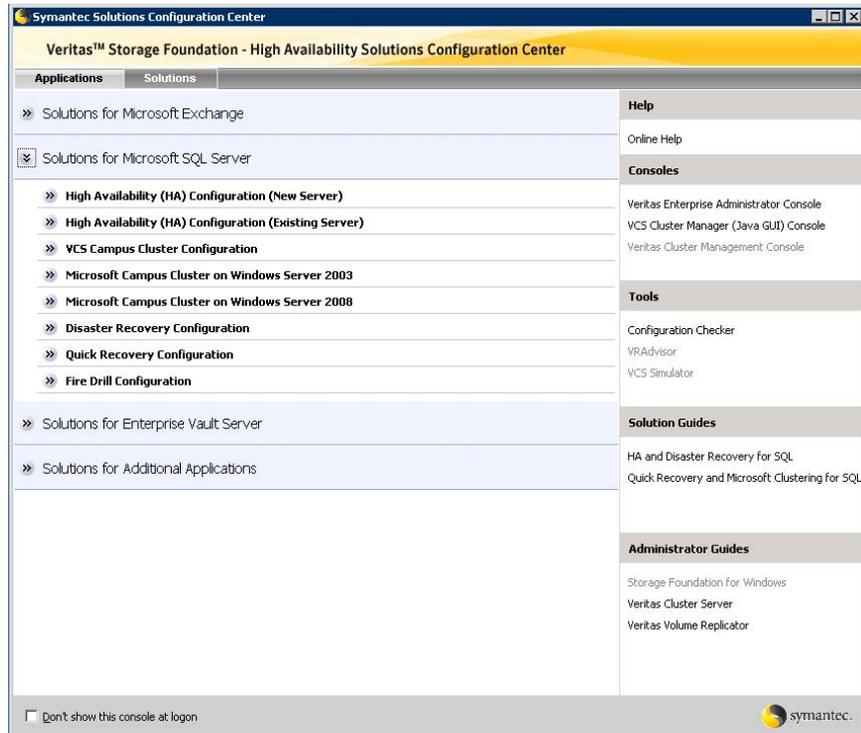


Figure 2-3 shows the choices available when you click Solutions for Enterprise Vault Server.

Figure 2-3 Solutions Configuration Center for Enterprise Vault Server

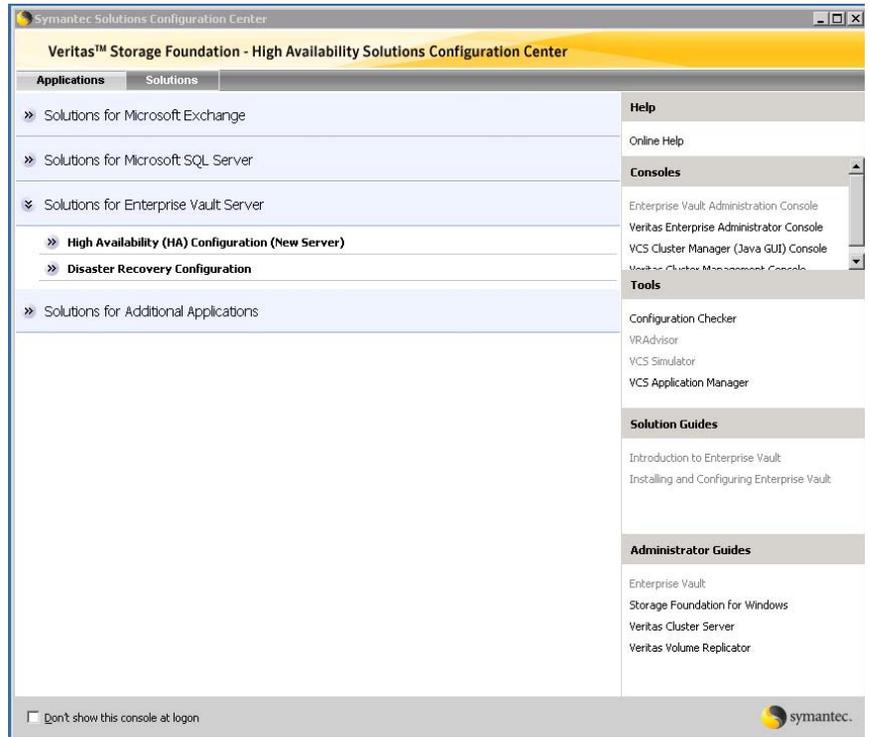
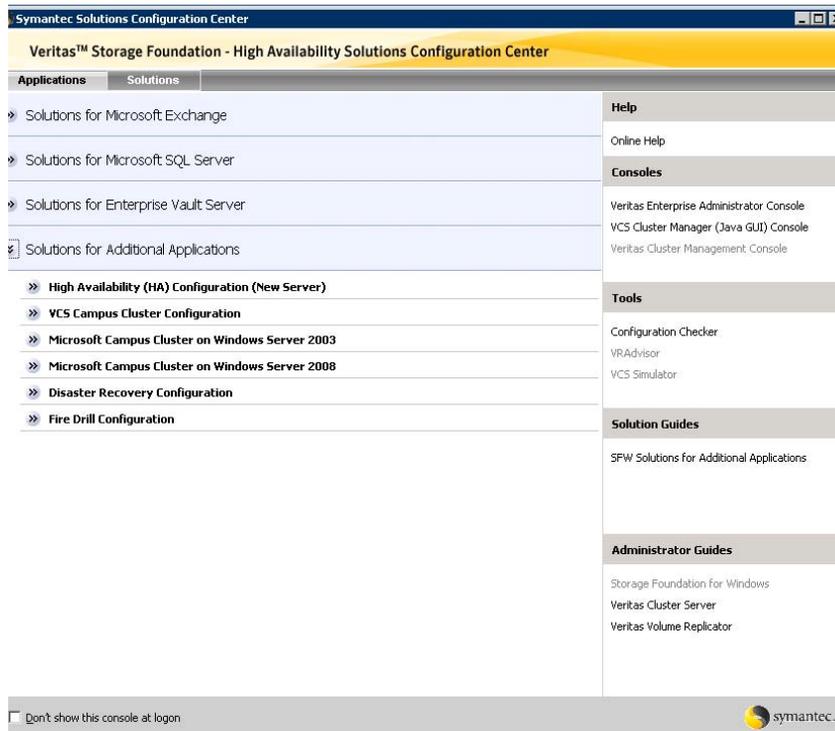


Figure 2-4 shows the choices available when you click Solutions for Additional Applications.

Figure 2-4 Solutions Configuration Center for additional applications



The submenu choices also vary by application. For example, different steps, information, or wizards are shown under High Availability (HA) Configuration for Exchange than those shown for SQL Server.

Figure 2-5 shows one of the steps for implementing high availability for Exchange.

Figure 2-5 Context-sensitive step for Exchange



Figure 2-6 shows one of the steps for implementing high availability for SQL Server.

Figure 2-6 Context-sensitive step for SQL Server



Figure 2-7 shows one of the steps for implementing high availability for Enterprise Vault Server.

Figure 2-7 Context-sensitive step for Enterprise Vault Server

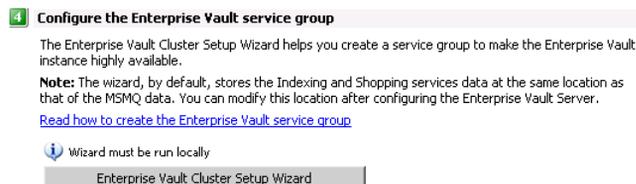


Figure 2-8 shows one of the steps for implementing high availability for additional applications.

Figure 2-8 Context-sensitive step for additional applications

4 Configure the service group

Create a service group to make your application or server role highly available.

[Read how to create a File Share service group](#)

Wizard must be run locally

File Share Configuration Wizard

[Read how to create a Print Share service group](#)

Wizard must be run locally

Print Share Configuration Wizard

[Read how to create an IIS Server service group](#)

Wizard must be run locally

IIS Configuration Wizard

[Read how to create a Microsoft Virtual Server virtual machine service group](#)

Wizard must be run locally

MSVirtual Machine Configuration Wizard

[Read how to create an Oracle service group](#)

Wizard must be run locally

Oracle Agent Configuration Wizard

[Read how to create a service group for an application, process or service](#)

Wizard must be run locally

Application Configuration Wizard

About running the Configuration Center wizards

You can run the wizards from the Applications tab if you are walking through the configuration steps on the Solutions Configuration Center. If you are already familiar with configuration, you can also go directly to a particular wizard by selecting the Solutions tab.

The Configuration Center and some wizards can be run from a remote system. Wizards that you can run remotely include the following:

VCS Configuration Wizard	Sets up the VCS cluster
Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster Also can configure Veritas Volume Replicator (VVR) replication or configure the VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication Requires first configuring high availability on the primary site
Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots
Fire Drill Wizard	Sets up a fire drill to test disaster recovery Requires configuring disaster recovery first

Wizards related to storage configuration and application installation must be run locally on the system where the process is occurring. Wizards that you must run locally include the following:

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Microsoft Exchange Server 2003 and 2007 for the high availability environment If Microsoft Exchange is already installed, refer to the documentation for further instructions. This wizard is not required for Exchange Server 2010.
Exchange Configuration Wizard	Configures the service group for Microsoft Exchange Server 2003 and 2007 high availability

Exchange 2010 Configuration Wizard	Configures the service group for Microsoft Exchange Server 2010 high availability
SQL Server Configuration Wizard	Configures the service group for SQL Server 2000 or SQL Server 2005 high availability You must first install SQL Server on each node according to the instructions in the documentation.
SQL Server 2008 Configuration Wizard	Configures the service group for SQL Server 2008 high availability You must first install SQL Server on each node according to the instructions in the documentation.
Enterprise Vault Cluster Setup Wizard	Configures the service group for Enterprise Vault Server high availability
MSDTC Wizard	Configures an MSDTC Server service group for SQL Server 2000, 2005, or 2008 environments
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group

The Additional Applications section of the Configuration Center provides wizards to be run locally for creating service groups for the following applications or server roles:

File Share Configuration Wizard	Configures FileShare for high availability
Print Share Configuration Wizard	Configures PrintShare for high availability
IIS Configuration Wizard	Configures IIS for high availability
MSVirtual Machine Configuration Wizard	Configures MS Virtual Machine for high availability
Oracle Agent Configuration Wizard	Configures Oracle for high availability
Application Configuration Wizard	Configures any other application service group for which application-specific wizards have not been provided

Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration

Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

Figure 2-9 shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

Figure 2-9 Workflow for configuring Exchange high availability



Figure 2-10 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

Figure 2-10 Workflow for configuring SQL Server high availability

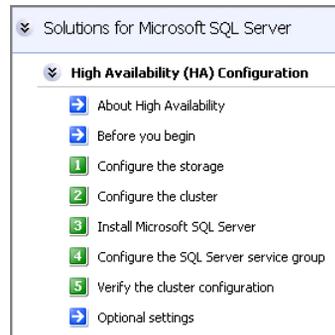


Figure 2-11 shows the high-level overview of the workflow steps for configuring high availability for Enterprise Vault Server from the Solutions Configuration Center.

Figure 2-11 Workflow for configuring high availability for Enterprise Vault Server

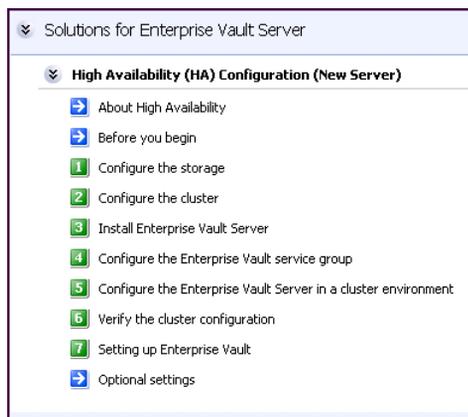
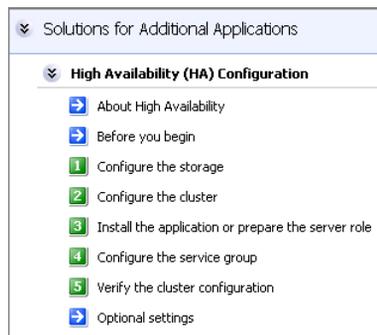


Figure 2-12 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

Figure 2-12 Workflow for configuring high availability for additional applications



Solutions wizard logs

The Solutions Configuration Center provides access to many wizards. However, three wizards are built in to the Solutions Configuration Center:

- Disaster Recovery Wizard
- Fire Drill Wizard
- Quick Recovery Configuration Wizard

These three Solutions wizards are launched only from the Solutions Configuration Center, whereas other wizards can be launched from product consoles or the Start menu.

Logs created by these three Solutions wizards are located in the following paths:

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application  
Data\VERITAS\winsolutions\log
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```


Quick Recovery concepts

This chapter includes the following topics:

- [About Quick Recovery snapshot solutions](#)
- [Advantages of Quick Recovery snapshots](#)
- [Quick Recovery process](#)
- [About the components used in Quick Recovery](#)

About Quick Recovery snapshot solutions

Veritas Storage Foundation for Windows (SFW) provides the capability to create point-in-time images of production Exchange databases. These images, called snapshot sets, contain snapshots of the volumes associated with the databases and transaction logs. A snapshot set is a complete copy of the database at the specific point in time the snapshots are taken.

SFW uses Veritas FlashSnap™ technology along with the Microsoft Volume Shadow Copy Service (VSS) framework to quiesce an Exchange database and ensure a persistent snapshot of the production data.

Quick Recovery is the term for creating and maintaining the snapshot sets on-host for use in quickly recovering missing or corrupt Exchange databases. The Quick Recovery solution provides fast recovery from logical errors and eliminates the time-consuming process of restoring data from tape. Databases can be recovered to the point in time when the snapshot was taken or, by using current logs, rolled forward to the point of failure.

Using the SFW Quick Recovery Configuration wizard, you can create multiple snapshot sets for each database and set up schedules for creating and refreshing the snapshot sets. The snapshot sets can be maintained on-host as a Quick Recovery solution.

SFW snapshot solutions use a split-mirror snapshot method. The snapshot is a separate persistent volume that contains an exact duplicate of all the data on

the original volume at the time the snapshot is taken. This type of persistent physical snapshot is also known as a Clone (HP) or a BCV (EMC). In contrast, copy-on-write snapshots, also known as metadata snapshots, only copy changed blocks to the snapshot and do not create a separate physical volume.

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the database, the snapshot set requires the same amount of space as the original volumes.

Veritas FlashSnap technology is also integrated into the Veritas NetBackup 6.0 Advanced Client Option and Symantec Backup Exec 10d Advanced Disk-based Backup Option. These products are the preferred solution for on and off host snapshot-assisted backup.

Advantages of Quick Recovery snapshots

A Quick Recovery solution serves as a first line of defense to recover corrupted or missing Exchange databases. Maintaining a snapshot set requires just the few seconds it takes to detach a split-mirror snapshot from its original volume. On-host snapshot recovery is faster than restoring a full backup from tape or other media; on-host snapshot recovery reduces downtime and helps meet service-level agreements for application availability.

In addition to the primary benefit of recovery from logical errors, snapshot sets can be moved over a SAN to another server and used for other purposes including:

- Application tuning and testing—data can be updated and modified in a realistic environment without impacting users.
- Business reporting and decision analysis—up-to-date data is available with minimal impact on the production environment.

Quick Recovery process

The Quick Recovery process can be broken down into the following phases:

- Creating an initial snapshot set
This has two stages:
 - Preparing the mirror for the snapshot set
This stage takes a while and should be scheduled for a time of low activity.
 - Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again, as needed or according to a pre-set schedule
This stage is automated by setting up snapshot schedules using the Quick Recovery wizard or VSS Exchange Snapshot Scheduler wizard.
- Using a snapshot set to recover a corrupted database

About the components used in Quick Recovery

SFW snapshot solutions use Veritas FlashSnap and FastResync technology along with the Microsoft Volume Shadow Copy Service framework.

FlashSnap and FastResync

Veritas FlashSnap provides the ability to create and maintain the on-host point-in-time copies of volumes that are integral to the snapshot solutions. Both the original and snapshot volume may consist of multiple physical devices, as in the case of RAID 0+1 (Mirrored Striped) volumes. FlashSnap cannot be used with software RAID-5 volumes.

FastResync is a FlashSnap feature that optimizes the resynchronization of a snapshot volume and its original volume. FlashSnap uses FastResync technology to track the changed blocks in an original volume after a snapshot is detached. A Disk Change Object (DCO) volume is automatically created to store a record of these changes. When the snapshot volume is resynchronized with the original volume, only the changed data blocks are written to the snapshot volume. This greatly reduces the time and performance impact of resynchronization which means that a snapshot set can be refreshed with minimal impact to production.

Integration with Microsoft Volume Shadow Copy Service

SFW integrates with the Microsoft Volume Shadow Copy Service (VSS) as both a VSS Requestor and a VSS Provider. This integration is provided by FlashSnap.

The Volume Shadow Copy Service (VSS) process allows the databases to be frozen before the snapshot operation occurs and then thawed immediately after it. This quiescing, allows for Microsoft supported and guaranteed persistent snapshots of your data.

FlashSnap integrates with VSS to create a snapshot set containing snapshot volumes of all the volumes associated with an Exchange database without taking the databases offline.

The VSS framework

There are four components to the VSS framework: Requestor, Writer, Provider, and the Volume Shadow Copy Service itself.

Table 3-1 VSS framework components

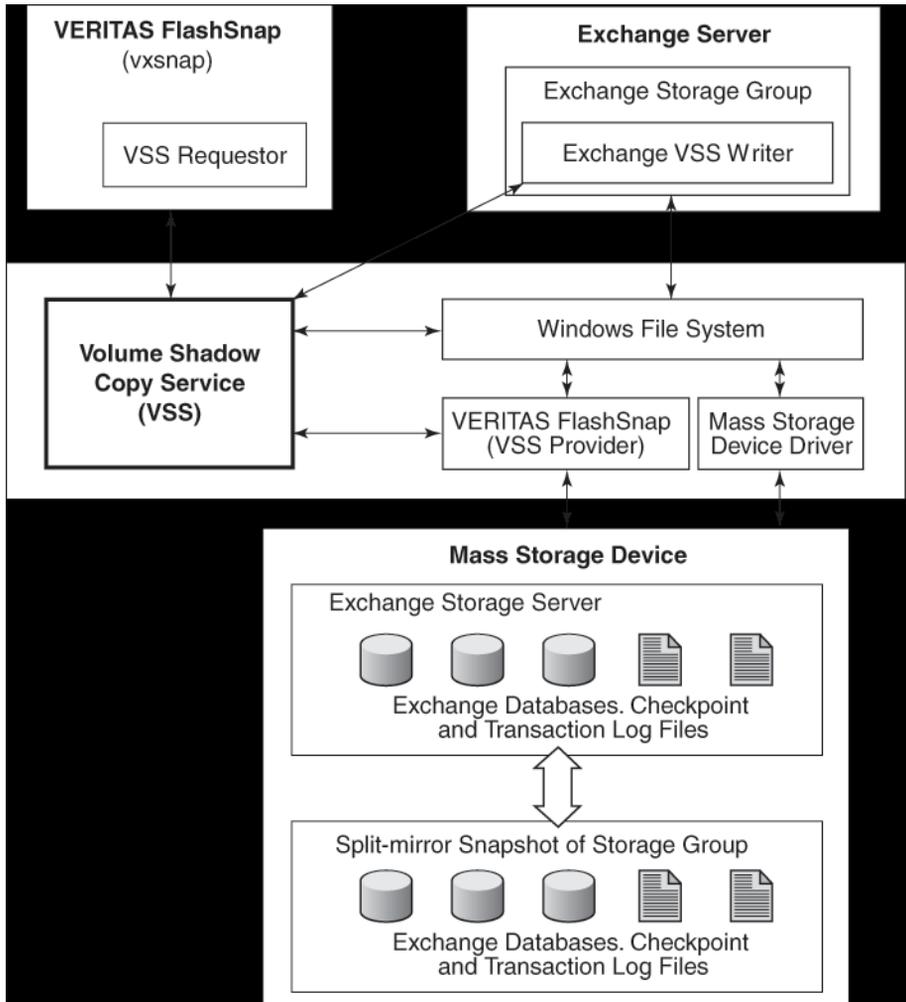
Component	Action
Volume Shadow Copy Service	Talks to and coordinates the Requestor, Provider, and Writer.
Requestor	As a Requestor, the vxsnap component of FlashSnap notifies the VSS coordinator to initiate the VSS request to prepare Exchange for quiescing and later requests that the snapshot process begin.

Table 3-1 VSS framework components

Component	Action
Writer	<p>As Writers, VSS-enabled applications such as Exchange respond to requests to prepare and participate in the generation of snapshots, provide success/failure status, and provide information about the application including what is to be backed up and restored, and restore strategy.</p> <p>The Exchange 2010 VSS Writer has two instances, the Store Writer and the Replication Writer.</p> <p>The Store Writer is available on any mailbox server and is responsible for backup/restore of active databases.</p> <p>The Replication Writer is available on the passive node of a DAG server. The Replication Writer supports backup functionality for a selected database where the snapshot is taken against the replicated instance of the database and transaction log files. The Replication Writer is supported by SFW Flashsnap for snapshot and schedule snapshot operations but not for restore operations.</p>
Provider	As a Provider, FlashSnap creates the persistent snapshot.

[Figure 3-1](#) shows the steps in the Volume Shadow Copy Service Process.

Figure 3-1 The Volume Shadow Copy Service process



FlashSnap integrates with Volume Shadow Copy Service as both a VSS Requestor and a VSS Provider.

The following steps occur, in the order presented, during the snapshot process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a split-mirror snapshot of an Exchange database.
- The VSS coordinator service calls the Exchange VSS Writer to find out which volumes contain the databases and transaction logs.

- The VSS coordinator service notifies the FlashSnap VSS Provider to prepare for the snapshot.
- Acting as a VSS Requestor, FlashSnap requests that the VSS coordinator service begin the snapshot call.
- The VSS coordinator service notifies the Exchange VSS Writer to quiesce and freeze the databases in preparation for the snapshot. When this is accomplished, the Exchange Writer informs the VSS coordinator service to proceed.
- The VSS coordinator service calls the FlashSnap Provider to create the split-mirror snapshot by detaching the snapshot volume from the original volume. The snapshot process takes a maximum of 10 seconds. After the snapshot volume is detached, the FlashSnap Provider informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the Exchange Writer to thaw (release the freeze) and resume normal I/O.

The following steps occur, in the order presented, during the restore process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a restore operation.
- The VSS coordinator service calls the Exchange VSS Writer, which prepares for the restore operation.
- The FlashSnap utility restores the snapshot volumes. After the snapback operation completes, the FlashSnap utility informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the Exchange Writer to process the post-restore operations.

Planning a Quick Recovery snapshot solution

This chapter includes the following topics:

- [System requirements](#)
- [Methods of implementing Quick Recovery snapshots](#)
- [Planning your Quick Recovery solution](#)
- [Backup types for snapshot sets](#)
- [About logs](#)
- [Recommendations and best practices](#)

System requirements

A Veritas Storage Foundation for Windows (SFW) Quick Recovery solution can be implemented on either a standalone system or a VCS clustered system. A Quick Recovery solution can also be implemented in an Exchange 2010 Database Availability Group (DAG).

Supported software

An SFW snapshot solution for Exchange 2010 requires that the following software be installed and configured:

- Veritas Storage Foundation 5.1 Service Pack 1 Application Pack 1 for Windows with the FlashSnap option or Veritas Storage Foundation HA 5.1 Service Pack 1 Application Pack 1 for Windows with the FlashSnap option. Refer to the *Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Application Pack 1 Release Notes* for information on upgrading to Application Pack 1.

Note: If using Quick Recovery with Exchange in a VCS cluster environment, Exchange must be installed on all cluster nodes.

- Microsoft Exchange servers and their operating systems (refer to the Microsoft documentation for any additional Exchange requirements):

Microsoft Exchange Server 2010
(Standard or Enterprise Edition)

- Microsoft Windows 2008 x64, Service Pack 2 Standard or Enterprise edition
- Microsoft Windows 2008 x64 R2 Standard or Enterprise edition

Storage requirements

The hardware for your SFW snapshot solution should include sufficient storage to be able to create dynamic volumes on separate disks for the following purposes:

- Databases
- Transaction logs
- Split-mirror snapshots of the database stores and transaction logs volumes

The system and boot volumes should reside on a separate disk (Harddisk0).

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the database, the snapshot set requires the same amount of space as the original volumes.

You can specify one or more snapshot sets for each volume with different disk locations for each. You can create more than one snapshot volume on the same disk as long as there is sufficient space available and as long as the snapshots are of different production volumes.

Disk groups must be of a Storage Foundation for Windows 4.0 or later version. It is necessary to upgrade any disk groups created using an earlier version of

Volume Manager for Windows before implementing SFW snapshot solutions. Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group. For more information, see *Veritas Storage Foundation Administrator's Guide*

Methods of implementing Quick Recovery snapshots

Veritas Storage Foundation for Windows provides more than one method of implementing Quick Recovery snapshots.

[Table 4-1](#) summarizes the methods and when you would use them.

Table 4-1 Methods of implementing Quick Recovery snapshots

Task you want to accomplish	Method	For more information
<ul style="list-style-type: none"> ■ Set up and schedule multiple snapshot sets for multiple databases <p>or</p> <ul style="list-style-type: none"> ■ Update an existing snapshot set schedule 	<p>From the Solutions Configuration Center: Quick Recovery Configuration Wizard</p>	<p>“About the Quick Recovery Configuration Wizard” on page 38</p>
<ul style="list-style-type: none"> ■ Add a new snapshot set schedule for one database 	<p>From the VEA: VSS Exchange Snapshot Scheduler Wizard</p> <p>Alternatively, you can use the Quick Recovery Configuration Wizard.</p>	<p>“About the VSS Exchange Snapshot Scheduler Wizard” on page 39</p>
<ul style="list-style-type: none"> ■ Set up a snapshot schedule for passive copies in an Exchange 2010 Database Availability Group (DAG) 	<p>From the VEA: VSS Exchange Snapshot Scheduler Wizard</p> <p>Note: Creating a snapshot schedule for passive copies in a DAG is not supported by the Quick Recovery Configuration Wizard.</p>	<p>“Exchange 2010 Database Availability Group (DAG) considerations” on page 47</p>
<ul style="list-style-type: none"> ■ Create a one-time snapshot as needed 	<p>From the VEA: VSS Exchange Snapshot Wizard and VSS Snapback Wizard</p> <p>Alternatively, you can use the Vxsnap utility from the CLI.</p>	<p>“About the VSS Snapshot and Snapback wizards and the vxsnap utility” on page 40</p>

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard provides the most complete method of configuring snapshot sets and is therefore recommended for initial

configuration. The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirror for the snapshot set
- Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again

The wizard enables you to set up and schedule multiple snapshot sets for one or more databases of the selected Exchange instance. You can set up one or more schedules for each snapshot set.

You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. For example, you can schedule mirror preparation, the most time-consuming step, for a time of low activity.

The scheduling capability automates the process of refreshing the snapshot sets. At the scheduled times, the snapshot volumes are automatically reattached, resynchronized, and then split again. Once configured, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The wizard also enables you to save all configuration settings to a .tpl file (an XML file) that can be imported for use on other systems.

See [Chapter 6, “Implementing snapshot sets with the configuration wizard”](#) on page 61.

Note: The Quick Recovery Configuration Wizard does not support scheduling snapshots of the passive copy in an Exchange 2010 DAG.

About the VSS Exchange Snapshot Scheduler Wizard

You can use the VSS Exchange Snapshot Scheduler Wizard to add a snapshot schedule for a selected database. This wizard uses the same scheduling process as the Quick Recovery wizard.

Unlike the Quick Recovery Configuration Wizard, the VSS Exchange Snapshot Scheduler Wizard does not prepare the snapshot mirror. To prepare the snapshot mirror you must use the VEA Prepare command (or the vxsnap command line utility) before running the VSS Exchange Snapshot Scheduler Wizard. Also, you can only schedule a snapshot for one database at a time.

You would typically use this wizard to add a schedule to the initial configuration that was created with the Quick Recovery Configuration Wizard. For example, you configured a daily snapshot for Quick Recovery use and now want to add a weekly snapshot on a different disk for backup use.

If replication is enabled for an Exchange 2010 database in a DAG, you can also use the VSS Snapshot Scheduler Wizard to schedule snapshots of the passive copy of an Exchange 2010 database.

Note: Adding a snapshot schedule using the VSS Exchange Snapshot Scheduler does not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

See [Chapter 7, “Scheduling or creating an individual snapshot set”](#) on page 81.

About the VSS Snapshot and Snapback wizards and the vxsnap utility

If you need to create an individual one-time snapshot, you can do so either from the command line, using the vxsnap command line utility, or from the VEA console using the VSS Snapshot and Snapback wizards.

You must prepare the mirror for the snapshot manually, using the Prepare command. In addition, you cannot use the Snapshot wizard or command line to schedule periodic refreshing of the snapshot. Instead you manually reattach the mirror, allowing it to resynchronize, and then create a snapshot again from the resynchronized mirror.

Therefore, these methods are best suited for a one-time special-purpose snapshot. If you need to maintain a snapshot set for future use, you should instead set up a snapshot schedule using the VSS Exchange Snapshot Scheduler Wizard or the Quick Recovery Configuration Wizard.

See [Chapter 7, “Scheduling or creating an individual snapshot set”](#) on page 81.

If replication is enabled for an Exchange 2010 database in a DAG, these SFW snapshot methods can also be used to create a snapshot of the passive copy of an Exchange 2010 database.

[Table 4-2](#) shows the corresponding vxsnap commands for the VSS wizards.

Table 4-2 Actions of VSS wizard and vxsnap command

Action	VSS wizard	vxsnap command
Creates a set consisting of snapshots of all the volumes in the specified Exchange database.	VSS Snapshot	<code>create</code>
Reattaches and resynchronizes a snapshot set to the production database volumes.	VSS Snapback	<code>reattach</code>

Note: The `vxsnap` commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode. For more information about the `vxsnap` command line utility, see [Chapter 10, “Vxsnap utility command line reference”](#) on page 153.

Planning your Quick Recovery solution

Consider the following questions as you plan your Quick Recovery solution:

- Should more than one snapshot set be created?
Symantec recommends maintaining a snapshot set that is refreshed after each full backup. A separate snapshot set that is refreshed after each incremental backup would reduce the time it takes to replay the logs in a point-of-failure recovery.
- How often and when will each snapshot set be refreshed?
- Is sufficient disk space available for multiple snapshot sets as well as the production volumes?
Each snapshot set contains split-mirror snapshots of all the volumes associated with a database and requires the same amount of space as the original volumes.
- Do you want to create a snapshot set schedule for use on multiple systems?
You can use the Quick Recovery Configuration Wizard to save a template to an XML file for use on other systems.

Backup types for snapshot sets

When creating a snapshot set, you can choose the backup type depending on how you plan to use the snapshot set:

- Copy backup is generally used for a snapshot set to be maintained on-host and used for quick recovery of a missing or corrupt Exchange database. This mode creates a snapshot or copy of the transaction log, but does not truncate it. You can choose whether or not to run the Eseutil consistency check.
- Full backup is generally used for a snapshot set to be used for snapshot-assisted backup. Full backup runs the Eseutil consistency check and truncates the transaction logs.

About logs

Microsoft Exchange employs “write-ahead” logging to improve performance. As transactions occur, they are appended to a transaction log. Transactions are committed to the database when resources permit. A checkpoint file (EOn.CHK) tracks the location of the last transaction in the log files that has been successfully committed to the database. Transaction logs can be used to roll forward a database to achieve a point-of-failure recovery.

Note: Do not enable the circular logging option. If circular logging is enabled, you will not be able to roll forward a database to achieve a point-of-failure recovery. Microsoft does not recommend circular logging in production environments. Refer to the Microsoft Exchange Server documentation for further details.

Recommendations and best practices

You should review the recommendations and best practices for the following:

- [“Recommendations for Exchange storage configuration”](#) on page 43
- [“Recommendations for maintenance and backups”](#) on page 44
- [“VCS cluster considerations”](#) on page 44
- [“Exchange 2010 Database Availability Group \(DAG\) considerations”](#) on page 47
- [“VVR considerations”](#) on page 48

Recommendations for Exchange storage configuration

To use the Quick Recovery snapshot functionality with Exchange databases, you must place the databases on Veritas Storage Foundation for Windows (SFW) dynamic volumes. The following recommendations enable you to take advantage of SFW storage configuration functionality as you manage your Exchange storage:

- Databases and transaction logs must be stored on disks contained within a single dynamic disk group.
- Each database should be in a separate volume, but the volumes may share the same dynamic disks.
- Mailbox stores and public stores must be stored on separate volumes in order to be able to recover each independently.
- Database stores and transaction logs must be in separate volumes in order to perform a roll-forward recovery to the point of failure.
- Database stores and transaction logs should be on separate disks so that disk failure does not affect both the database stores and transaction logs.
- Transaction logs should always be configured in a redundant layout. The preferred software layout is RAID 0+1 (mirrored striped) volumes as this provides better read and write performance than RAID 1 (mirrored) alone. The transaction log will generate the most I/O and thus should use the highest performance disks available.
- The preferred layout for the database stores is hardware RAID 5, software RAID 1 (mirrored with logging enabled) or software RAID 0+1 (mirrored striped).

Note: FlashSnap is not supported for software RAID 5 volumes.

- When configuring volumes on a replication node of an Exchange 2010 DAG, use the same drive letters as the volumes on the active node.
- By default, during installation, a database is created in the install location, which is the boot drive. Snapshots cannot be created on the boot drive. In order to use VSS snapshots on the database, you must move the database components from the boot drive to an SFW dynamic volume on another drive.

See [“Pointing the database and log paths to the SFW volumes”](#) on page 59

Recommendations for maintenance and backups

- Locate the snapshot volumes for each database on separate disks from snapshots of other databases. This is recommended so that the process of creating the snapshot of one database doesn't interfere with any operations on another database.
- Each snapshot set is associated with a metadata XML file that is needed for Quick Recovery restoration. Create a unique name for the metadata XML file of each snapshot set you create and maintain. The snapshot XML files should be stored separately from the volumes that are included in snapshots. Snapshot schedules have additional XML files associated with them. By default, the snapshot XML files are created in the following paths. For Windows Server 2008:

`C:\ProgramData\Veritas\VSSXML\application name`

- For Quick Recovery solutions, Symantec recommends that you create or refresh a snapshot set immediately after a Full Backup just after the database has been checked for corruption and the transaction logs have been truncated. Thus, you are assured an image of a clean database. Additionally, you may wish to create another snapshot set after an Incremental Backup. Create this snapshot set on a separate set of disks rather than refreshing the snapshot set taken after the Full Backup. This practice ensures you are not overwriting a snapshot set of a clean database with an image of a potentially corrupted database.

VCS cluster considerations

In a VCS cluster environment, observe the following precautions:

- The XML metadata file for each snapshot set along with files that store snapshot schedule information are created in a folder on the local drive by default. In a cluster environment, store these files on shared storage so that the files are available from all nodes in the cluster. The snapshot XML files should be stored separately from the volumes that are included in snapshots.
 - If you use the Quick Recovery Configuration Wizard to create the snapshot set, you can use the wizard to specify the file path to the appropriate volume.
 - If you use a VSS wizard to create the snapshot set, you can store the XML files in a location of your choice using one of the following methods:
 - Edit the directory path in the Directory field in the VSS wizard.

- Use a text editor to create a text file named **redirect.txt**. This text file should contain a single text line specifying the full path to the location of the metadata file, for example, `G:\BackupSets`. Save the `redirect.txt` file in the default VSS XML file directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML` on each node of the cluster.
- When using `vxsnap` utility commands that require the filename attribute, specify the full path to the location of the XML metadata file.
- If an Exchange 2010 mailbox database is configured under a VCS cluster, you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. In the case of a scheduled snapshot, even though schedules are visible on all nodes, the snapshot will happen only on a node where the mailbox database is present. There can be a large number of mailbox databases in Exchange 2010 and each database may have independent schedules. Hence the number of different paths required for storing the snapshot metadata files is higher, which is achieved by configuring a file share resource. You can also specify the file share path in the `redirect.txt` file, for example, `\\MySystemName1\share1\QRDir`.
 In the case of a VCS clustered setup with more than one node, configure the Veritas Scheduler Service with any user account (other than the Local System account) which is valid on all the nodes of the cluster. This user account should have read-write permissions to the file share path.
 If a schedule is created on a file share, it is visible on all nodes of the cluster and can be deleted from any node irrespective of where the Exchange mailbox database component is online.
- If you plan to use the Quick Recovery or VSS Snapshot Scheduler wizard to specify scripts (commands) to be run before or after a snapshot, store the scripts on shared storage so that they are available to all nodes.
- If you need to restore an Exchange 2010 database from a snapshot, running the VSS Restore Wizard brings the VCS resource for the database offline, dismounting the database before beginning the restoration. You can if necessary bring the resource offline in VCS manually to dismount the database. Once the wizard successfully restores the database, the wizard brings the resource back online, mounting the database; if not you can bring the resource online manually. If you are using the CLI, specifying the `-a` option in the `vxsnap restore` command offlines and onlines the resource automatically. See [Chapter 8, “Recovering Exchange mailbox databases”](#) on page 105.

Recommendations and best practices

- If you set up a snapshot schedule with the Quick Recovery wizard and later add a node to the cluster, you can run the wizard again to synchronize schedules on the existing nodes with the new node.
- If you set up a snapshot schedule with the VSS Snapshot Scheduler wizard, before adding or removing a node from a VCS cluster setup, delete the schedules and then recreate the schedules on the required node.
- Exchange must be installed on all cluster nodes.

Exchange 2010 Database Availability Group (DAG) considerations

See the following topics for using snapshots with an Exchange 2010 DAG:

- [“Using snapshot schedules in a DAG”](#) on page 47
- [“Guidelines for creating passive copy snapshots in a DAG”](#) on page 48
- [“Recovery in a DAG”](#) on page 48

Using snapshot schedules in a DAG

In an Exchange 2010 DAG, you may want to create snapshot schedules on the replication (passive) nodes of the DAG as well as the active node. The VSS Schedule Snapshot wizard supports scheduling snapshots on passive copies.

Note: The Quick Recovery Configuration Wizard does not support snapshots for a passive copy. If the active database copy switches to another node, a Quick Recovery Wizard schedule on the node that is now passive is no longer valid and scheduled snapshot operations on the passive node will fail.

For example, if the DAG contains an active database on NodeA that is being replicated to NodeB, you can set up a snapshot schedule on NodeB as well as on NodeA. Setting up a snapshot schedule on the replication node ensures that snapshots continue as scheduled if a database fails on the active node.

Exchange 2010 provides two VSS Writers as follows:

- The Store Writer is available on any mailbox server and is used for backup and restore of active databases.
- The Replication Writer is available on the passive node of a DAG server. The Replication Writer supports backup functionality for a selected database where the snapshot is taken against the replicated instance of the database and transaction log files. The Replication Writer is supported by SFW Flashsnap for snapshot and schedule snapshot operations but not for restore operations.

If you have set up schedules on both the active and passive nodes, the appropriate VSS writer is used automatically for the scheduled snapshot. For example, if the database on NodeA changes from an active to a passive copy, the snapshot schedule is maintained with the Replication Writer if the database remains in a healthy state. If you have set up a snapshot schedule on NodeB and the database on NodeB changes from a passive to an active copy, the snapshot schedule on NodeB continues using the Store Writer.

Guidelines for creating passive copy snapshots in a DAG

Creating a snapshot of a passive database copy requires using the Microsoft Exchange Replica Writer to prepare the snapshot mirror and to create the snapshot. If replication is enabled:

- The Microsoft Exchange Replication Writer instance is displayed in the tree view of the VEA under the Applications node. Right-clicking the Exchange Replica node displays a context menu that shows VSS Schedule Snapshot, VSS Snapshot, VSS Snapback, and VSS Refresh selections.
- Using the `vxsnap` CLI command, you can specify Microsoft Exchange Replica Writer (instead of Microsoft Exchange Writer) for the `writer` option in the required commands.

See [Chapter 7, “Scheduling or creating an individual snapshot set”](#)

Recovery in a DAG

Before beginning a recovery in a DAG environment, ensure that you review the recovery procedures in [Chapter 8, “Recovering Exchange mailbox databases”](#). Ensure that you complete all steps in the sequence shown. See the following recovery procedures:

- [“Recovering using the VSS Restore Wizard”](#) on page 108
- [“Recovering using the vxsnap utility”](#) on page 112

Although you cannot restore a snapshot to a passive copy, you can use a passive copy snapshot to restore a database after first manually performing a "Move-ActiveMailboxDatabase" on the database. For steps on restoring using a passive copy snapshot, see the following:

- [“Recovery using an Exchange 2010 passive copy snapshot in a Database Availability Group \(DAG\)”](#) on page 107

VVR considerations

In a VVR environment, observe the following precautions:

- Store the XML metadata file and other snapshot related files on a volume that is included in the replicated data set so that the metadata file is available at the secondary site. Additionally, if VCS is used, store the metadata file on a volume in the cluster disk group associated with the Exchange database so that the metadata file is available from all nodes in the cluster.

See [“VCS cluster considerations”](#) on page 44.

- During a point-in-time recovery, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data

consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

See “[Post-recovery steps](#)” on page 116 in [Chapter 8, “Recovering Exchange mailbox databases”](#).

Configuring Exchange for Quick Recovery snapshots

This chapter includes the following topics:

- [Tasks for configuring Exchange storage](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Creating dynamic disk groups](#)
- [Creating dynamic volumes](#)
- [Pointing the database and log paths to the SFW volumes](#)

Tasks for configuring Exchange storage

[Table 5-1](#) outlines the high-level objectives and the tasks to complete each objective for configuring Exchange storage for Quick Recovery snapshots.

Table 5-1 Tasks for configuring Exchange storage for Quick Recovery snapshots

Objective	Tasks
“Reviewing the prerequisites” on page 52	■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 53	■ Reviewing the sample configuration
“Creating dynamic disk groups” on page 54	■ Creating dynamic disk group for each Exchange database
“Creating dynamic volumes” on page 55	■ Creating volumes for the database stores and transaction logs in each database
“Pointing the database and log paths to the SFW volumes” on page 59	■ Using Exchange to set the database and log paths for new or existing databases to point to the SFW volumes

Reviewing the prerequisites

This solution assumes that the required SFW software is already installed and configured.

Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation and configuration information.

Exchange 2010 requires upgrading to Veritas Storage Foundation for Windows 5.1 Service Pack 1 Application Pack 1 or Veritas Storage Foundation HA for Windows 5.1 Service Pack 1 Application Pack 1. For upgrade information, refer to the *Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Application Pack 1 Release Notes*.

The solution also assumes that the required Exchange software is installed.

You should also familiarize yourself with the system requirements and recommendations and best practices for configuring Exchange storage for snapshot solutions.

See [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).

Reviewing the configuration

Table 5-2 shows the object names that are used to illustrate the tasks you perform when configuring Exchange storage with Veritas Storage Foundation for Windows:

Table 5-2 Object names for Exchange

Name	Drive Letter	Object
Exch1		Exchange Server host name
Exch-DG1		Dynamic disk group
Harddisk1, Harddisk2, Harddisk3		Disks included in Exch-DG1
DB1-tlogs	I:	Transaction logs volume
DB1	G:	Mailbox database volume
DB1-tlogssnap		Snapshot volume of DB1-tlogs
DB1snap		Snapshot volume of DB1
Image1.xml		Name for the metadata file that contains information about the snapshot set

Creating dynamic disk groups

Create one dynamic disk group for each Exchange database.

If your Exchange production server is in a clustered environment, choose the cluster disk group option.

See *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange* for further configuration information on the SFW HA clustered environment.

Note: Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before implementing SFW snapshot solutions. Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group. For more information, see *Veritas Storage Foundation Administrator's Guide*.

To create a dynamic disk group from the VEA console

- 1 Click **Start > Programs > Symantec > Veritas Storage Foundation>Veritas Enterprise Administrator** and if prompted to select a profile, select a profile (or Default).
- 2 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
For the local system you can specify **localhost**.
- 3 If prompted to do so, specify the user name, password, and domain for the system.
- 4 In the tree expand the system name and expand the storage agent.
- 5 Right-click **Disk Groups**, and click **New Dynamic Disk Group**.
- 6 On the **Welcome** screen of the New Dynamic Disk Group Wizard, click **Next**.
- 7 Enter a name for the disk group (for example, Exch-DG1).
- 8 For an off-host or cluster environment, choose from the following:
 - For a cluster environment, check the **Create cluster group** check box.
 - For an off-host environment, check the **Add private group protection** check box.
- 9 Select the appropriate disks in the Available disks list and click the **Add** button to move them to the Selected disks list.
- 10 Click **Next**.

- 11 Click **Next** to upgrade the selected disks.
- 12 Click **Finish** to create the new disk group.

To create a dynamic disk group from the command line

- Type the following command:

For an on-host environment

```
vxdg -gExch-DG1 init Harddisk1 Harddisk2 Harddisk3
```

For an off-host environment

```
vxdg -gExch-DG1 -R init Harddisk1 Harddisk2 Harddisk3
```

For a cluster environment

```
vxdg -gExch-DG1 -s init Harddisk1 Harddisk2 Harddisk3
```

where Exch-DG1 is the name of the dynamic disk group you want to create and Harddisk1, Harddisk2, and Harddisk3 are the disks included in the dynamic disk group.

For the complete syntax of the `vxdg init` command, see the *Veritas Storage Foundation Administrator's Guide*.

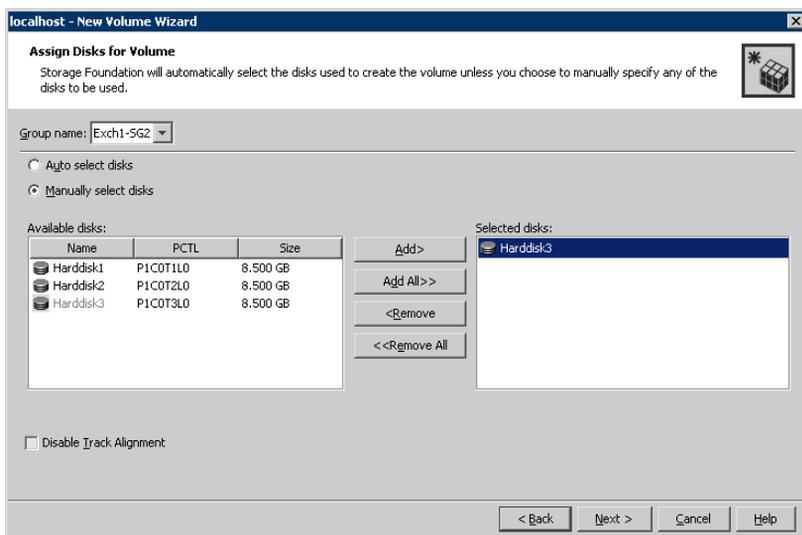
Creating dynamic volumes

Create a volume for the log files and an additional volume for each database.

To create a dynamic volume from the VEA console

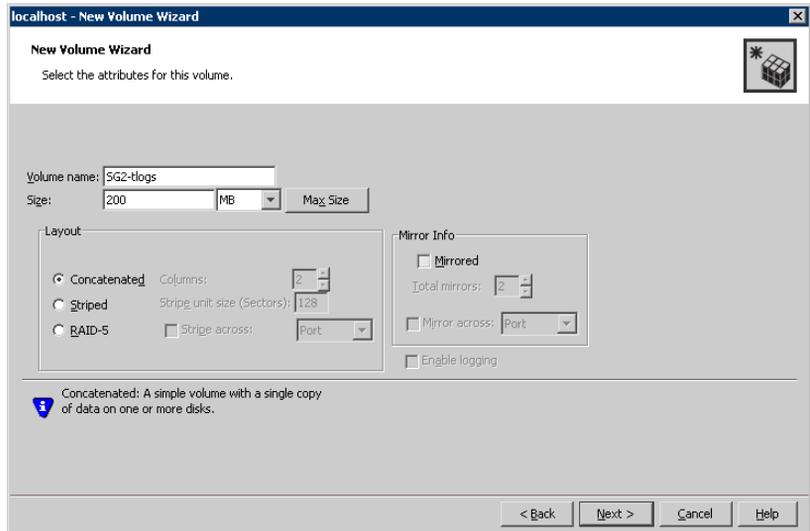
- 1 Start the VEA and connect to the appropriate host.
- 2 In the tree, expand the system name, expand the storage agent, and then expand **Disk Groups**.
- 3 Right-click on the disk group in which to create the volumes (for example, Exch-DG1) and click **New Volume**.
- 4 In the Welcome panel of the New Volume Wizard, click **Next**.

5 Select the disks for the volume:



- Confirm that the **Group name** is the correct disk group (Exch-DG1). If necessary, select the correct disk group name from the drop-down menu.
- Specify automatic or manual disk selection. Symantec recommends using the **Manually select disks** option.
- Select the appropriate disks in the Available disks list, and click the **Add** button to move them to the Selected disks list.
- You may also check Disable Track Alignment to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

6 Specify the parameters of the volume:



- Enter the volume name (for example, DB1-tlogs).
 - Enter the size.
 - Select the layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 7 Assign a drive letter to the volume (for example, I: for the DB1-tlogs volume) and click **Next**.
- If you are assigning a drive letter to a replication volume in an Exchange 2010 DAG, use the same drive letter as the volume on the active node.
- 8 Create an NTFS file system:
- Accept the default **Format this volume**.
 - Click **NTFS**.
 - Select an allocation size.
 - Accept the default file system label, which is the same as the volume name you entered previously or enter a file system label.
 - If desired, select **Perform a quick format**.
 - Click **Next**.
- 9 Review the volume specifications, then click **Finish** to create the new volume.

- 10 Repeat the previous steps as necessary to create volumes for the databases. For example, create volumes DB1 and DB2.

To create a volume from the command line

- 1 Type the CLI command, as in the following example:

```
vxassist [-b] -gExch-DG1 make DB1-tlogs 200 DriveLetter=I
```

This example command will create a 200 MB volume named DB1-tlogs on drive letter I: in the Exch-DG1 dynamic disk group.

Note: This command does not format the volume (a file system is not created on the volume). You must use the operating system format command to format the volume.

- 2 Modify and repeat this command as necessary to create volumes for the databases.

For the complete syntax of the `vxassist make` command, see *Veritas Storage Foundation Administrator's Guide*.

Pointing the database and log paths to the SFW volumes

For each Exchange 2010 database, set the database and log files paths to point to the appropriate SFW volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new databases or change the locations for existing databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

Note: When Exchange Server 2010 is installed, a database is created on the boot drive by default. The boot drive cannot be snapshotted. If you want to implement a snapshot solution for that database, you must move it to new volumes not on the boot drive.

The following procedure can be used to change paths for existing databases to point to the SFW volumes you created for them.

To use the Exchange Management Console to change file paths for an existing database

- 1 Start the Exchange Management Console (**Start > Administrative Tools > Exchange Management Console**) on the server on which the database is located.
- 2 In the console tree, navigate to **Organization Configuration > Mailbox**.
- 3 In the result pane, on the Database Management tab, select the database you want to configure.
- 4 In the work pane, click **Move Database Path**.
- 5 In the Move Database Path wizard, edit the **Database file path** field to specify the SFW volume location for the database. Edit the **Log folder path** field to specify the SFW volume location for the log files.
- 6 Click **Move**.
- 7 View the status of the move operation. The wizard moves the database file path and the log folder path to the new location. Click **Back** if you want to make configuration changes.

- 8 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Implementing snapshot sets with the configuration wizard

This chapter covers the following topics:

- [About the Quick Recovery Configuration Wizard](#)
- [Tasks for implementing snapshot sets with the configuration wizard](#)
- [Reviewing the prerequisites](#)
- [Scheduling and creating snapshot sets](#)
- [Viewing the status of scheduled snapshots](#)
- [Troubleshooting scheduled snapshots](#)
- [Deleting or modifying schedules](#)
- [Synchronizing schedules after adding a cluster node](#)

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirrors
- Creating the initial snapshot sets
- Periodically refreshing the snapshot sets

The wizard enables you to set up and schedule multiple snapshot sets for one or more Exchange components. For Microsoft Exchange 2003 and 2007, Exchange

components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

You can set up different schedules for each snapshot set. You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. The scheduling capability automates the process of refreshing the snapshot sets. At the scheduled time, the snapshot volumes are automatically reattached, resynchronized, and then split again.

The snapshot creation process integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes. The snapshot is done while the databases are online and without disrupting the email flow.

Once configured and applied, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The wizard includes the following settings:

- Which Exchange component to snapshot.
- Number of snapshot sets for each Exchange component
- Volume layout and disk location for each snapshot set
- When to prepare the selected volumes for the snapshots
- When and how often snapshots occur
You can schedule snapshot run days using days of the week, days of the month, and include or exclude dates. You can specify that a schedule recurs daily or uses some other recurrence interval. You can also schedule a daily time window in which the snapshots occur, and the intervals for refreshing snapshot sets within the time window.
- Backup type (Full, Copy, or Copy plus Eseutil)

Optionally, you can also specify scripts to run before and after snapshots.

After you configure the settings, you can do one or both of the following:

- Save the settings in a template file for reuse on other systems. Even if you do not save the settings to a template file, they are still stored for the Exchange component to which they apply. You can access them later by selecting the same instance and Exchange component in the wizard.
- Implement the settings immediately.

About snapshot sets

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the Exchange component, the snapshot set requires the same amount of space as the original volumes. For Microsoft Exchange 2003 and

2007, Exchange components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

Each snapshot set can be created as a different backup type: full copy including Eseutil consistency check and transaction log truncation, copy with Eseutil, or copy without Eseutil or transaction log truncation.

Snapshot set creation has the following stages, all of which you can schedule using the Quick Recovery Configuration Wizard:

- Preparing the mirror for the snapshot set. This stage takes a while and should be scheduled for a time of low activity. The remaining stages take little time.
- Creating the initial snapshot by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again. This stage occurs for each snapshot set at the times or frequency that you specify. For example, you can schedule one snapshot to occur every 30 minutes starting at 9 A.M. and ending at 6 P.M.

About snapshot templates

A snapshot template contains all the settings that you implemented with the Quick Recovery Configuration Wizard for the selected Exchange component. For Microsoft Exchange 2003 and 2007, Exchange components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

After configuring settings with the wizard, you can save the template settings to a template (.tpl) file (an XML file) for use on another Exchange component. This enables you to re-use the same schedule or use it with minor modifications. If you apply the settings without saving, they are stored in the registry so that you can view or modify them in the wizard, but they are not available to import for another instance or Exchange component.

Before you can use the template settings on another system, you must first copy the template (.tpl) file to that system. You can then import the file while using the Quick Recovery Configuration Wizard.

If settings already exist for an Exchange component that you selected in the wizard, you cannot import a template to overwrite those settings. You can only modify the settings in the wizard. However, you can import a template for another Exchange component for which settings do not exist.

Templates store the following information:

- The selected Exchange component

- The number of snapshot sets for each Exchange component
- When to prepare the snapshot mirrors
- The snapshot schedule and rules related to the schedule, including backup type and the names of optional scripts
- The current date and time when the template is created
- The application name

Templates and multiple components

When you apply a template that has multiple components, the wizard first attempts to match the names of components (storage group for Microsoft Exchange 2003 and 2007, or database for Microsoft Exchange 2010) in the template to the Exchange component you selected in the wizard.

If it matches a name, it applies the information stored under that name in the template to that Exchange component. For example, if you select Exch_SG1 for a Microsoft Exchange 2003 or 2007 component or Exch1-DB1 for a Microsoft Exchange 2010 component in the wizard and the template has settings for these Exchange components, it applies those settings.

An Exchange component selected in the wizard may not match the names of any Exchange component in the template. In that case, the wizard applies the information for the first unapplied component in the template to the first selected Exchange component that does not match any name in the template. It continues in that sequence.

If you selected more Exchange components in the wizard than the number of components in the template, the wizard prompts you to fill in any required information for the remaining Exchange components.

Templates and schedule start dates

Templates contain settings for one or more schedule start dates, including a mirror preparation date, a "Start on" date, and a "Schedule takes effect on" date.

If you import a template after a start date has elapsed, the wizard tries to maintain the same delta between the template import date and a start date as the delta between the template creation date and a start date. It shifts the dates forward to maintain this delta.

Therefore, after importing the template, use the wizard to review the settings and adjust dates as necessary to meet your needs.

If you import a template that uses the current date as the "takes effect on" date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, a schedule includes two daily snapshots, one in the morning and one at night, and you import the

schedule in the afternoon. In this case neither of the snapshots will occur on the current date. Instead the effective date is shifted ahead one day.

Tasks for implementing snapshot sets with the configuration wizard

[Table 6-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 6-1 Tasks for implementing snapshot sets with the configuration wizard

Objective	Tasks
“Reviewing the prerequisites” on page 65	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Scheduling and creating snapshot sets” on page 66	<ul style="list-style-type: none"> ■ Running the Quick Recovery Configuration Wizard to configure and schedule snapshot sets ■ Optionally, saving the configuration and schedules to a template file ■ Applying the configuration
“Viewing the status of scheduled snapshots” on page 75	<ul style="list-style-type: none"> ■ Viewing the status of scheduled snapshots to determine whether they were successful
“Deleting or modifying schedules” on page 78	<ul style="list-style-type: none"> ■ Deleting schedules that are no longer needed or modifying schedules.
“Synchronizing schedules after adding a cluster node” on page 79	<ul style="list-style-type: none"> ■ In a clustered environment, synchronizing schedules after adding a cluster node

Reviewing the prerequisites

Before running the Quick Recovery Configuration Wizard, you should do the following:

- Ensure that your system hardware and software meets the requirements. See [“System requirements”](#) in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).

- Review the recommendations and best practices.
See [“Recommendations and best practices”](#) on page 42 in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).
- Set up your Exchange component (storage group or database depending upon the version of Microsoft Exchange you are using) for use with Storage Foundation for Windows (SFW).
See [Chapter 5, “Configuring Exchange for Quick Recovery snapshots”](#) on page 51.
- For Microsoft Exchange 2003 or 2007, ensure that the mailbox stores for which you plan to schedule snapshot operations are mounted. For Microsoft Exchange 2010, ensure that the mailbox databases for which you plan to schedule snapshot operations are mounted.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- Ensure that you are logged in as a domain administrator or as a member of the Domain Admins group.
- If a firewall exists between the wizard and any systems it needs access to, set the firewall to allow both ingoing and outgoing TCP requests on port 7419.
- In a clustered server environment, ensure that the Veritas Scheduler Service is configured and running with domain administrator privileges on all nodes in the cluster.
- To use a previously created template, copy it to the server that you are configuring.
- Ensure that the Microsoft Software Shadow Copy Provider service is running. For a clustered server environment, ensure that the service is running on all systems in the cluster.

Scheduling and creating snapshot sets

You schedule and create snapshot sets using the Quick Recovery Configuration Wizard. You can also use the wizard to modify or delete existing schedules or import a schedule template.

See [“About the Quick Recovery Configuration Wizard”](#) on page 61.

You should ensure that you meet the prerequisites before you begin using the wizard.

See [“Reviewing the prerequisites”](#) on page 65.

In addition, if you plan to reuse settings for the Exchange component (storage group or database depending upon the version of Microsoft Exchange you are using) on other instances, you should be familiar with the information about snapshot templates.

See “[About snapshot templates](#)” on page 63.

Note: For Exchange 2007, if using local continuous replication (LCR), you can schedule snapshots only for the active copy, not the passive copy.

Note: For Exchange 2010, if using replication in a DAG, you can schedule snapshots only for the active copy, not the passive copy.

To schedule and create snapshot sets

- 1 Start the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**), expand the application on the Solutions tab, expand **Quick Recovery Configuration**, and start the Quick Recovery Configuration Wizard.
- 2 In the Welcome panel, review the information to ensure you meet the requirements and click **Next**.
- 3 In the System Selection panel, specify the fully qualified host name or IP address of the system that is running the application (or specify localhost for the local system), and click **Next**.
 For failover of an application virtual server in a clustered server (high availability) environment, Symantec recommends that you specify the virtual server name or virtual IP address.
 Ensure that the disk groups for the application are imported to the selected system.
- 4 In the Instance Selection panel, specify the following options and click **Next**:

Set up Quick Recovery for	For Exchange 2003 or 2007, select the Exchange virtual server name. For Exchange 2010, select the system name of the Exchange mailbox server. The selection details are displayed in the panel.
---------------------------	--

Select objects or view details Depending upon your version of Microsoft Exchange, the Components list shows either storage groups or databases. Select the component for which you want to configure the snapshot sets.

For Exchange 2010, you can select either VCS clustered databases or standalone, but not both together. If you are selecting clustered databases, they must be configured to fail over to the same list of systems.

The wizard checks for existing schedules for the selected component. The wizard cleans up any obsolete schedule files and entries. In a clustered server environment, the wizard will synchronize schedules between cluster nodes to ensure that all existing schedules are available on the selected node.

- 5 The wizard validates that the volumes containing the schedule configuration files are mounted. If a volume is not mounted, the Mount Details panel displays the information about the missing drive. Mount the missing drive and click **Refresh**, then click **Next**. Otherwise, the schedule is deleted.
- 6 For existing schedules, the Synchronizing Schedules panel displays the status of schedule synchronization and cleanup. If schedule synchronization fails in a clustered environment, restart any cluster nodes that are down and then restart the wizard. Otherwise, click **Next**.
- 7 In the Template Selection panel, select from the following options and click **Next**:

Create or modify a template Create or modify the Quick Recovery settings for the selected components.

If at least one of the selected components has existing settings, this option is the only choice. In this case, the existing settings are displayed for you to view or modify as you continue through the wizard panels.

If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.

Import a template Import Quick Recovery settings from a template (.tpl) file created by the Quick Recovery wizard earlier. Browse to the file location and select the file. The file must be located on the system for which Quick Recovery is being configured. The file also must be of the same application type as the application you are configuring.

This option is available only if no settings exist for a selected component.

- 8 In the Number of Snapshot Sets panel, select how many sets of snapshots to create for each database. Remember that each snapshot set uses the same amount of space as the original volumes.
- 9 In the Snapshot Volume Assignment panel, expand the snapshot set and make the following selections:

Snapshot set In the box under the Snapshot Set column heading, optionally edit the snapshot set name. If you edit the name, ensure that the name is unique among all snapshot sets for all databases.

If you are modifying an existing schedule, the snapshot set name cannot be changed.

XML Metadata File Name Specify a name for the XML file that will be associated with the snapshot set. This file is used for recovery, so you may want to assign a name to easily identify it for that purpose. Ensure that the XML file for each snapshot set has a unique name.

If you are modifying an existing schedule, the XML file name cannot be changed.

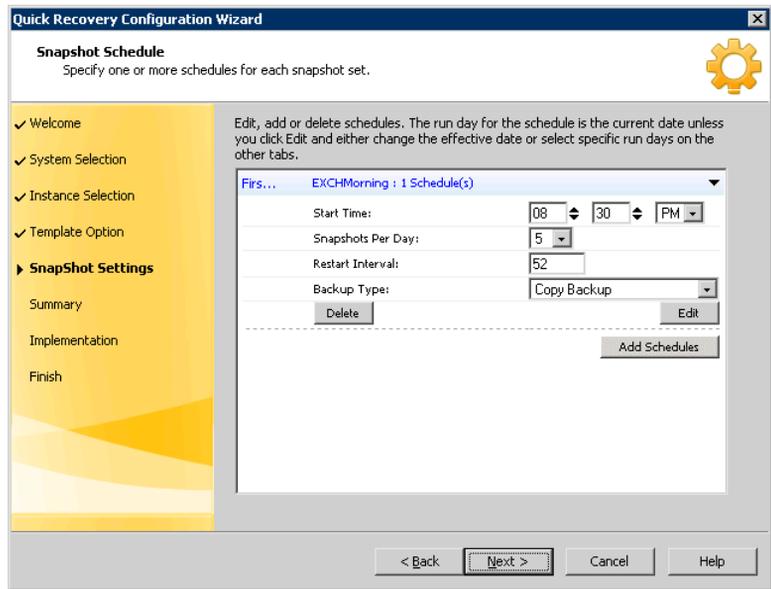
Mirror Preparation Time Click **Edit** and in the Mirror Preparation dialog box, specify the date and time for mirror preparation and click **OK**.

Preparing the snapshot mirrors slows down performance, so you may want to choose a time of low usage.

Snapshot Disks Assign one or more disks to each snapshot volume. Click the icon next to the disk column to select from available disks and to select a concatenated or striped volume layout.

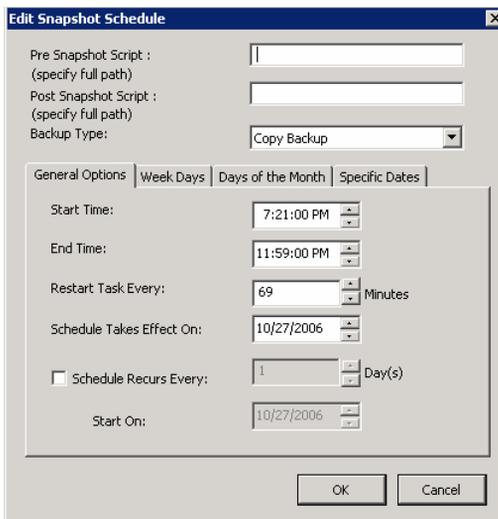
Mount	<p>Optionally, click Set to set a drive letter or mount path.</p> <p>The drive letter specified may not be available when the snapshot operation is performed. When this occurs, the snapshot operation is performed, but no drive letter is assigned.</p>
File path for snapshot XML file	<p>File path for the snapshot XML files. The path specified applies to all the snapshot sets that you are currently creating, for all currently selected databases. If the field is not editable, you are modifying an existing snapshot set schedule, so the value cannot be changed.</p> <p>For a non-clustered environment, the default location on Windows Server 2003 is:</p> <pre>C:\Documents and Settings\All Users\ Application Data\VERITAS\VSSXML\ application name</pre> <p>and on Windows Server 2008:</p> <pre>SystemDrive:\ProgramData\VERITAS\VSSXML\ application name</pre> <p>If a <code>redirect.txt</code> file has been created, the path in that file is shown in the field.</p> <p>For new schedules for a clustered server environment, if there is no <code>redirect.txt</code> file, the field is empty. Enter the full path to a location on shared storage, for example: <code>G:\SnapshotSets</code>.</p> <p>For Exchange 2010 configured in a VCS cluster, you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. You can also specify the file share path in the <code>redirect.txt</code> file.</p> <p>You must store the XML files on a separate volume from the volumes that are included in snapshots.</p> <p>Note: If the snapshot XML files are not stored separately from the volumes that are included in the snapshots, a restore will fail.</p> <p>Ensure that you use a consistent location. That way, you can easily find the XML metadata file, which is required for recovery.</p>

- 10 Repeat [step 9](#) for any additional snapshot sets and when you finish, click **Next**.
- 11 In the Snapshot Schedule panel, choose one of the following:



- To specify a simple schedule, edit the default schedule settings shown on this panel: the time to start the snapshot (at least 1 hour later than the scheduled mirror preparation time), the number of snapshots to take, the interval in minutes between snapshots, and the type of snapshot. By default, the simple schedule takes effect on the current date and does not recur. Continue with [step 16](#).
- To specify a recurring schedule, a time window, a different run date, or other schedule details, click **Edit**.
 Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.

- 12 If you clicked **Edit**, in the Edit Snapshot Schedule dialog box, choose from the following settings and then make additional selections on the dialog box tabs:



Pre Snapshot Script

Optionally, specify the full path of a script to run before the scheduled snapshot occurs.

In a cluster environment, script files should be located on shared storage.

For security purposes, ensure that script files are not accessible to any user with fewer privileges than the user account under whose context the scheduler service is running. Otherwise, a user with fewer privileges might change or replace the file that is launched by the scheduler service.

Post Snapshot Script

Optionally, specify the full path of a script to run after the snapshot is complete.

Backup Type

Select a backup type:

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the Exchange component
- Runs Eseutil to check for consistency before truncating the logs
- Truncates the transaction logs

Copy Backup is typically used for Quick Recovery. It creates a copy of the Exchange component, but does not truncate the transaction logs.

Copy Backup + Eseutil is used to create a copy and check the snapshot for consistency.

13 In the Edit Snapshot Schedule dialog box, on the General Options tab, you can specify the following:

Start Time

The time of the day to begin taking snapshots

End Time

The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.

Restart task every

The interval between snapshots, in minutes.

For example, if the interval is 360 minutes and you schedule a snapshot start time of 12 P.M. and an end time of 7 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.

Schedule takes effect on

The date on which the specified schedule takes effect. If you specify no other run day information on the other tabs, this date is used as the only run date.

If you specify run days on the other tabs, the schedule starts running on the earliest day specified. Therefore, if you want the current date included, specify it as an Include date on the Specific Dates tab.

Schedule recurs every	Enable this option to have the snapshot schedule continue to recur. Otherwise the schedule applies only for one day. Specify the number of days before restarting the snapshot schedule. For example, 1 day would mean the schedule takes effect daily, 2 days would mean every other day.
Start On	If you specify the Every option, specify the starting date.

- 14 In the Edit Snapshot Schedule dialog box, optionally make selections on the Week Days, Days of Month, and Specific Dates tabs as follows:

Week Days	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking 1st schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also specify the last day of the month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules for the same snapshot set overlap for the same snapshot, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

- 15 When you are done making selections on the Edit Snapshot Schedule dialog box for this schedule, click **OK**.
- 16 In the Snapshot Schedule panel, choose from the following and when scheduling is complete for all snapshot sets, click **Next**:
- Edit schedules for any remaining snapshot sets
 - Click **Add Schedules** if you want to add a new schedule

- Click **Delete** if you want to remove a schedule
- 17 In the Summary panel, choose one or both of the following:
 - If you want to save the settings to a template file for reuse, click **Save** and save the template to a file location of your choice.
If you do not save the settings to a file, you can still view or modify them by launching the wizard and selecting the same instance and Exchange component.
 - If you are ready to implement the template with its current settings, click **Apply**.
If you click **Apply** without saving the template, you are prompted to confirm.
If you have saved the settings to a template file and want to exit the wizard without applying the template, click **Cancel**.
 - 18 In the Template Implementation panel, wait until the wizard shows that Quick Recovery configuration is complete and click **Next**.
 - 19 Click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot. The scheduled snapshots are listed in the pane on the right.
If you have just configured the schedules and they are not yet displayed, right-click the Storage Agent node and click **Refresh** to update the display.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click the selected application and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.

- 6 In the Job History dialog box, view the schedule information. You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.

Situation

Resolution

You want to use the VSS Scheduler wizard but the Applications node is not shown in the VEA.

You may need to refresh the VEA display to see the node.

You are unable to locate a snapshot set XML file.

The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Drive letters assignments for snapshot volumes do not occur.

If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.

A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.

When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails.

To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.

Situation	Resolution
Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.	You can choose from the following ways to avoid these problems: <ul style="list-style-type: none"><li data-bbox="644 343 1115 395">■ Include both databases in the snapshot when running the wizard the first time.<li data-bbox="644 404 1158 456">■ Select a different disk for the second snapshot mirror when running the wizard the second time.<li data-bbox="644 465 1129 539">■ Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.
The following problems can occur:	
<ul style="list-style-type: none"><li data-bbox="329 618 629 791">■ When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.<li data-bbox="329 800 629 848">■ Mirror preparation for the second database fails.	

In addition, when troubleshooting, you may want to review the following logs.

If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting or modifying schedules

You can delete or modify a schedule that you created with the Quick Recovery Configuration Wizard by running the wizard again and deleting or editing the schedule on the Snapshot Schedule panel.

You cannot use the Quick Recovery Configuration Wizard to delete or modify schedules created from the VEA console.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

To delete a schedule from the VEA

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the **System** node and **StorageAgent** node. Select and click to expand the **Scheduled Tasks > Exchange** node.
The scheduled snapshots are listed on the right pane.
- 3 Right-click the name of the snapshot schedule and select **Delete Schedule**.

Synchronizing schedules after adding a cluster node

In a cluster environment, you may add a cluster node after you set up snapshot schedules with the Quick Recovery Configuration Wizard.

In such a case, you can ensure that the schedules are available on the new node by running the Quick Recovery Configuration Wizard again.

To synchronize schedules after adding a node

- 1 Start the Quick Recovery Configuration Wizard from the Solutions Configuration Center.
- 2 Continue through the wizard until the Synchronizing Schedules panel shows that synchronization between cluster nodes is complete.
- 3 Click **Cancel** to exit the wizard.

Scheduling or creating an individual snapshot set

This chapter includes the following topics:

- [About scheduling or creating an individual snapshot set](#)
- [Tasks to schedule a new snapshot](#)
- [Tasks to create a one-time snapshot set](#)
- [Reviewing the prerequisites](#)
- [Preparing the snapshot mirrors](#)
- [Scheduling a snapshot set](#)
- [Troubleshooting scheduled snapshots](#)
- [Creating a one-time snapshot set](#)
- [Refreshing a snapshot set manually](#)

About scheduling or creating an individual snapshot set

Typically you set up your initial snapshot schedules with the Quick Recovery Configuration Wizard.

See [Chapter 6, “Implementing snapshot sets with the configuration wizard”](#) on page 61.

However, later, you may want to schedule an additional snapshot. For example, you configured a daily snapshot for Quick Recovery use and now want to add a weekly snapshot on a different disk for backup use.

You can use the VSS Exchange Snapshot Scheduler Wizard to add a snapshot schedule for a selected mailbox database. Like the Quick Recovery Configuration

Wizard, the scheduler wizard enables you to automate the refreshing of snapshots according to the schedule that you define.

However, unlike the Quick Recovery Configuration Wizard, the VSS Exchange Snapshot Scheduler Wizard does not prepare snapshot mirrors. You must use the Prepare command to prepare the snapshot mirrors before running the VSS Exchange Snapshot Scheduler Wizard for that mailbox database. In addition, you can use the scheduler wizard to schedule only one snapshot set for one mailbox database at a time.

Note: Adding a snapshot set using the VSS Exchange Snapshot Scheduler will not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

See “[Scheduling a snapshot set](#)” on page 85.

At times you may want to create a one-time snapshot set. You can do so using either the vxsnap command line utility or from the VEA console using the VSS Exchange Snapshot and Snapback wizards.

See “[Creating a one-time snapshot set](#)” on page 97.

Note: To snapshot only a single volume rather than multiple volumes, you can use the VEA Snapshot Volume Wizard. See the *Veritas Storage Foundation Administrator's Guide*.

Tasks to schedule a new snapshot

[Table 7-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 7-1 Tasks for scheduling a new snapshot set

Objective	Tasks
“ Reviewing the prerequisites ” on page 83	■ Verifying hardware and software prerequisites
“ Preparing the snapshot mirrors ” on page 84	■ Creating snapshot mirrors using the using the VEA Prepare command or the vxsnap utility
“ Scheduling a snapshot set ” on page 85	■ Using the VSS Exchange Snapshot Scheduler Wizard to create the initial snapshot and set up the schedule for keeping it refreshed.

Tasks to create a one-time snapshot set

[Table 7-2](#) outlines the high-level objectives and the tasks to complete each objective.

Table 7-2 Tasks for creating a one-time snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 83	■ Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 84	■ Creating snapshot mirrors using the VEA Prepare command or the vxsnap utility
“Creating a one-time snapshot set” on page 97	■ Creating the one-time snapshot set using the VEA or the vxsnap utility.
“Refreshing a snapshot set manually” on page 101	■ Refreshing the one-time snapshot set if necessary.

Reviewing the prerequisites

Before implementing an individual snapshot set, you should do the following:

- Ensure that your system hardware and software meets the requirements. See [“System requirements”](#) on page 35 in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).
- Set up your Exchange mailbox databases for use with Storage Foundation for Windows. See [Chapter 5, “Configuring Exchange for Quick Recovery snapshots”](#) on page 51.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- The vxsnap commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Preparing the snapshot mirrors

To prepare the snapshot mirrors in order to create a snapshot from the VEA or from the `vxsnap` command line, you can use either of the following methods:

- The Prepare command from the VEA
You repeat the VEA console Prepare operation for each volume in the database.
- The `vxsnap prepare` command from the CLI
Use the `vxsnap prepare` command to prepare a mirror for each of the volumes associated with the databases or transaction logs of your Exchange database.

The snapshot mirrors remain attached to the original volumes and continue to be updated until you use the VSS Exchange Snapshot Wizard, the `vxsnap create` command, or the VSS Exchange Snapshot Scheduler Wizard to create the snapshot set.

For the snapshot volumes, make sure to select disks or LUNs that are not used for production data. You can create more than one snapshot volume on the same disk as long as there is sufficient space available and as long as the snapshots are of different production volumes.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare** option from the context menu.
- 2 Choose one of the following.

If the volume is not mirrored

Choose **Manually select disks**, use the **Add** and **Remove** buttons to move the desired disks to the Selected disks box, and click **OK**.

If the volume is mirrored and no additional disks are available to create a new mirror

Click on an existing plex and click **OK**.

If the volume is mirrored and there are additional disks available on your system

Choose either to use an existing mirror for the snapshot or to create a new mirror.

- To create a new mirror, click **Select Disk**, use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box, and click **OK**.
- To use an existing mirror, click **Select existing mirror for snap**, select the desired mirror, and click **OK**.

- 3 Repeat [step 1](#) and [step 2](#) to create a snapshot mirror for each volume associated with the Exchange databases and transaction logs. Verify that the lower pane of the VEA console indicates that the resynchronization process is complete before creating the snapshot set.

To create the snapshot mirrors using the command line

Type the command, as in the following example:

```
vxsnap prepare component=DB1/writer="Microsoft Exchange Writer"  
source=G:/harddisk=harddisk4  
source=H:/harddisk=harddisk4  
source=I:/harddisk=harddisk5
```

In this example, snapshot mirrors for the database volumes mounted at G: and H: are created on disk 4 and a snapshot mirror for the log volume mounted at I: is created on disk 5.

Note: For Exchange 2010, if replication is enabled and you want to create a snapshot mirror for the passive copy, specify the writer option as "Microsoft Exchange Replica Writer" or "Microsoft Exchange Writer Replica".

The complete syntax of the `vxsnap prepare` command is:

```
vxsnap prepare component=<componentName>/writer=<writerName>  
[-b] [source=<volume>/harddisk=<harddisk,...>] ...]
```

See "[vxsnap prepare](#)" on page 156.

Scheduling a snapshot set

Before you run the VSS Exchange Snapshot Scheduler Wizard to schedule a snapshot set for a database, you must prepare snapshot mirrors attached to the volumes in the database.

See "[Preparing the snapshot mirrors](#)" on page 84.

You can then use the VSS Exchange Snapshot Scheduler Wizard to schedule the initial snapshot and to set up the schedule for keeping it refreshed.

When the scheduled snapshot occurs, the snapshot mirrors are detached from the original volumes, creating separate on-host snapshot volumes as well as an XML file to store the Exchange and snapshot volume metadata. The scheduled process integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes in the database. This snapshot is done while the databases are online and without disrupting the email flow.

Note that if an Exchange 2010 mailbox database is configured under VCS, then you can store the snapshot set file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over.

In the case of a scheduled snapshot, even though schedules are visible on all nodes, the snapshot will happen only on a node where the mailbox database is present. There can be a large number of mailbox databases in Exchange 2010 and each database may have independent schedules. Hence the number of different paths required for storing the snapshot metadata files is higher, which is achieved by configuring a file share resource. You can also specify the file share path in the `redirect.txt` file, for example,

```
\\MySystemName1\share1\QRDir.
```

In the case of a VCS clustered setup with more than one node, configure the Veritas Scheduler Service with any user account with administrative privileges (other than the Local System account) which is valid on all the nodes of the cluster. This user account should have read-write permissions to the file share path.

For Exchange 2010, if replication is enabled in a DAG, you can choose to schedule a snapshot set for either the active or passive copy.

Note: Before adding or removing a node from a VCS cluster setup, delete the schedules and then recreate the schedules on the required node.

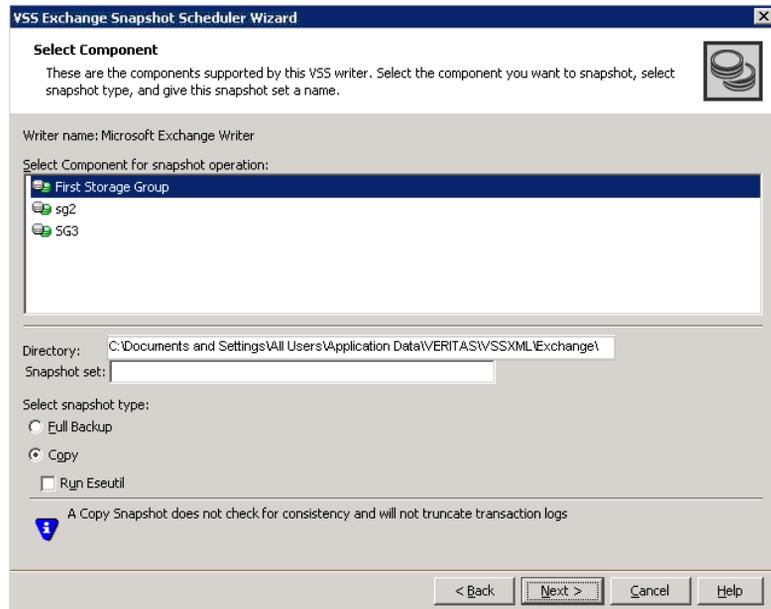
Note: For a VCS cluster setup if a schedule is created with snapshot set files stored on a file share path, then it is visible on all nodes of the cluster and can be deleted from any node irrespective of where the Exchange mailbox database component is online.

To schedule a snapshot

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.

- 3 Choose one of the following:
 - Right-click **Exchange** and select **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a schedule on the replication node, right-click **Exchange Replica** and click **Schedule VSS Exchange Snapshot**.
- 4 In the Welcome panel, review the information and click **Next**.

- 5 Specify the snapshot set parameters as follows and then click **Next**:



Select component for
snapshot operation

Select the component for the snapshot set.

For Exchange 2010, the component is a mailbox database.

Directory

Accept or enter a directory location for the XML file.

The XML file is stored by default in the directory shown on the screen.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore will fail.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, either edit the directory path in the Directory field for this wizard screen or use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the full path to the location of the XML file, for example, `G:\BackupSets`. Save the `redirect.txt` file in the default directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML` on each node of the cluster.

Note: You must not use the volume name or volume path in the `redirect.txt` file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.

Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. You can also specify the file share path in the `redirect.txt` file.

Snapshot set

Enter a name for the snapshot set.

Select snapshot type

Select the snapshot type.

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the selected component
- Runs **Eseutil** to check for consistency before truncating the logs
- Truncates the transaction logs

Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally, check **Run Eseutil** with the **Copy** option to check the snapshot for consistency.

You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database.

- 6 In the Change Attributes panel, optionally change the attributes for the snapshot volumes and click **Next**:

Snapshot Volume Label

Displays the read-only label for the snapshot volume.

Drive Letter

Optionally, click a drive letter and select a new choice from the drop-down menu.

The drive letters specified may not be available when the snapshot is taken. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

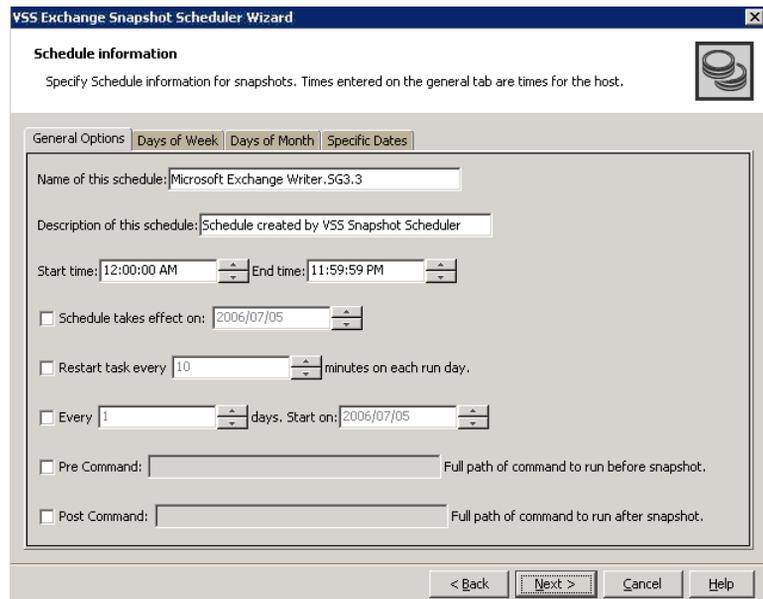
Plex

Optionally, click a plex and select a new choice from the drop-down menu.

- 7 Optionally, in the Synchronized Snapshot panel (VVR only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.
This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

See *Veritas Volume Replicator Administrator's Guide*.

- 8 In the Schedule Information panel, on the General Options tab, specify the following:



Name of this schedule	Enter a unique name for the snapshot set schedule. This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the VSS writer name, the component name and a numbered suffix that increments with each schedule.
Description of this schedule	Optionally, enter a description to help you identify the schedule when you view information about the snapshot status
Start Time	The time of the day to begin taking snapshots.
End Time	The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.
Schedule takes effect on	The date on which the specified schedule takes effect. The default is the current date.

Restart task every	The interval between snapshots, in minutes. For example, if the interval is 360 minutes and you schedule a snapshot start time of 12 P.M. and an end time of 7 P.M., the snapshot occurs twice.
Every	Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day. Specify the number of days before restarting the snapshot schedule. For example, 1 day would mean the schedule takes effect daily, 2 days would mean every other day.
Start On	If you enable the Every option, specify the starting date.
Pre Command	Optionally, specify the full path of a command script to run before the scheduled snapshot occurs. Note: Precommands which launch windows or any other GUI related entities are not allowed in the Scheduler.
Post Command	Optionally, specify the full path of a command script to run after the snapshot is complete. Note: Postcommands which launch windows or any other GUI related entities are not allowed in the Scheduler.

9 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day checkbox to schedule the snapshot for the last day of each month.

Specific Dates

Select one or more specific dates to include in or to exclude from the schedule.

Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

10 Click **Next**.

11 Review the snapshot set and schedule details and click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot. The scheduled snapshots are listed in the pane on the right.
If you have just configured the schedules and they are not yet displayed, right-click the Storage Agent node and click **Refresh** to update the display.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click the selected application and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the Job History dialog box, view the schedule information.
You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.
You want to use the VSS Scheduler wizard but the Applications node is not shown in the VEA.	You may need to refresh the VEA display to see the node.
You are unable to locate a snapshot set XML file.	The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Situation

Drive letters assignments for snapshot volumes do not occur.

A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.

Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.

The following problems can occur:

- When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.
- Mirror preparation for the second database fails.

Resolution

If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.

When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails.

To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.

You can choose from the following ways to avoid these problems:

- Include both databases in the snapshot when running the wizard the first time.
- Select a different disk for the second snapshot mirror when running the wizard the second time.
- Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.

In addition, when troubleshooting, you may want to review the following logs.

If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting schedules

You can delete (but not modify) a schedule from the VEA console. To modify a schedule, run the wizard again and select the same instance and component.

Note: Deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

To delete a schedule from the VEA

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the **System** node and **StorageAgent** node. Select and click to expand the **Scheduled Tasks > Exchange** node.
The scheduled snapshots are listed on the right pane.
- 3 Right-click the name of the snapshot schedule and select **Delete Schedule**.

Creating a one-time snapshot set

Creating a one-time snapshot set is a two-step process:

- The first step is to prepare snapshot mirrors attached to all the original volumes in the specified database. If you are creating a snapshot set after a snapback to refresh existing snapshot mirrors, you can skip this step. See [“Preparing the snapshot mirrors”](#) on page 84.
- The second step uses either the `vxsnap create` command or the VSS Exchange Snapshot Wizard to create the snapshot set by detaching the snapshot mirrors from the original volumes. This step creates separate on-host snapshot volumes as well as an XML file to store the Exchange and snapshot volume metadata.

The `vxsnap create` command and VSS Exchange Snapshot Wizard integrate with VSS to quiesce the databases and then simultaneously snapshot the volumes in the database. This snapshot is done while the databases are online and without disrupting the email flow. The resulting snapshot set provides a complete picture of the database at the point in time the command is issued.

For Exchange 2010, if replication is enabled, you have the option to create a snapshot set of the passive copy instead of the active database. For additional information on the use of passive copy snapshots in a DAG, see:

- [“Exchange 2010 Database Availability Group \(DAG\) considerations”](#) on page 47 in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#)

The VSS Exchange Snapshot Wizard can be run from either a local system or a remote node. The `vxsnap` utility must be run from the local system.

Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over.

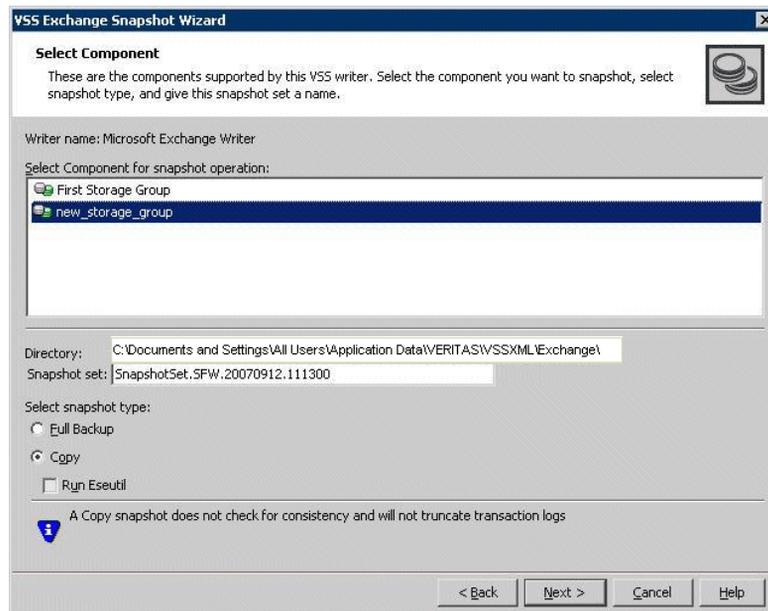
You can also specify the file share path in place of the `redirect.txt` path, for example, `\\MySystemName1\share1\QRDir`. There can be a large number of mailbox databases in Exchange 2010 and each database may have independent schedules. Hence the number of different paths required for storing the snapshot metadata files is higher, which can be achieved by configuring a file share resource.

See [“Scheduling a snapshot set”](#) on page 85.

To create the snapshot set from the VEA console

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.

- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a snapshot on the replication node, right-click **Exchange Replica** and click **VSS Exchange Snapshot**.
- 4 In the wizard, review the Welcome page and click **Next**.
- 5 Specify the snapshot set parameters as follows and then click **Next**:



Select Component for snapshot operation

Select the component for the snapshot set.
For Exchange 2010, the component is a mailbox database.

Directory

Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory shown on the screen.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore will fail.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, either edit the directory path in the Directory field for this wizard screen or use a text editor to create a text file named "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, G:\BackupSets. Save the **redirect.txt** file in the default directory C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML on each node of the cluster.

Note: You must not use the volume name or volume path in the `redirect.txt` file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.

Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. You can also specify the file share path in the `redirect.txt` file.

Snapshot set

Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.

Select snapshot type

Select the snapshot type.

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the selected component
- Runs **Eseutil** to check for consistency before truncating the logs
- Truncates the transaction logs

Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally check **Run Eseutil** with the **Copy** option to check the snapshot for consistency.

- 6 In the Change Attributes panel, optionally change the attributes for the snapshot volumes as follows and click **Next**:

Snapshot Volume Label

Displays the read-only label for the snapshot volume.

Drive Letter

Optionally, click a drive letter and select a new choice from the drop-down menu.

Plex

Optionally, click a plex and select a new choice from the drop-down menu.

- 7 Optionally, in the Synchronized Snapshot panel, select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

To create the snapshot set from the command line

- Type the command, as in the following example:

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange Writer"  
component=DB1 backupType=COPY -E -O
```

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]  
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<n  
ewVolName>] [/Plex=<plexName>] . . . writer=<writerName>  
component=<componentName> [backuptype=<backuptype>] [-E] [-O] [-C] [  
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Creating the snapshot set with the copy backup type does not automatically run the Eseutil consistency check nor truncate the transaction logs. You can check the snapshot for consistency by specifying the -E flag.

Note: Any text string that contains spaces must be enclosed in quotation marks.

See “[vxsnap create](#)” on page 157.

Refreshing a snapshot set manually

Once a snapshot set has been created, it can be refreshed quickly since the time-consuming step of preparing the mirrors is not required.

Normally, if you want to periodically refresh a snapshot set, you set up the snapshot schedule using the VSS Exchange Snapshot Scheduler Wizard or the Quick Recovery Configuration Wizard.

However, if you should need to manually refresh a snapshot set, you can do so. To refresh the snapshot set requires the following tasks:

- “[Reattaching the split-mirror snapshots](#)” on page 101
- “[Creating the refreshed snapshot set](#)” on page 103

Note: The VSS Refresh option available in the VEA console from the Microsoft Exchange Writer object refreshes the display of the VSS Writer and components. It does not refresh the snapshot set.

Reattaching the split-mirror snapshots

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot set so that it matches the current state of its original Exchange mailbox database. The wizard is available in the context menu of the VSS Writer object.

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot set.
- 2 From the VEA console URL bar, select the *<host name>* which is the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click on the node of the application and click **VSS Snapback**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.
The XML metadata file contains all required information needed to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.
- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click Yes to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command, as in the following example:

```
vxsnap -x Image1.xml reattach writer="Microsoft Exchange Writer"
```

The complete syntax for the `vxsnap reattach` command is:

```
vxsnap -x <filename> [-f][-b] reattach writer=<writername>  
[secHosts=<secondary hosts>]
```

See "[vxsnap reattach](#)" on page 161.

Creating the refreshed snapshot set

Create a new snapshot set of the database using either the VSS Exchange Snapshot Wizard or the `vxsnap create` command.

See “[Creating a one-time snapshot set](#)” on page 97.

Recovering Exchange mailbox databases

This chapter includes the following topics:

- [About recovery using Quick Recovery snapshots](#)
- [Tasks for recovery using Quick Recovery snapshots](#)
- [Prerequisites for recovery](#)
- [Recovery using an Exchange 2010 passive copy snapshot in a Database Availability Group \(DAG\)](#)
- [Recovering using the VSS Restore Wizard](#)
- [Recovering using the vxsnap utility](#)
- [Post-recovery steps](#)

About recovery using Quick Recovery snapshots

The on-host snapshot set can be used to quickly recover a database after logical corruption. You can restore a database either to the Point in Time (PIT) that the `vxsnap restore` command was last refreshed or, using current transaction logs, to the Point of Failure (POF) of the database.

You can use either the VSS Restore Wizard from the VEA console or the `vxsnap restore` command. Both the VSS Restore wizard and `vxsnap restore` command integrate with VSS to notify the Exchange VSS Writer to prepare for the restore before the snapback operation and then to complete post-restore processes afterwards.

After completing the recovery, you refresh the snapshot set.

Tasks for recovery using Quick Recovery snapshots

[Table 8-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 8-1 Tasks for recovery using Quick Recovery snapshots

Objective	Tasks
“Prerequisites for recovery” on page 107	■ Verifying hardware and software prerequisites
“Recovering using the VSS Restore Wizard” on page 108	■ Running the VSS Restore Wizard from the VEA console to restore a database to Point in Time (PIT) or restore a database Point of Failure (POF)
“Recovering using the vxsnap utility” on page 112	■ Using the vxsnap utility to restore a database to point in time or restore a database point of failure
“Post-recovery steps” on page 116	■ Refreshing the snapshot set ■ Performing additional tasks in a VVR environment

Prerequisites for recovery

You can use the VSS Restore Wizard or vxsnap utility to recover a single database in cases where a snapshot set is available and where there has been no hardware failure.

Caution: Do not attempt to use the VSS Restore Wizard for recovery after hardware failure. Instead, refer to the special procedures for recovery after hardware failure.

See [Chapter 9, “Recovering after hardware failure”](#) on page 117.

Recovery using an Exchange 2010 passive copy snapshot in a Database Availability Group (DAG)

For Exchange 2010 mailbox databases, if you used SFW to create a snapshot of a passive copy of the database, the passive copy snapshot can be used to restore an Exchange mailbox database.

Refer to the Microsoft Exchange documentation for additional information on Microsoft Exchange Management Shell and cmdlets.

To recover using an Exchange 2010 passive copy snapshot, you follow this process:

An example of a PIT recovery procedure from a snapshot of a passive replica copy of an Exchange mailbox database **MB1** on an active mailbox database **TestServer1** and replica copy on **TestServer2** would be as follows:

- Take a snapshot of the replica copy of the mailbox database on **TestServer2**.
- Use the `Move-ActiveMailboxDatabase` cmdlet to perform the active mailbox database **MB1** switchover from **Test Server 1** to **Test Server 2**.
`Move-ActiveMailboxDatabase MB1 -ActivateOnServer TestServer2`
Now **TestServer2** has the active mailbox database and **TestServer1** has the replica copy.
- Use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).
`Suspend-MailboxDatabaseCopy -Identity MB1\TestServer1 [-Confirm:$False]`
Use `-Confirm` to prevent user intervention.
- Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for “**This database can be overwritten by Restore**”.

```
set-mailboxdatabase -identity MB1\TestServer2  
-allowfilerestore:$true
```

- Use the `dismount-Database` cmdlet to dismount the database

```
dismount-Database -Identity MB1\TestServer2 [-Confirm:$False]
```

Use `-Confirm` to prevent user intervention.

For a VCS cluster setup, offline the VCS resource for the Exchange 2010 database manually if automatic offlining fails.

Specifying the `-a` option with the `vxsnap restore` utility dismounts the database automatically. For a cluster setup, it offlines the resource automatically.

- Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.

```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft  
Exchange Writer"
```

For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer  
Replica" component=MB1 backupType=COPY -E -O
```

- Use the `mount-Database` cmdlet to mount the database.

```
mount-Database -Identity MB1\TestServer2
```

For a VCS cluster setup, online the VCS resource for the Exchange 2010 database manually if automatic onlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` utility mounts the database automatically. For a cluster setup, it onlines the resource automatically.

- Use the `Update-MailboxDatabaseCopy` cmdlet.

```
Update-MailboxDatabaseCopy -Identity MB1\TestServer2  
-SourceServer TestServer1 -DeleteExistingFiles [-Confirm:$False]
```

Use `-Confirm` to prevent user intervention.

Recovering using the VSS Restore Wizard

Using the VSS Restore Wizard, you can do the following:

- Restore a database to the point in time of the snapshot set
- Perform a roll-forward recovery of a database to the point of failure

Note: Recovery Database for Exchange 2010 is not supported by SFW.

Before you begin, review the prerequisites.

See “[Prerequisites for recovery](#)” on page 107.

SFW automatically dismounts the stores before beginning the restore operation. If it fails to dismount the stores, the restore operation fails. In that case, manually dismount the stores. In the case of a VCS cluster setup, it offlines and onlines the resource for the database automatically. Manual intervention is required if the resource fails to offline automatically.

The VSS Restore Wizard helps you to restore a VSS snapshot of the mailbox database components using the snapshot set files. The mailbox database can be restored with the restore wizard. Before you proceed with the restore operation, ensure that you have prepared and created a snapshot set with VSS Snapshot Wizard or the vxsnap utility.

After completing the POF restore operation on a snapshot set file, you can refresh it by using either the VSS Snapback wizard or vxsnap reattach command to snapback any remaining database volumes. Then, you can use the VSS Exchange Snapshot Wizard or vxsnap create command to create a new snapshot set of the mailbox database.

For an Exchange 2010 database configured under VCS cluster, offline the corresponding VCS resource for the database manually before the restore operation, if the resource is not offlined automatically. Once the restore operation completes successfully, online the VCS resource. Specifying the `-a` option in the vxsnap restore command offlines and onlines the VCS resource automatically.

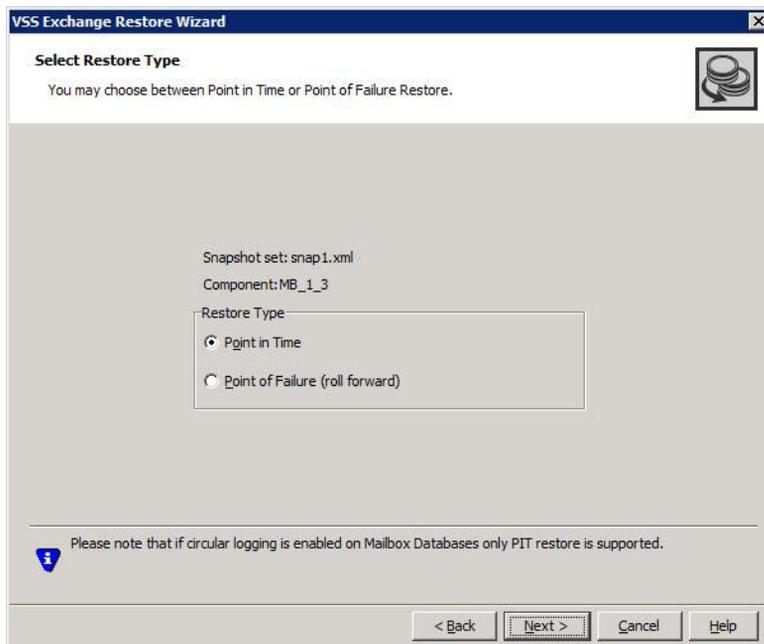
See “[Recovering using the vxsnap utility](#)” on page 112.

To restore Point in Time (PIT) snapshot set (Exchange 2010)

- 1 Click **Start > Exchange Management Console**. Expand the **Organization Configuration** node and select **Mailbox**.
- 2 If replication is enabled for the selected database, you should first suspend continuous replication on this database. In the result pane, on the Database Management tab, select the database that you want to suspend.
- 3 In the work pane, on the Database Copies tab, right-click the database for which you want to suspend continuous replication. Select **Suspend Database Copy** from the context menu. Click **Yes** to suspend continuous replication.
- 4 From the work pane of the snap-in, right-click on the mailbox database and select **Properties > Maintenance**. Enable the checkbox **This database can be overwritten by a restore option**.
- 5 Now proceed with the VSS restore operations on the selected mailbox database.

Close the database application GUI, all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot volumes.

- 6 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 7 Expand the system icon and the **Applications** node.
- 8 Right-click **Exchange** and select **VSS Restore** from the context menu.
- 9 Review the Welcome page. Click **Next**.
- 10 Select the snapshot set you wish to restore and click **Next**.
The XML metadata file contains all required information needed to restore the snapshot set, including the names of the database and transaction log volumes.
Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 11 In the Select Restore Type panel, select **Point in Time**.
Point in Time restores to the point-in-time that particular snapshot set was created or refreshed based on the information contained in the metadata file.



Click **Next**.

- 12 Verify that the restore specifications are correct on the summary page and Click **Finish** to exit the wizard.
- 13 If you have suspended replication for your clustered mailbox database as shown in [step 2](#), you should now update it from the work pane of the Exchange Management Console. In the work pane, on the Database Copies tab, right-click the mailbox database and select **Update Database Copy**.

If you want to do a roll forward or point of failure (POF) recovery, refer to the steps shown below. A roll forward recovery restores the mailbox databases to the point of failure.

Note: Do not enable the circular logging option. If circular logging is enabled, you will not be able to roll forward a database to achieve a point-of-failure recovery. Microsoft does not recommend circular logging in production environments. Refer to the Microsoft Exchange Server documentation for further details.

To do a roll forward recovery to the Point of Failure (POF) of mailbox databases, perform the following:

To perform a roll forward recovery the point-of-failure (POF) (Exchange 2010)

- 1 Click **Start > Exchange Management Console**. Expand the **Organization Configuration** node and select **Mailbox**.
- 2 If replication is enabled for the selected database, you should first suspend continuous replication on this database. In the result pane, on the Database Management tab, select the database that you want to suspend.
- 3 In the work pane, on the Database Copies tab, right-click the database for which you want to suspend continuous replication. Select **Suspend Database Copy** from the context menu. Click **Yes** to suspend continuous replication.
- 4 From the work pane of the snap-in, right-click on the mailbox database and select **Properties > Maintenance**. Enable the checkbox for **This database can be overwritten by a restore option**.
- 5 Now proceed with the VSS restore operations on the selected mailbox database.
Close the database application GUI, all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot volumes.

- 6 From the VEA console, navigate to the system where the production volumes and snapshot set files are located.
- 7 Expand the system icon and the **Applications** node.
- 8 Right-click **Exchange** and select **VSS Restore** from the context menu.
- 9 Review the information on the Welcome page and click **Next**.
- 10 Select the name of the XML metadata file for the snapshot set you wish to restore and click **Next**.
The XML metadata file contains all required information about the Exchange mailbox database, including the names of the database and transaction log volumes.
Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 11 In the Select Restore Type Panel, select **Point of Failure (roll forward)**.
Roll forward restores the database to the most recent state before a failure occurred.
Click **Next**.
- 12 Verify the restore specifications on the summary page and click **Finish** to exit the wizard.
- 13 If you have suspended replication for your clustered mailbox database as shown in [step 2](#), you should now update it from the work pane of the Exchange Management Console. In the work pane, on the Database Copies tab, right-click the mailbox database and select **Update Database Copy**.

Recovering using the vxsnap utility

You can also recover exchange mailbox databases by using the vxsnap utility from the Command Line Interface (CLI).

Refer to Microsoft Exchange Shell cmdlets for details.

Note: SFW automatically dismounts the stores before beginning the restore operation. If it fails to dismount the stores, the restore operation fails. In that case, manually dismount the stores and repeat the restore operation. For a VCS cluster setup, it offlines and onlines the resource for the database automatically. If automatic offlining of the resource fails, then manually offline the resource.

Using the vxsnap utility, you can do the following:

- Restore a database to the point in time (PIT) of the snapshot set
- Perform a roll-forward recovery of a database to the point of failure (POF)

Before you begin, review the prerequisites.

See “[Prerequisites for recovery](#)” on page 107.

You use the `vxsnap restore` command. The complete syntax of the `vxsnap restore` command is:

```
vxsnap -x <filename>[-f][-b][-r] [-a] restore
restoreType=<PIT|POF>writer=<writename> [subComponent=<subComponentName>] [RSG=<Yes|No>]
```

Note: The `subComponent` and `RSG=YES|NO` options are not valid for Exchange 2010.

To restore a database to the point in time of the snapshot set using the `vxsnap restore` command

- If replication is enabled for a cluster setup, then use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).

```
Suspend-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
```

Use `-Confirm` to prevent user intervention.
- Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for “**This database can be overwritten by Restore**”.

```
set-mailboxdatabase -identity MB1 -allowfilerestore:$true
```
- Use the `dismount-Database` cmdlet to dismount the database.

```
dismount-Database -Identity MB1[-Confirm:$False]
```

Use `-Confirm` to prevent user intervention.

For a VCS cluster setup, offline the VCS resource for the database manually if automatic offlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` command dismounts the database automatically. In the case of a VCS cluster setup, it offlines the resource automatically.
- Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.

```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft Exchange Writer"
```

For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer Replica" component=MB1 backupType=COPY -E -O
```
- Use the `mount-Database` cmdlet to mount the database.

```
mount-Database -Identity MB1
```

For a VCS cluster setup, online the VCS resource for the database manually if automatic onlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` command utility mounts the database automatically. In the case of a VCS cluster setup, it onlines the resource for the database.

- 6 If you have suspended replication as shown in [step 1](#), then use the `Update-MailboxDatabaseCopy` cmdlet.


```
Update-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
Use -Confirm to prevent user intervention.
```

To perform a roll-forward recovery of a database to the point of failure using the `vxsnap restore` command

- 1 If replication is enabled for a cluster setup, then use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).


```
Suspend-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
Use -Confirm to prevent user intervention.
```
- 2 Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for **“This database can be overwritten by Restore”**.


```
set-mailboxdatabase -identity MB1 -allowfilerestore:$true
```
- 3 Use the `dismount-Database` cmdlet to dismount the database.


```
dismount-Database -Identity MB1[-Confirm:$False]
Use -Confirm to prevent user intervention.
```

For a VCS cluster setup, offline the VCS resource for the database manually if automatic offlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` command dismounts the database automatically. In the case of a VCS cluster setup, it offlines the resource.
- 4 Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.


```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft Exchange Writer"
```

For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer Replica" component=MB1 backupType=COPY -E -O
```
- 5 Use the `mount-Database` cmdlet to mount the database.


```
mount-Database -Identity MB1
```

For a VCS cluster setup, online the VCS resource for the database manually if automatic onlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` command utility mounts the database automatically. In the case of a VCS cluster setup, it onlines the resource.

- 6 If you have suspended replication as shown in [step 1](#), then use the

`Update-MailboxDatabaseCopy` cmdlet.

```
Update-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
```

Use `-Confirm` to prevent user intervention.

Post-recovery steps

After you have performed any of the recovery methods, whether point-in-time or roll forward, you should refresh the snapshot set, by performing a snapback to reattach the snapshot mirrors, and then create a new snapshot set.

See “[Refreshing a snapshot set manually](#)” on page 101 in [Chapter 7, “Scheduling or creating an individual snapshot set”](#).

In a VVR environment, there is an additional post-recovery step. During a point-in-time recovery in a VVR environment, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

To resynchronize the node at the secondary site with the node at the primary site:

- 1 Right-click on the primary RVG and click **Resynchronize Secondaries**.
- 2 Click **Yes** to resynchronize the nodes.

Recovering after hardware failure

This chapter includes the following topics:

- [About recovery after hardware failure](#)
- [Tasks for recovering after hardware failure](#)
- [Reviewing the prerequisites](#)
- [Reviewing the sample configuration](#)
- [Scenario I: Database and transaction logs volumes are missing](#)
- [Scenario II: Database volumes missing, transaction logs are available](#)
- [Refreshing the snapshot set](#)
- [Refreshing the snapshot set on the current disks](#)
- [Moving the production volumes to different disks and refreshing the snapshot set](#)

About recovery after hardware failure

A system failure that results in the loss of either database or transaction log volumes leads to unplanned downtime. If the failure does not affect the disk or disks where the snapshot set resides, you can recover from the failure and resume normal email operations faster by using the snapshot set than you could by restoring from your backup media.

You can use the `vxsnap` utility to recover a database after production volumes are lost due to a hardware failure.

Note: The methods described in this chapter are the simplest possible and assumes that a volume snapshot set is already available before proceeding with the restore operation. However, if the snapshot set is not available, you will need to recover it from your tape backup.

Recovery scenarios

[Table 9-1](#) lists the database recovery scenarios that you choose from if the complete snapshot set (including the XML metadata file) is available:

Table 9-1 Available recovery type after missing production volume

Scenario	Database Volumes	Transaction Logs Volume	Recovery
Scenario I	One or more volumes are missing	Missing	point in time
Scenario II	All volumes are missing	Available	point in time or point of failure

Caution: If you choose to restore the database to the point in time, you cannot later restore it to the point of failure. You can only perform one recovery procedure on a mailbox database.

Recovery tasks

Perform the recovery tasks in the order shown below.

For a VSS-integrated recovery, you need the following information:

- Prepare for the recovery
 - Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
 - Delete the missing volumes from Storage Foundation for Windows.
 - Replace the failed hardware and add the new disks to the dynamic disk group.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
- Use the `vxsnap restore` command to recover the databases.
- Refresh the snapshot set.

Tasks for recovering after hardware failure

Table 9-2 outlines the high-level objectives and the tasks to complete each objective.

Table 9-2 Tasks for recovering after hardware failure

Objective	Tasks
“Reviewing the prerequisites” on page 120	<ul style="list-style-type: none"> ■ Reviewing the prerequisites for a recovery after system failure
“Scenario I: Database and transaction logs volumes are missing” on page 121	<ul style="list-style-type: none"> ■ Preparing for the recovery ■ Changing drive letters or mount points of all the snapshot volumes ■ Restoring the database to the point in time ■ Refreshing the snapshot set
“Scenario II: Database volumes missing, transaction logs are available” on page 128	<ul style="list-style-type: none"> ■ Preparing for the recovery ■ Removing the drive letter or mount point of the transaction logs volume ■ Changing drive letters or mount points of all the snapshot database volumes ■ Restoring the database to the point in time or Recovering the database to the point of failure ■ Refreshing the snapshot set
“Refreshing the snapshot set” on page 137	<ul style="list-style-type: none"> ■ Understanding your current configuration ■ Choosing the disks for the refreshed snapshot sets.
“Refreshing the snapshot set on the current disks” on page 137	<ul style="list-style-type: none"> ■ Reattaching healthy snapshot volumes ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Creating the refreshed snapshot set
“Moving the production volumes to different disks and refreshing the snapshot set” on page 144	<ul style="list-style-type: none"> ■ Reattaching healthy snapshot volumes ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Adding mirrors to volumes whose drive letters or mount points were reassigned ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Creating the refreshed snapshot set

Reviewing the prerequisites

Ensure that you have fulfilled the prerequisites cited below before proceeding with the restore operation.

Observe the following prerequisites:

- Make sure to correct the hardware or software issues and repair the failed volumes prior to using the procedures described below. Note that the `vxsnap restore -r` recovery method should be used only after other options have proved unsuccessful.
 See the Troubleshooting section of the *Veritas Storage Foundation Administrator's Guide* for more information.
- The complete snapshot set including the corresponding XML metadata file must be available before you perform the restore operation.

Reviewing the sample configuration

[Table 9-3](#) shows the objects that relate to an Exchange dynamic disk group, Exch-DG1, and are used to illustrate the tasks for hardware recovery.

Table 9-3 Objects for the Exchange mailbox database

Volume	Drive letter or mount point	Object
DB1-tlogs	I:\Logs	transaction logs volume
DB1	G:	mailbox database volume

The dynamic group Exch-DG1 has an associated snapshot set with a metadata file TestDB.xml and the volumes shown in [Table 9-4](#).

Table 9-4 Objects for mailbox DB1 snapshot set

Volume	Drive letter or mount point	Object
DB1-tlogssnap	not assigned	snapshot volume of DB1-tlogs
DB1snap	not assigned	snapshot volume of DB1

Scenario I: Database and transaction logs volumes are missing

If the metadata file is available but the transaction logs volume and database volume are missing, you can use the snapshot set to restore the mailbox database to the point in time that the snapshot set was created or last refreshed.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
See [“Identifying the missing volumes”](#) on page 122.
- Delete the missing volumes from Storage Foundation for Windows.
See [“Deleting missing volumes from Storage Foundation for Windows”](#) on page 123.
- Replace the failed hardware and add the new disks to the dynamic disk group.
See [“Replacing hardware and adding disks to the dynamic disk group”](#) on page 123.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
See [“Changing the drive letter or mount points of the snapshot volumes”](#) on page 124.
- Use the `vxsnap restore` to recover the databases.
See [“Restoring the mailbox database to the point in time”](#) on page 126.
- Refresh the snapshot set.
See [“Refreshing the snapshot set”](#) on page 127.

Identifying the missing volumes

Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. The drive letter or mount point information is available either through the VEA console or Exchange.

For example, if DB1-tlogs and DB1 are missing due to hardware failure, you would need the information in [Table 9-5](#).

Table 9-5 Information for missing volumes

Status	Volume	Drive letter or mount point	Object
Missing	DB1-tlogs	I:\Logs	transaction logs volume
Missing	DB1	G:	mailbox database volume
Available	DB1-tlogssnap	not assigned	snapshot volume of DB1-tlogs
Available	DB1snap	not assigned	snapshot volume of DB1

Deleting missing volumes from Storage Foundation for Windows

You must delete missing volumes from Storage Foundation for Windows. This makes the drive letter or mount point available and removes information about the missing volume that is no longer relevant.

Note: Before deleting the missing production volumes, note the volume name and drive letter or mount point. You will need this information later on in order to assign the same drive letter to the snapshot volume associated with each missing production volume.

To delete a missing volume using the VEA console

- 1 Right-click on the missing volume and select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete the volume.

To delete a missing volume using the command line

- ◆ Type the command as in the following example, which deletes the DB1volume:

```
> vxassist -gExch-DG1 delete DB1
```

For complete syntax of `vxassist delete` command, refer to the *Veritas Storage Foundation Administrator's Guide*.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk.

Note the new disk number(s). You will need the information later on to add the disks to the dynamic disk group and for the recovery operation.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 In the Actions menu, click **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).

or

If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console and is marked with (No Signature), for example, `harddisk5 (No signature)`. Proceed to the next step.

Scenario I: Database and transaction logs volumes are missing

- 4 Right-click on a new, unsigned disk and click **Write Signature**.
- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange mailbox database. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 In the Welcome panel, click **Next**.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Changing the drive letter or mount points of the snapshot volumes

If the production volume is missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

[Table 9-6](#) shows the changes that would be made in the sample configuration.

Table 9-6 Changes to make in the sample configuration

Volume	Drive letter or mount point	Object
DB1-tlogsnap	I:\Logs	snapshot volume of DB1-tlogs
DB1snap	G:	snapshot volume of DB1

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To change a snapshot volume drive letter to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.

- 2 Select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.

To change a snapshot volume mount point to a production volume drive letter

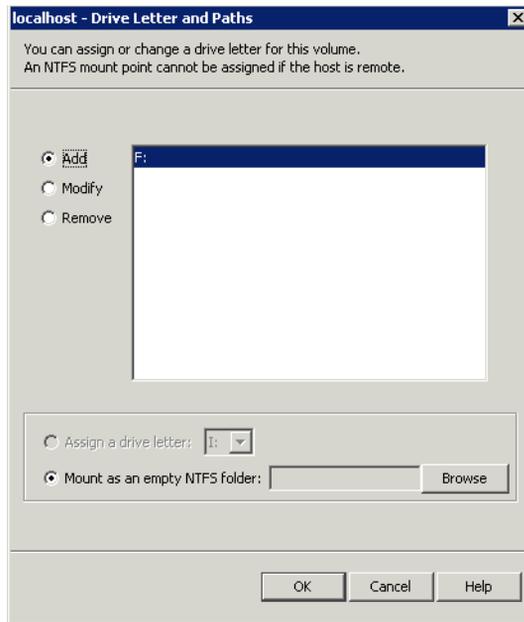
- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Click **Mount as an empty NTFS folder** and click **Browse**.

Scenario I: Database and transaction logs volumes are missing

- 8 Double-click the volume where the production volume was mounted.



- 9 Select the drive folder where the production volume was mounted.
- 10 Click **OK** to assign the mount point.

Restoring the mailbox database to the point in time

Restore the mailbox database to the point in time and mount the Exchange databases.

When restoring after hardware failure, you must use the command line and make sure to specify the `-r` option.

See “[vxsnap restore](#)” on page 162.

Caution: You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.

To restore the mailbox database to the point in time of the snapshot set

- 1 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.

- 2 If replication is enabled for a cluster setup, then use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).


```
Suspend-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
- 3 Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for **“This database can be overwritten by Restore”**.


```
set-mailboxdatabase -identity MB1 -allowfilerestore:$true
```
- 4 Use the `dismount-Database` cmdlet to dismount the database.


```
dismount-Database -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
 For a VCS cluster setup, offline the VCS resource for the database manually if automatic offlining of the resource fails.
 Specifying the `-a` option with the `vxsnap restore` command dismounts the database automatically. In the case of a VCS cluster setup, it offlines the resource for the database automatically.
- 5 Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.


```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft Exchange Writer"
```

 For this example, assume that the snapshot of the replica was performed with


```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer Replica" component=MB1 backupType=COPY -E -O
```
- 6 Use the `mount-Database` cmdlet to mount the database.


```
mount-Database -Identity MBI
```

 For a VCS cluster setup, online the VCS resource for the database manually if automatic onlining of the resource fails.
 Specifying the `-a` option with the `vxsnap restore` command utility mounts the database automatically. In the case of a VCS cluster setup, it onlines the resource for the database.
- 7 If you have suspended replication as shown in [step 2](#), then use the `Update-MailboxDatabaseCopy` cmdlet.


```
Update-MailboxDatabaseCopy -Identity MB1 -SourceServer -DeleteExistingFiles [-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.

Refreshing the snapshot set

After recovering from a hardware failure, refresh the snapshot set.

See [“Refreshing the snapshot set”](#) on page 137.

In the example, the DB1-tlogsnap and DB1snap volumes will reside on the disks where the original snapshot volumes were.

Scenario II: Database volumes missing, transaction logs are available

If all the database volumes are missing but the transaction logs volume and metadata file are available, you can either use the snapshot set to restore the mailbox database to the Point in Time (PIT) that the snapshot set was created or last refreshed or you can use the snapshot set to perform a roll-forward recovery to the Point of Failure (POF).

The tasks for either a point in time or point of failure recovery are the same except for the actual `vxsnap restore` command.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
See [“Identifying the missing volumes”](#) on page 129.
- Delete the missing volumes from Storage Foundation for Windows.
See [“Deleting missing volumes from Storage Foundation for Windows”](#) on page 129.
- Replace the failed hardware and add the new disks to the dynamic disk group.
See [“Replacing hardware and adding disks to the dynamic disk group”](#) on page 130.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
See [“Changing the drive letter or mount points of the snapshot volumes”](#) on page 131.
- Use the `vxsnap restore` to recover the databases.
See [“Restoring the mailbox database to the Point in Time \(PIT\)”](#) on page 133.
See [“Recovering the mailbox database to the Point of Failure \(POF\)”](#) on page 135.
- Refresh the snapshot set.
See [“Refreshing the snapshot set”](#) on page 137.

Identifying the missing volumes

Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. The drive letter or mount point information is available either through the VEA console or Exchange.

For example, if DB1 is missing due to hardware failure, you would need the information in [Table 9-7](#).

Table 9-7 Information to identify missing volumes

Status	Volume	Drive letter or mount point	Object
Available	DB1-tlogs	I:\Logs	transaction logs volume
Missing	DB1	G:	mailbox database volume
Available	DB1-tlogssnap	not assigned	snapshot volume of DB1-tlogs
Available	DB1snap	not assigned	snapshot volume of DB1

Deleting missing volumes from Storage Foundation for Windows

You must delete missing volumes from Storage Foundation for Windows. This will make the drive letter or mount point available and will remove information about the missing volume that is no longer relevant. For example delete the DB1 volume. Do not delete the DB1-tlogs volume.

Note: Before deleting the missing production volumes, note the volume name and drive letter or mount point. You will need this information later on in order to assign the same drive letter to the snapshot volume associated with each missing production volume.

To delete a missing volume using the VEA console

- 1 Right-click on the designated volume and select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete the volume.

To delete a missing volume using the command line

Type the command as in the following example, which deletes the DB1 volume:

```
vxassist -gExch-DG1 delete DB1
```

For the complete syntax of the `vxassist delete` command, see *Veritas Storage Foundation Administrator's Guide*.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk. Note the new disk number(s). You will need the information to add the disks to the dynamic disk group and for the recovery operation.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 From the **Actions** menu, select **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).
or
If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console and is marked with (No Signature), for example, `harddisk5 (No signature)`. Proceed to the next step.
- 4 Right-click on a new, unsigned disk and click **Write Signature**.
- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange mailbox database. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 Click **Next** at the welcome screen.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Changing the drive letter or mount points of the snapshot volumes

If the production volume is missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

Table 9-8 shows the changes that would be made in the sample configuration.

Table 9-8 Changing the drive letter or mount points

Volume	Drive letter or mount point	Object
DB1-tlogssnap	Do not change, production volume is healthy	snapshot volume of DB1-tlogs
DB1snap	G:	snapshot volume of DB1

Note: Assign drive letters to database snapshot volumes only. Do not change the transaction logs volume.

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To change a snapshot volume drive letter to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.
- 5 Repeat this procedure for all snapshot volumes of database volumes only.

To change a snapshot volume mount point to a production volume drive letter

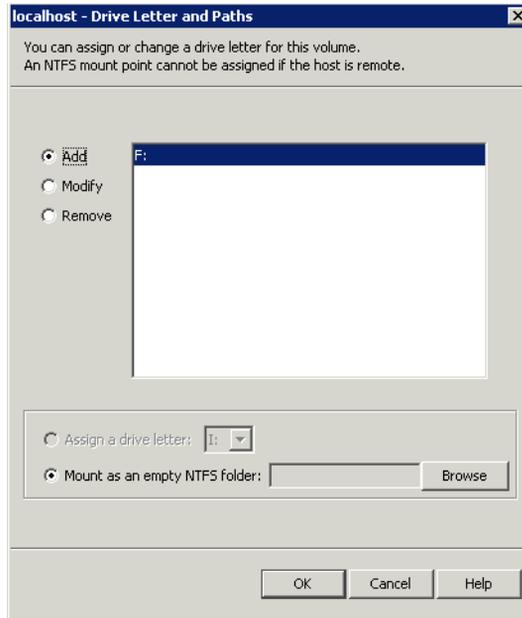
- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.

- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.
- 9 Repeat this procedure for all snapshot volumes of database volumes only.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Click **Mount as an empty NTFS folder** and click **Browse**.

- 8 Double-click the volume where the production volume was mounted.



- 9 Select the drive folder where the production volume was mounted.
- 10 Click **OK** to assign the mount point.
Repeat this procedure for all snapshot volumes of database volumes only.

Restoring the mailbox database to the Point in Time (PIT)

Restore the mailbox database to the point in time that the snapshot set was created or last refreshed and mount the Exchange databases.

Refer to Microsoft Exchange Shell cmdlets for details.

Note: When restoring after hardware failure, you must use the command line and make sure to specify the `-r` option.

Caution: If you choose to restore the mailbox database to the point in time, you cannot later restore it to the point of failure. You can only perform one recovery procedure on a database.

Caution: You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.

To restore the mailbox database to the point in time of the snapshot set

- 1 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.
- 2 If replication is enabled for a cluster setup, then use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).

```
Suspend-MailboxDatabaseCopy -Identity MB1 [-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
- 3 Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for **“This database can be overwritten by Restore”**.

```
set-mailboxdatabase -identity MB1-allowfilerestore:$true
```
- 4 Use the `dismount-Database` cmdlet to dismount the database.

```
dismount-Database -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
 For a VCS cluster setup, offline the VCS resource for the database manually if automatic offlining of the resource fails.
 Specifying the `-a` option with the `vxsnap restore` command dismounts the database automatically. In the case of a VCS cluster setup, it offlines the resource for the database automatically.
- 5 Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.

```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft Exchange Writer"
```

 For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer Replica" component=MB1 backupType=COPY -E -O
```
- 6 Use the `mount-Database` cmdlet to mount the database.

```
mount-Database -Identity MB1
```

 For a VCS cluster setup, online the VCS resource for the database manually if automatic onlining of the resource fails.

Specifying the `-a` option with the `vxsnap restore` command utility mounts the database automatically. In the case of a VCS cluster setup, it online the resource for the database.

- 7 If you have suspended replication as shown in [step 2](#), then use the `Update-MailboxDatabaseCopy` cmdlet.

```
Update-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.

Recovering the mailbox database to the Point of Failure (POF)

Use the following procedure to recover the mailbox database to the point of failure.

Refer to the Microsoft Exchange Shell cmdlets.

Note: When restoring after a hardware failure, you must use the command line and make sure to specify the `-r` option.

Caution: If you choose to restore the mailbox database to the point of failure, you cannot later restore to the point of time. You can only perform one recovery procedure on a mailbox database.

Caution: You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.

To recover the mailbox database to the point of failure

- 1 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database.
- 2 If replication is enabled for a cluster setup, then use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay).

```
Suspend-MailboxDatabaseCopy -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
- 3 Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set the checkbox for **“This database can be overwritten by Restore”**.

```
set-mailboxdatabase -identity MB1-allowfilerestore:$true
```

- 4 Use the `dismount-Database` cmdlet to dismount the database.


```
dismount-Database -Identity MB1[-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.
 For a VCS cluster setup, offline the VCS resource for the database manually if automatic offlining of the resource fails.
 Specifying the `-a` option with the `vxsnap restore` command dismounts the database automatically. In the case of a VCS cluster setup, it offlines the resource automatically.
- 5 Now perform the VSS restore operation using either the VEA wizard or the `vxsnap` CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot.


```
vxsnap -x snapdata.xml restore RestoreType=PIT writer="Microsoft Exchange Writer"
```

 For this example, assume that the snapshot of the replica was performed with


```
vxsnap -x snapdata.xml create writer="Microsoft Exchange writer Replica" component=MB1 backupType=COPY -E -O
```
- 6 Use the `mount-Database` cmdlet to mount the database.


```
mount-Database -Identity MB1
```

 For a VCS cluster setup, online the VCS resource for the database manually if automatic onlining of the resource fails.
 Specifying the `-a` option with the `vxsnap restore` command utility mounts the database automatically. In the case of a VCS cluster setup, it onlines the resource for the database.
- If you have suspended replication as shown in [step 2](#), then use the `Update-MailboxDatabaseCopy` cmdlet.


```
Update-MailboxDatabaseCopy -Identity MB1-SourceServer -DeleteExistingFiles [-Confirm:$False]
```

 Use `-Confirm` to prevent user intervention.

Refreshing the snapshot set

After recovering from a hardware failure, refresh the snapshot set.

See “[Refreshing the snapshot set](#)” on page 137.

In the example, DB1snap will reside on the disks where the original snapshot volumes were and DB1-tlogs will reside on the disk(s) where the original production volume was.

Refreshing the snapshot set

Refresh your snapshot set after recovering from the hardware failure. You must decide if the disks which now store the data are suitable for production or if you want to move the production volumes to different disks which you earlier added to the dynamic disk group as replacement disks.

Note: Note that during the process of refreshing the snapshot set any volumes that were originally snapshot volumes of missing production volumes will become master (production) volumes. These volumes currently reside on the disks on which the original snapshot volumes were created.

Choose one of the following options for your snapshot set:

- [“Refreshing the snapshot set on the current disks”](#) on page 137.
- [“Moving the production volumes to different disks and refreshing the snapshot set”](#) on page 144.

Refreshing the snapshot set on the current disks

Perform the following tasks to refresh the snapshot set on the current disks. In this case the volumes that were originally snapshot volumes of missing production volumes will become production volumes and will reside on the disks on which the original snapshot volumes were created.

Complete the following tasks to refresh the snapshot set on the current disks:

- Snap back the snapshot transaction logs volumes if you recovered using a roll-forward recovery to the point of failure. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them). See [“Reattaching healthy snapshot volumes”](#) on page 138.
- Snap clear each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. The snap clear removes the association of the snapshot volume with the original production volume. See [“Clearing the snapshot association for volumes whose drive letters or mount points were reassigned”](#) on page 139.
- Prepare the snapshot mirror for each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. See [“Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned”](#) on page 139.

- Use the `vxsnap create` command to create a new snapshot set of all the volumes in the mailbox database.
See “[Creating the new snapshot set](#)” on page 140.

Reattaching healthy snapshot volumes

If you recovered using a roll-forward recovery to the Point of Failure (POF), you must snap back the snapshot transaction logs volume using the **Resynchronize using the original volume** option. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them),.

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot snapshot set so that it matches the current state of its original Exchange mailbox database.

For the complete syntax of the `vxassist snapback` command, see *Veritas Storage Foundation Administrator’s Guide*.

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot set.
- 2 From the VEA console URL bar, select the `<host name>` which is the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click on the node of the application and click **VSS Snapback**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.
The XML metadata file contains all required information needed to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.
- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click Yes to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close all open handles on the snapshot volumes.
- 2 Type the following command, replacing the example values:

```
> vxassist -gExch-DG1 snapback DB1-tlogssnap
```

If there is an open handle on the volume you will get the following error message “One of the volumes of this dynamic disk group is currently in use. Please close applications and try again.”
- 3 Repeat the command for any other snapshot volumes in the mailbox database that were not affected by the disk failure.

Clearing the snapshot association for volumes whose drive letters or mount points were reassigned

The clear removes the association of the snapshot volume with the original production volume. Perform the snap clear operation only on those volumes whose drive letter or mount point was reassigned prior to the restore process. Volumes that remained healthy throughout the disk failure have an existing snapshot mirror that will continue to be used as part of the snapshot set.

To snap clear a volume

- 1 Right-click the snapshot volume and select **Snap > Snap Clear**.
- 2 Click **Yes** to confirm the snap clear operation for the specified volume.
- 3 Repeat this process for each volume whose drive letter or mount point was reassigned prior to the restore process.

Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned

Use the Prepare command to prepare each of the volumes whose drive letter or mount point was reassigned prior to the restore process. Make sure to select disks or LUNs that are not used for production data. However, you may create more than one snapshot volume on the same disk or LUN as long as there is sufficient space available.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare**.
- 2 In the **Prepare** dialog box:
 - Choose **Manually select disks**.

- Use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box.
 - Click **OK**.
- 3 Repeat the procedure to create a snapshot mirror for each volume whose drive letter or mount point was reassigned prior to the restore process. Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before continuing with the `vxsnap create` command.

To create snapshot mirrors from the command line

- 1 Type the command as in the following example, which prepares volume G: and selects `harddisk3` for the split-mirror snapshot volume:

```
vxassist prepare G: Harddisk3
```
- 2 Repeat [step 1](#) to create snapshot mirrors of each volume associated whose drive letter or mount point was reassigned prior to the restore process. Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before continuing with the `vxsnap create` command.

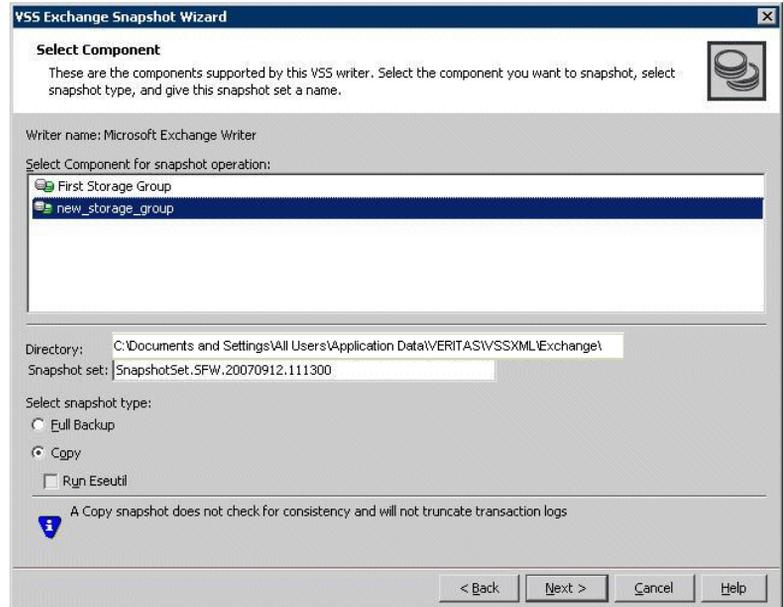
Creating the new snapshot set

Create the new snapshot set from the VEA console or from the command line.

To create the snapshot set from the VEA console

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a snapshot on the replication node, right-click **Exchange Replica** and click **VSS Exchange Snapshot**.
- 4 In the wizard, review the Welcome page and click **Next**.

5 Specify the snapshot set parameters as follows and then click **Next**:



Select Component for snapshot operation

Select the component for the snapshot set.

For Exchange 2010, the component is a mailbox database.

Directory

Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory shown on the screen.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore will fail.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, either edit the directory path in the Directory field for this wizard screen or use a text editor to create a text file named "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, `G:\BackupSets`. Save the **redirect.txt** file in the default directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML` on each node of the cluster.

Note: You must not use the volume name or volume path in the `redirect.txt` file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.

Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. You can also specify the file share path in the `redirect.txt` file.

Snapshot set

Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.

Select snapshot type

Select the snapshot type.

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the selected component
- Runs **Eseutil** to check for consistency before truncating the logs
- Truncates the transaction logs

Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally check **Run Eseutil** with the **Copy** option to check the snapshot for consistency.

- 6 In the Change Attributes panel, optionally change the attributes for the snapshot volumes as follows and click **Next**:

Snapshot Volume Label

Displays the read-only label for the snapshot volume.

Drive Letter

Optionally, click a drive letter and select a new choice from the drop-down menu.

Plex

Optionally, click a plex and select a new choice from the drop-down menu.

- 7 Optionally, in the Synchronized Snapshot panel, select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

To create the new snapshot set from the command line

- ◆ Type the `vxsnap create` command. The exact command will vary depending on your environment.

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<n
ewVolName>] [/Plex=<plexName>] .. writer=<writerName>
component=<componentName> [backuptype=<backuptype>] [-E] [-O] [-C] [
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Any text string that contains spaces must be enclosed in quotation marks.

Note: If you have scripted the refresh process, you will need to change the snapshot volume names assigned by your script using the option `Newvol=NewVolName`. For instance, if DB1 volume was missing, and the volume DB1snap was used in the restore process, it will retain the volume name DB1snap. Your script will not be able to assign DB1snap as the new snapshot volume name in step 4 when using the `vxsnap create` command.

See “[vxsnap create](#)” on page 157.

Moving the production volumes to different disks and refreshing the snapshot set

During the process of refreshing the snapshot set any volumes that were originally snapshot volumes of missing production volumes will become master (production) volumes. These volumes currently reside on the disks on which the original snapshot volumes were created.

Perform the following tasks to move the master (production) volumes to different disks:

- Snap back the snapshot transaction logs volumes if you recovered using a roll-forward recovery to the point of failure. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them) .
See “[Reattaching healthy snapshot volumes](#)” on page 145.
- Snap clear each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. The snap clear removes the association of the snapshot volume with the original production volume.
See “[Clearing the snapshot association for volumes whose drive letters or mount points were reassigned](#)” on page 146.
- Add a mirror to each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. For each new mirror (plex), choose the disks where you want the production volumes to reside.

Moving the production volumes to different disks and refreshing the snapshot set

See “[Adding mirrors to volumes whose drive letters or mount points were reassigned](#)” on page 146.

- Prepare each of the new mirror volumes and convert an existing mirror into a snap plex. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex).
See “[Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned](#)” on page 147.
- Use the `vxsnap create` command to create a new snapshot set of all the volumes in the mailbox database.
See “[Creating the new snapshot set](#)” on page 148.

Reattaching healthy snapshot volumes

If you recovered using a roll-forward recovery to the point of failure, use the VSS Snapback wizard or `vxsnap reattach` command to reattach the snapshot transaction logs volume and the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them).

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot set.
- 2 From the VEA console URL bar, select the `<host name>` which is the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click on the node of the application and click **VSS Snapback**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.
The XML metadata file contains all required information needed to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.
- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click Yes to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command, replacing the example file name:

```
vxsnap -x Image1.xml reattach writer="Microsoft Exchange Writer"
```

The complete syntax for the `vxsnap reattach` command is:

```
vxsnap -x <filename> [-f][-b] reattach writer=<writername>
[secHosts=<secondary hosts>]
```

Clearing the snapshot association for volumes whose drive letters or mount points were reassigned

The `snap clear` removes the association of the snapshot volume with the original production volume. Perform the `snap clear` operation only on those volumes whose drive letter or mount point was reassigned prior to the restore process. Volumes that remained healthy throughout the disk failure have an existing snapshot mirror that will continue to be used as part of the snapshot set.

To snap clear a volume

- 1 Right-click the snapshot volume and select **Snap > Snap Clear**.
- 2 Click **Yes** to confirm the Snap Clear operation for the specified volume.
- 3 Repeat this process for each volume whose drive letter or mount point was reassigned prior to the restore process.

Adding mirrors to volumes whose drive letters or mount points were reassigned

Add a mirror to each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. For each new mirror (plex), choose the disks where you want the production volumes to reside.

To add a mirror to a volume

- 1 Right-click on the volume you want to mirror, click **Mirror** and click **Add**.
- 2 Enter the number of mirrors to create.
- 3 Choose to manually assign the destination disks for the mirrors.
- 4 Select the disks in the Available disks list where you want the production volume to reside, and use the **Add** button to move them to the Selected disks list.

- 5 Click **OK** to add the mirror.

Note: If you create more than one mirror at a time, you may see inconsistent information on the progress bar. Also, the generation of multiple mirrors does affect system resources. After creating a mirror, you may want to wait until it has finished generating before creating another mirror.

Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned

Prepare each of the new mirror volumes and choose to convert an existing mirror into a snap plex. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex).

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare**.
- 2 The result of this command varies depending on whether or not there are additional disks available to create a new mirror. Choose from the following:

If the volume is mirrored and no additional disks are available to create a new mirror

- Click the mirror where you want the snapshot volume to be created. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex). The remaining mirror will become the master (production) volume.
- Click **OK**.

If the volume is mirrored and there are additional disks available on your system

- Choose an existing mirror for the snapshot
- Click **Select existing mirror for snap**.
- Click the mirror where you want the snapshot volume to be created. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex). The remaining mirror will become the master (production) volume.
- Click **OK**.

- 3 Repeat the procedure to create a snapshot mirror for each volume associated with the Exchange mailbox database and transaction logs.

Caution: Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before creating the new snapshot set.

Creating the new snapshot set

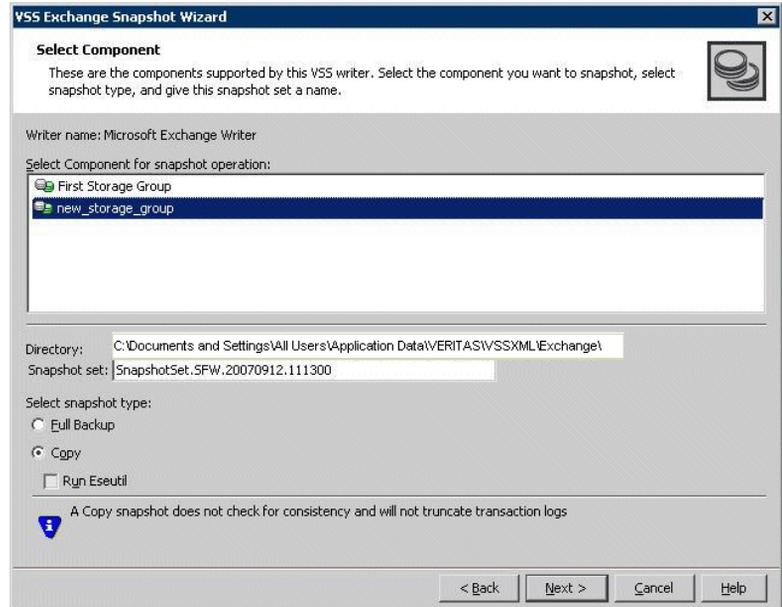
Create the new snapshot set from the VEA console or from the command line.

To create the snapshot set from the VEA console

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a snapshot on the replication node, right-click **Exchange Replica** and click **VSS Exchange Snapshot**.
- 4 In the wizard, review the Welcome page and click **Next**.

Moving the production volumes to different disks and refreshing the snapshot set

- 5 Specify the snapshot set parameters as follows and then click **Next**:



Select Component for
snapshot operation

Select the component for the snapshot set.

For Exchange 2010, the component is a mailbox
database.

Directory

Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory shown on the screen.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore will fail.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, either edit the directory path in the Directory field for this wizard screen or use a text editor to create a text file named "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, `G:\BackupSets`. Save the **redirect.txt** file in the default directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML` on each node of the cluster.

Note: You must not use the volume name or volume path in the `redirect.txt` file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.

Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. You can also specify the file share path in the `redirect.txt` file.

Snapshot set

Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.

Moving the production volumes to different disks and refreshing the snapshot set

Select snapshot type

Select the snapshot type.

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the selected component
- Runs **Eseutil** to check for consistency before truncating the logs
- Truncates the transaction logs

Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally check **Run Eseutil** with the **Copy** option to check the snapshot for consistency.

- 6 In the Change Attributes panel, optionally change the attributes for the snapshot volumes as follows and click **Next**:

Snapshot Volume Label

Displays the read-only label for the snapshot volume.

Drive Letter

Optionally, click a drive letter and select a new choice from the drop-down menu.

Plex

Optionally, click a plex and select a new choice from the drop-down menu.

- 7 Optionally, in the Synchronized Snapshot panel, select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

To create the new snapshot set from the command line

- Type the `vxsnap create` command. The exact command will vary depending on your environment. For example:

```
> vxsnap -x TestDB.xml create source=G:/Newvol=DB1snap2
source=I:/Newvol=DB1-tlogssnap2
writer="Microsoft Exchange Writer" component=Exch-DG1
backuptype=copy
```

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]  
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<n  
ewVolName>] [/Plex=<plexName>] ...writer=<writerName>  
component=<componentName>[backuptype=<backuptype>] [-E] [-O] [-C] [  
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Any text string that contains spaces must be enclosed in quotation marks.

Note: If you have scripted the refresh process, you will need to change the snapshot volume names assigned by your script using the option `Newvol=NewVolName`. For instance, if DB1 volume was missing, and the volume DB1snap was used in the restore process, it will retain the volume name DB1snap. Your script will not be able to assign DB1snap as the new snapshot volume name in step 4 when using the `vxsnap create` command.

See “[vxsnap create](#)” on page 157.

Vxsnap utility command line reference

This chapter includes the following topics:

- [About the vxsnap utility](#)
- [Vxsnap keywords](#)

About the vxsnap utility

The command line utilities are available in the Veritas Storage Foundation for Windows installation directory.

The vxsnap utility integrates with the Windows Volume Shadow Copy Service (VSS) as a VSS Requester. This allows for the simultaneous snapshot of all volumes associated with an Exchange Server database.

Note the following requirements:

- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before using the vxsnap utility
- The CLI commands run only on the server. They will not run on the Veritas Storage Foundation for Windows client.
- The vxsnap commands must be invoked on a local system.
- For Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Note: Any of the writer names 'Microsoft Exchange Replica Writer' or 'Microsoft Exchange Writer Replica' can be specified in the vxsnap CLI command while performing operations on a replica.

Vxsnap keywords

Type the following sequence to view a description of the command syntax:

```
vxsnap keyword -?
```

vxsnap has the following keywords:

<code>prepare</code>	<p>Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes</p> <p>Either the <code>prepare</code> or <code>start</code> keyword may be used in the CLI; however <code>prepare</code> is recommended.</p>
<code>create</code>	<p>Creates simultaneous snapshots of all volumes in the specified component, providing a point-in-time snapshot set.</p>
<code>reattach</code>	<p>Reattaches and resynchronizes an existing snapshot set to the original database volumes.</p>
<code>restore</code>	<p>Restores an Exchange database from a snapshot set. Recovery can be either a point-in-time restore or a roll-forward recovery to the point of failure.</p> <p>Note: For Exchange 2010, Restore to Recovery Database is not supported by SFW.</p>

vxsnap prepare

Creates snapshot mirrors of the volumes in the specified component and eliminates the need for multiple `vxassist prepare` commands. The snapshot mirrors remain attached to and synchronized with the original volumes.

Syntax

```
vxsnap prepare component=<componentName>/writer=<writerName>  
[-b] [source=<volume>/harddisk=<harddisk,...>] ...]
```

Attributes

The following attributes apply:

<code>component=<ComponentName></code>	Name of the component; for Exchange 2010, this is the mailbox database name, for example, "DB1".
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer.
<code>-b</code>	Resynchronizes the volume in the background. A snapshot cannot be made until the resynchronization is complete.
<code>source=<Volume></code>	Indicates the source volume for the snapshot mirror specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\".
<code>harddisk=<Harddisk></code>	Name of the disk where the snapshot mirror is prepared, for example, <code>harddisk2</code> .

Example

```
vxsnap prepare component=DB1/writer="Microsoft Exchange Writer"  
source=L:/harddisk=harddisk2 source=M:/harddisk=harddisk3
```

This command creates snapshot mirrors of the volumes in an Exchange mailbox database (DB1). The snapshot mirror of the volume mounted on drive letter L: will be created on disk 2 and the snapshot mirror of the volume mounted on drive M: will be created on disk 3.

vxsnap create

Creates simultaneous snapshots of all volumes in a specified component, providing a point-in-time snapshot set.

Separate source volumes and attributes with forward slashes, not spaces. Source and snapshot volume attributes are paired. You must specify the source volume if you choose to specify the snapshot volume plex, drive letter, drive path, label, or volume name.

Note that if Exchange 2010 mailbox database is configured under VCS, then you can store the snapshot set file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over.

There can be large number of mailbox databases in Exchange 2010 and each database may have independent schedules. Hence number of different paths required for storing the snapshot metadata files are higher, which is achieved by configuring a file share resource. You can also specify the file share path in the `redirect.txt` path, for example, `\\MySystemName1\share1\QRDir`.

In case of schedule snapshot even though schedules are visible on all nodes, snapshot will happen only on a node where mailbox database is present.

See [“Creating a one-time snapshot set”](#) on page 97.

Syntax

```
vxsnap [-x <filename>] create [source=<volume>]  
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>]  
[/Newvol=<newVolName>] [/Plex=<plexName>]...writer=<writerName>  
component=<componentName>[backuptype=<backuptype>] [-E] [-O] [-C]  
[secHosts=<secondary hosts>]
```

Attributes

The following attributes apply:

- `-x <Filename>` Indicates the name to be assigned to the XML metadata file that will be created with the command. The file name must include the ".xml" extension. The default path to the file is under the SFW program files directory (normally C:\Documents and Settings\All Users\Application Data\Veritas\VSSXML\Exchange). If you wish to place the file in another directory, specify a full path before the file name, for example J:\XML\Image1.xml. For a VCS cluster setup, you can configure a file share resource and give the file share resource path in the redirect.txt. See ["Creating a one-time snapshot set"](#) on page 97.
- `source=<Volume>` Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified component (for example, Exchange storage group).
- `[/plex=<PlexName>]` Specifies the name of the mirror or plex that is to be detached. Use this parameter if there are multiple snap plexes available to be snapshotted.
- `[/DriveLetter=<DriveLetter>]` The drive letter to be assigned to the new snapshot volume.
- `[/DrivePath=<DrivePath>]` The drive path to be assigned to the new snapshot volume. The drive path must reference an empty local NTFS folder, which was created beforehand. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.

[/Newvol=<NewVolName>]	Specifies the name of the new snapshot volume that is created. If the name is not specified using this option, the form "SnapVolume01" is created. The full device path becomes: \\Device\HarddiskDmVolumes\ <DiskGroupName>\<NewVolName>
writer=<WriterName>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer. If you plan to be able to do a roll-forward recovery to the point of failure, you must specify the writer.
component=<ComponentName>	Name of the component; for Exchange 2010, this is the mailbox database name, for example, DB1. If you plan to be able to do a roll-forward recovery to the point of failure, you must use this component.
backuptype=<Backuptype>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default. Copy backup creates a copy of the database and transaction logs volumes. Full backup creates a copy of the database and transaction logs volumes, runs Eseutil to check for consistency, and if consistent, truncates the transaction logs.
-E	Runs the Eseutil consistency check for database and log files. Eseutil is run automatically with a Full backup, but must be optionally specified for a copy backup.
-o	Allows an existing XML file of the same name to be overwritten. If -O is not specified the vxsnap create command does not overwrite an existing XML file of the same name and the operation fails.

`sechosts=<SecondaryHosts>` Applies to Veritas Volume Replicator (VVR) environment only. Comma separated list of secondary hosts on which a synchronized snapshot is to be taken.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Examples

```
vxsnap -x backupdoc.xml create writer="Microsoft Exchange  
Writer" component="DB1" source=L:/DriveLetter=O  
source=M:/DriveLetter=P backuptype=full
```

This example creates a snapshot set based on the component “DB1,” which contains volume L, the log volume, and volume M, the database volume.

The snapshots are assigned drive letters O and P respectively. The XML file backupdoc.xml is used to store the VSS metadata that identifies the snapshot set. The Full backup creates the copy of the database and transaction logs volumes, runs Eseutil to check for consistency, and then truncates the transaction logs.

```
vxsnap -x snapdata.xml  
createsource=E:\Data\DB1/DrivePath=E:\Backup\DB1
```

This command creates snapshots from the volume mounted on E:\Data\DB1. The resulting snapshot volume is assigned mount point E:\Backup\DB1. The metadata involved in this operation is stored in snapdata.xml.

vxsnap reattach

Reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

Note: Make sure that the snapshot volumes are not in use before using this command.

Syntax

```
vxsnap -x <filename> [-f][-b] reattach writer=<writername>  
[secHosts=<secondary hosts>]
```

Attributes

The following attributes apply:

<code>-x <Filename></code>	The file created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file. Note: This file is deleted after the reattach operation has completed successfully.
<code>-f</code>	Forces the reattach. Make sure the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>writer=<writername></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer. If you plan to be able to do a roll-forward recovery to the point of failure, you must specify the writer.
<code>secHosts=<secondary hosts></code>	Applies to Veritas Volume Replicator (VVR) environment only. Comma separated list of secondary hosts on which a synchronized snapshot is to be taken.

Example

```
vxsnap -x snapdata.xml reattach writer="Microsoft Exchange  
Writer"
```

This command uses the information in the `snapdata.xml` file to reattach and resynchronize all the volumes in the snapshot set. This `xml` file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the `vxsnap create` command is issued.

vxsnap restore

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to restore data, for example, after an original volume has become corrupted. You can restore the data either to the point in time that the snapshot set was last refreshed or to the point of failure.

Before using this command, make sure that the source volumes and the snapshot volumes are not in use. Use the [-a] attribute to dismount and mount the database automatically or use exchange to dismount and mount the database before and after the restore operation respectively.

For a VCS cluster setup, specifying the -a option offlines the VCS resource database before a restore operation and onlines the resource after restore operation.

Note: After completing a point of failure (POF) recovery of a database, Symantec recommends using the **vxsnap reattach** command to reattach and resynchronizing the log volumes. You can then use the **vxsnap create** command to create a new snapshot set.

Syntax

```
vxsnap -x <filename>[-f][-b][-r] [-a] restore  
restoreType=<PIT|POF>writer=<writername> [subComponent=<subComponentName>] [RSG=<Yes|No>]
```

Note: The Subcomponent and RSG=Yes | No options for restore are not valid for Exchange 2010.

Attributes

The following attributes apply:

-x <Filename>	The file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file.
-f	Forces the snapback. Make sure the volume is not in use by another application before using this command. Use this option with care.
-b	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
-r	Recover one or more of the original volumes that are missing. Example below shows additional required steps.

<code>-a</code>	Use the <code>[-a]</code> attribute to dismount and mount the database automatically. For a VCS cluster setup, specifying the <code>-a</code> option offlines the VCS resource database before a restore operation and onlines the resource after restore operation.
<code>restoreType=<PIT POF ></code>	PIT specifies a restore to the point in time that the snapshot set was created or last refreshed. POF specifies a roll-forward recovery to the point of failure.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer. If you specify the POF option, you must specify the writer.
<code>subComponent=<subComponentName></code>	Note: The subcomponent attribute is not valid for Exchange 2010.
<code>[RSG=<Yes No>]</code>	Exchange 2007 only. <code>Yes</code> restores the snapshot set to a recovery storage group (RSG). In order to restore to an RSG, the RSG for that storage group must already exist, and databases created inside the RSG must have the same names as that of the source storage group. When performing a PIT recovery for an RSG, all subcomponents of the storage group are recovered. When performing a POF recovery for an RSG, only the subcomponents that you specify in the command are recovered.

Examples

Restore to the Point in Time

```
vxsnap -x snapdata.xml restore restoreType=PIT writer="Microsoft Exchange Writer"
```

This command uses the information in the `snapdata.xml` file to restore all the volumes in the snapshot set identified in that file to the point in time the snapshot set was created or last refreshed.

Roll-Forward Recovery to the Point of Failure

```
vxsnap -x snapdata.xml restore restoreType=POF writer="Microsoft Exchange Writer"
```

This command uses the information about the database specified in the `snapdata.xml` file to snapback the database volumes and then use current transaction logs to roll forward to the point of failure.

Recovery after hardware failure

You can use the `-r` switch to perform a VSS-integrated recovery after a hardware failure. The following recovery scenarios are possible if the complete snapshot set including the XML metadata file is available:

Table 10-1 Available recovery type after missing production volume

Scenario	Database Volumes	Transaction Logs Volume	Recovery
Scenario I	One or more volumes are missing.	Missing	point in time
Scenario II	All volumes are missing.	Available	point in time or point of failure

Caution: Before using the `vxsnap restore` command, verify that you have correctly assigned the drive or mount point to each volume.

It is important that each of the tasks be performed exactly as described.

See [Chapter 9, “Recovering after hardware failure”](#) on page 117 for details.

Complete the following tasks to perform a VSS-integrated recovery:

- Prepare for the recovery
 - Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
 - Delete the missing volumes from Storage Foundation for Windows.
 - Replace the failed hardware and add the new disks to the dynamic disk group.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.

- Use the `vxsnap restore` command to recover the databases by including the `-r` switch in the `vxsnap restore` command.

For example, run the following command:

```
vxsnap -x snapdata.xml -r restore restoreType=PIT  
writer="Microsoft Exchange Writer"
```

This command uses the information in the `snapdata.xml` file to restore all the volumes in the snapshot set identified in that file to the point in time the snapshot set was created or last refreshed.

- Refresh the snapshot set.

Index

C

- commands for snapshots
 - prepare 145
 - snap back 144
 - snap clear 139, 144, 146
 - vxsnap 153
 - vxsnap create 157
 - vxsnap prepare 156
 - vxsnap reattach 161
 - vxsnap restore 162
- configuration
 - Exchange storage for Quick Recovery 53

D

- Data Availability Group (DAG) considerations 47
- disk groups
 - creating
 - Exchange Quick Recovery 54
 - Exchange with Microsoft cluster 54
- drive letters 124, 131

E

- Exchange databases
 - dismounting for hardware recovery 129
 - setting paths for Quick Recovery 59
- Exchange storage configuration for Quick Recovery 53
- Exchange transaction log paths for Quick Recovery 59

F

- FastResync 29
- FlashSnap 29

H

- hardware recovery
 - database and transaction logs volumes
 - missing 121
 - deleting missing volumes 123
 - identifying the missing volumes 122, 129
 - prerequisites 120
 - replacing hardware 123
 - sample configuration 120
 - tasks 118
 - using snapshots 117

M

- Microsoft Exchange Replica Writer 48
- mirrors
 - adding to volumes whose drive letters or mount points were reassigned 146
 - preparing from VEA or vxsnap 84
 - scheduling preparation with Quick Recovery Configuration Wizard 66
- missing volumes
 - deleting 123, 129
- mount points
 - changing 124
 - snapshot volumes 131

P

- passive copies
 - recovery 107
- passive copies in a DAG 47, 48
 - recovery 48
- pre and post snapshot script file locations in VCS
 - cluster 45
- prepare command 145
- prerequisites
 - hardware recovery 120
 - Quick Recovery 52
 - see also requirements

Q

Quick Recovery

- best practices 42
- components 29
- configuring Exchange storage 51
- methods 38
- overview 27
- planning 35, 41
- prerequisites 52
- process 29
- recommendations 42
- recovery 105
- storage requirements 36
- system requirements 35
- VCS considerations 44
- VVR considerations 48

Quick Recovery Configuration Wizard

- overview 61
- prerequisites 65
- running 66

R

Recovery

- using passive copy snapshot in a DAG 107

recovery

- adding disks to the dynamic disk group 130
- after hardware failure 117, 166
- prerequisites 120
- replacing hardware 130
- tasks for hardware recovery 118
- using passive copy in a DAG 107
- vxsnap utility 112

replication copies in a DAG 47, 48

Replication Writer 47

requirements

- Quick Recovery 35

roll-forward recovery

- to point of failure 166
- to point of failure of a single database 166

S

schedules for snapshot sets

- creating with Quick Recovery Configuration Wizard 66
- creating with VSS Snapshot Scheduler Wizard 85

deleting 78, 79

in a DAG 47

modifying 78

passive copies 47

synchronizing after adding a cluster node 79

troubleshooting 76

snap back 144

snap clear 139, 144, 146

snapshot volumes

changing drive letter 124, 131

changing mount points 124

reattaching healthy volumes 138, 145

snapshots

clearing association for volumes 139, 146

creating after hardware recovery 140, 148

creating mirrors after hardware recovery 139, 147

methods of implementing 38

modifying schedules 78

overview 62

planning 35

reattaching split-mirror 101

recovery storage group (RSG) 106

refreshing after hardware recovery 127, 136, 137

refreshing manually 101

refreshing on the current disks 137

scheduling with VSS Snapshot Scheduler Wizard 85

storage requirements 36

system requirements 35

templates 63

troubleshooting 76

VCS considerations 44

viewing status 75

VVR considerations 48

Solutions Configuration Center

context sensitivity 15

overview 13

running wizards remotely 21

starting 14

wizard descriptions 21

storage group

recovering to the point of failure 135

restoring to the point in time 126, 133

storage requirements for Quick Recovery 36

Store Writer 47

synchronizing schedules in a cluster 79

T

- templates for snapshot sets
 - description 63
 - multiple components 64
 - schedule start dates 64

V

- VCS cluster
 - Quick Recovery considerations 44
- Volume Shadow Copy Service 30
- volumes
 - creating
 - Quick Recovery 55
- VSS
 - defined 30
 - framework 30
- VSS Snapback Wizard 101
- VSS Snapshot Scheduler Wizard 85
- VVR
 - Quick Recovery considerations 48
- vxsnap create 157
- vxsnap prepare 156
- vxsnap reattach 161
- vxsnap restore 162
- vxsnap utility command reference 153
- vxsnap utility, using for recovery 112

X

- XML file locations 44

