

Veritas Storage Foundation™ and High Availability Solutions Quick Recovery Solutions Guide for Microsoft SQL 2008 and 2008 R2

Windows Server 2003,
Windows Server 2008

5.1 Service Pack 2



Veritas Storage Foundation and HA Solutions Quick Recovery Solutions Guide for Microsoft SQL 2008 and 2008 R2

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 2

Document version: 5.1.SP2.0

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

Contents

Section 1 Introduction

Chapter 1	Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL	
	About the solutions guides	13
	Quick Recovery	14
Chapter 2	Using the Solutions Configuration Center	
	About the Solutions Configuration Center	15
	Starting the Configuration Center	16
	Available options from the Configuration Center	16
	About running the Configuration Center wizards	18
	Solutions wizard logs	19
Chapter 3	Quick Recovery overview for SQL Server	
	About Quick Recovery snapshot solutions	21
	About snapshot-assisted backups	22
	Advantages of Quick Recovery snapshots	22
	Quick Recovery process	23
	Methods of implementing Quick Recovery snapshots for SQL Server	24
	About the Quick Recovery Configuration Wizard	24
	About the VSS Snapshot Scheduler Wizard	25
	About the VSS Snapshot and Snapback wizards and the vxsnap utility	26
	About the components used in Quick Recovery	27
	FlashSnap and FastResync	27
	Integration with Microsoft Volume Shadow Copy Service	27
	VCS, Microsoft clustering, and VVR considerations	29

Section 2 Quick Recovery

Chapter 4	Preparing to implement Quick Recovery for SQL Server	
-----------	--	--

Tasks for preparing to implement Quick Recovery for SQL Server	33
Reviewing the prerequisites	34
Supported software	34
Storage requirements and best practices	38
Configuration requirements and best practices	39
Reviewing the configuration	40
Configuring SQL Server storage with Veritas	
Storage Foundation for Windows	41
Creating dynamic disk groups	41
Creating dynamic volumes	43
Pointing the databases and log paths to the SFW volumes	45

Chapter 5 Implementing Quick Recovery for SQL Server with the configuration wizard

About the Quick Recovery Configuration Wizard	47
Backup types for snapshot sets	49
About snapshot templates	49
Tasks for implementing snapshot sets with the configuration wizard	51
Reviewing the prerequisites	52
Scheduling and creating snapshot sets	52
Viewing the status of scheduled snapshots	61
Troubleshooting scheduled snapshots	62
Deleting or modifying schedules	64
Synchronizing schedules after adding a cluster node	65

Chapter 6 Scheduling or creating an individual snapshot set for SQL Server

About scheduling or creating an individual snapshot set	67
Tasks to schedule a new snapshot set	68
Tasks to create a one-time snapshot set	69
Reviewing the prerequisites	70
Preparing the snapshot mirrors	70
Scheduling a new snapshot set	72
Viewing the status of scheduled snapshots	78
Troubleshooting scheduled snapshots	79
Deleting or modifying schedules	81
Creating a one-time snapshot set	82
Refreshing a snapshot set manually	86
Reattaching the split-mirror snapshots	86
Creating the refreshed snapshot set	87

Chapter 7	Recovering a SQL Server database	
	About recovering a SQL Server database	89
	Tasks for recovering a SQL Server database	90
	Prerequisites for recovering a SQL Server database	91
	Types of recovery	91
	Recovering using snapshots without log replay	92
	Recovering using snapshots and log replay	94
	Restoring snapshots and manually applying logs	97
	Recovering missing volumes	99
	Preparing for the recovery	100
	Performing the recovery	103
	Post-recovery steps	105
	Vxsnap restore command reference	106
Chapter 8	Vxsnap utility command line reference for SQL Server	
	About the Vxsnap utility	109
	Vxsnap keywords	110
	vxsnap prepare	110
	vxsnap create	111
	vxsnap reattach	114
	vxsnap restore	115
Index.....		119

Introduction

- [Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL](#)
- [Using the Solutions Configuration Center](#)

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL

This chapter covers the following topics:

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

About the solutions guides

Solutions for Microsoft SQL for Quick Recovery are found in this manual, *Veritas Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft SQL 2008*.

Solutions for Microsoft SQL for High Availability and Disaster Recovery are in *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2008*.

Solutions for Microsoft SQL for Microsoft clustering are in *Veritas Storage Foundation and High Availability Solutions Microsoft Clustering Solutions Guide for Microsoft SQL 2008*.

Separate guides are available for Microsoft Exchange solutions and for other application solutions.

Quick Recovery

Quick Recovery is the process of creating and maintaining on-host point-in-time copies of production data that can be used to quickly recover SQL Server databases that have been subject to accidental or malicious updates or other data corruption.

A Quick Recovery solution serves as a first line of defense to recover SQL Server databases that have been subject to accidental or malicious updates. 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](#)
- [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
- Microsoft SQL Server
- Enterprise Vault Server (high availability solution only)
- Additional applications

You can use the Configuration Center and its wizards to set up your environment for Quick Recovery. Quick Recovery provides on-host recovery from logical errors in application data for Microsoft SQL Server.

Note: For information about the Configuration Center solutions for Microsoft Exchange Server and other applications, refer to the appropriate Symantec documentation.

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 two ways:

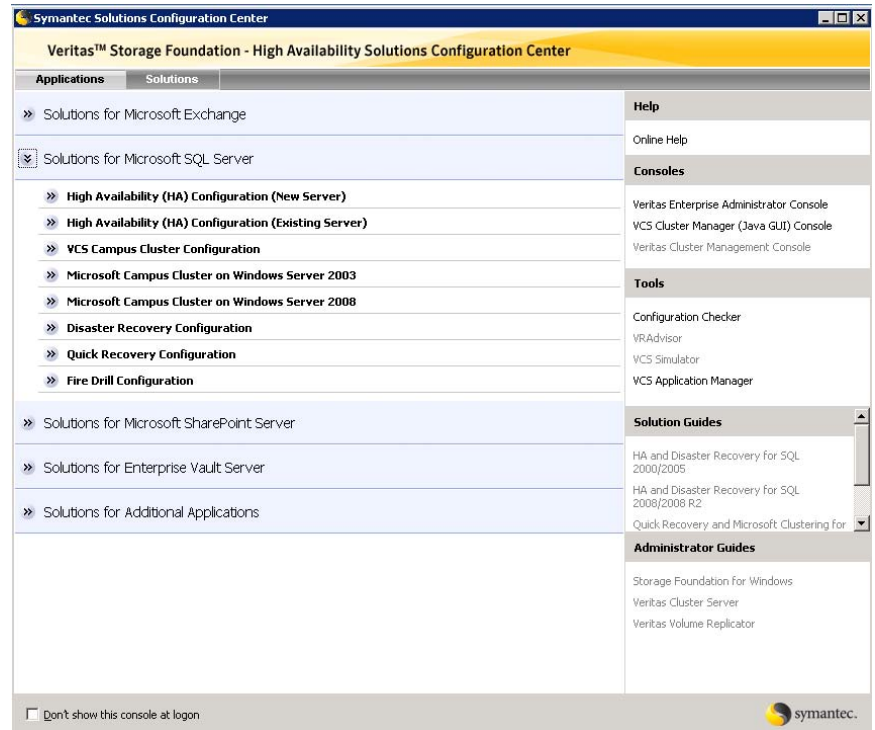
- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

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.

[Figure 2-1](#) on page 17 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-1 Solutions Configuration Center for Microsoft SQL Server



The submenu choices vary by application (Microsoft Exchange, Microsoft SQL Server, etc.). For example, [Figure 2-2](#) on page 18 shows one of the steps for implementing high availability for SQL Server from an accessible submenu.

Figure 2-2 Context-sensitive step for SQL Server



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 the Quick Recovery Configuration Wizard can be run from a remote system. The Quick Recovery Configuration Wizard schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots.

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

Solutions wizard logs

The Solutions Configuration Center provides access to many wizards. The Quick Recovery Configuration Wizard is built into the Solutions Configuration Center. The Quick Recovery Configuration Wizard is launched only from the Solutions Configuration Center, whereas other wizards can be launched from product consoles or the Start menu.

Logs created by this Solution wizard 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 overview for SQL Server

This chapter covers the following topics:

- [About Quick Recovery snapshot solutions](#)
- [About snapshot-assisted backups](#)
- [Advantages of Quick Recovery snapshots](#)
- [Quick Recovery process](#)
- [Methods of implementing Quick Recovery snapshots for SQL Server](#)
- [About the components used in Quick Recovery](#)
- [VCS, Microsoft clustering, and VVR considerations](#)

About Quick Recovery snapshot solutions

Veritas Storage Foundation for Windows (SFW) provides the capability to create a point-in-time image of all the volumes associated with a SQL Server instance. This image, called a snapshot set, is a complete copy of the SQL Server instance 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 the database and ensure a persistent snapshot of the production data.

Quick Recovery is the term for creating and maintaining the SFW snapshot sets on-host for use in quickly recovering databases in the event of corruption or an accidental or malicious update. 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 SQL Server instance and set up schedules for creating and refreshing the snapshot sets. The snapshot sets can be maintained on-host as a Quick Recovery solution.

If you are using Veritas Volume Replicator (VVR) for replication, you can also synchronize snapshot sets on the secondary site. See *Veritas Volume Replicator, Administrator's Guide*.

SFW snapshots 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.

About snapshot-assisted backups

Since a snapshot set created using SFW is an exact copy of the production data, it can be used with a backup application in place of the production volumes to produce regular backups to tape or other media. A snapshot set can be used for backup either on-host or moved to a secondary server.

A snapshot-assisted backup provides the following advantages:

- Reduces backup I/O load on the production volumes in an on-host environment, allowing backup to be completed with less impact to the users
- Eliminates backup I/O load on the production server in an off-host environment

Advantages of Quick Recovery snapshots

A Quick Recovery solution serves as a first line of defense to recover databases that have been subject to accidental or malicious updates. Quick Recovery is designed to augment your traditional backup methodology. 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 Snapshot Scheduler wizard.
- Using a snapshot set to recover a corrupted SQL Server instance or single database

Methods of implementing Quick Recovery snapshots for SQL Server

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

- Quick Recovery Configuration Wizard
- VSS Snapshot Scheduler
- VSS Snapshot and Snapback wizards and the vxsnap utility

[Table 3-1](#) summarizes the methods and when you would use them.

Table 3-1 Methods of implementing Quick Recovery snapshots

Tasks	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"> ■ Perform major updates to 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 24</p>
<ul style="list-style-type: none"> ■ Add a new snapshot set schedule for one database 	<p>From the VEA: VSS Snapshot Scheduler Wizard</p> <p>You can also use the Quick Recovery Configuration Wizard.</p>	<p>“About the VSS Snapshot Scheduler Wizard” on page 25</p>
<ul style="list-style-type: none"> ■ Create a one-time snapshot as needed 	<p>From the VEA: VSS SQL Snapshot Wizard and VSS Snapback Wizard</p> <p>You can also use the Vxsnap utility from the CLI.</p>	<p>“About the VSS Snapshot and Snapback wizards and the vxsnap utility” on page 26</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 implementation. 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 in a SQL Server instance. Each database can have different filegroups, such as Database, Log, and FILESTREAM filegroups (for SQL Server 2008).

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

The snapshot creation process integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. The snapshot is done while the database is online and without disrupting the database operations.

The wizard also enables you to save all configuration settings to an XML file that can be imported for use on other systems or instances.

See [Chapter 5, “Implementing Quick Recovery for SQL Server with the configuration wizard”](#) on page 47.

About the VSS Snapshot Scheduler Wizard

You can use the VSS Snapshot Scheduler Wizard to add a snapshot schedule for a selected database. This wizard uses the same scheduling process as the Quick Recovery wizard. However, you can only schedule one database at a time.

Unlike the Quick Recovery Configuration Wizard, which enables you to perform all aspects of snapshot configuration, the VSS Snapshot Scheduler Wizard does not include preparing the snapshot volume. Therefore, you must use the Prepare command before running the VSS Snapshot Scheduler Wizard to schedule a new snapshot set. Also, you can only schedule one database at a time.

You might use this wizard to add a schedule to the initial configuration set up 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.

Note: Adding a snapshot schedule using the VSS 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 [Chapter 6, “Scheduling or creating an individual snapshot set for SQL Server”](#) on page 67.

About the VSS Snapshot and Snapback wizards and the vxsnap utility

If you need to create an individual snapshot set, 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.

Unlike using the Quick Recovery Configuration Wizard, which enables you to perform all aspects of snapshot configuration, you must prepare the mirror for the snapshot manually, using the Prepare command (or vxsnap utility). In addition, these methods do not include the capability of scheduling periodic refreshing of the snapshot set. Instead you would need to manually reattach the mirror, allowing it to resynchronize, and then create a snapshot set again from the resynchronized mirror.

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

For more information about the VSS Snapshot and snapback wizards, see [Chapter 6, “Scheduling or creating an individual snapshot set for SQL Server”](#) on page 67.

For more information about the vxsnap command line utility, see [Chapter 8, “Vxsnap utility command line reference for SQL Server”](#) on page 109.

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.

Table 3-2 shows the vxsnap commands that perform the same actions as the VSS wizards.

Table 3-2 Actions of VSS wizard and vxsnap command

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

About the components used in Quick Recovery

SFW Quick Recovery snapshots 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 Windows 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 of a SQL Server instance 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 a SQL Server instance without taking the databases offline.

VSS framework

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

Table 3-3 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 the SQL Server instance for quiescing and later requests that the snap shot process begin.
Writer	As Writers, VSS-enabled applications such as SQL Server 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.
Provider	As a Provider, FlashSnap creates the persistent snapshot.

VSS 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 a SQL Server instance.
- The VSS coordinator service calls the SQL Server 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 SQL Server VSS Writer to quiesce and freeze the databases in preparation for the snapshot. When this is accomplished, the SQL Server 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 SQL Server VSS 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 SQL Server VSS Writer, which prepares for the restore operation.
- The FlashSnap utility restores the snapshot volumes. After the snapshot operation completes, the FlashSnap utility informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the SQL Server VSS Writer to process the post-restore operations.

VCS, Microsoft clustering, and VVR considerations

Certain requirements apply when implementing a snapshot solution with a VCS, Microsoft clustering, and VVR environment.

In a VCS or Microsoft clustering 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.

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

- If you use the Quick Recovery Configuration Wizard or the VSS SQL Snapshot Wizard to create the snapshot set, you can use either wizard to specify the file path to the appropriate volume.

- Use either the Quick Recovery Configuration Wizard or the VSS SQL Snapshot Wizard to create the snapshot set. After creating the snapshot set, you can store the XML files in a location of your choice using either the Quick Recovery Configuration Wizard which provides an option for specifying the path (by modifying the default populated path) or by the following method: 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.

Warning: You must not use the volume name or volume path in `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.

- When using `vxsnap` utility commands that require the filename attribute, specify the full path to the location of the XML metadata file.
- 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 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.

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 database so that the metadata file is available from all nodes in the cluster.
- 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.

Quick Recovery

This section includes the following chapters:

- [Quick Recovery overview for SQL Server](#)
- [Preparing to implement Quick Recovery for SQL Server](#)
- [Implementing Quick Recovery for SQL Server with the configuration wizard](#)
- [Scheduling or creating an individual snapshot set for SQL Server](#)
- [Recovering a SQL Server database](#)
- [Vxsnap utility command line reference for SQL Server](#)

Preparing to implement Quick Recovery for SQL Server

This chapter covers the following topics:

- [Tasks for preparing to implement Quick Recovery for SQL Server](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring SQL Server storage with Veritas Storage Foundation for Windows](#)

Tasks for preparing to implement Quick Recovery for SQL Server

[Table 4-1](#) outlines the high-level objectives and the tasks to complete each objective for preparing to implement SQL Server Quick Recovery.

Table 4-1 Tasks for preparing to implement SQL Server Quick Recovery

Objective	Tasks
“Reviewing the prerequisites” on page 34	■ Verifying hardware and software prerequisites and storage configuration best practices
“Reviewing the configuration” on page 40	■ Reviewing the sample configuration

Table 4-1 Tasks for preparing to implement SQL Server Quick Recovery

Objective	Tasks
“Configuring SQL Server storage with Veritas Storage Foundation for Windows” on page 41	<ul style="list-style-type: none">■ Creating a dynamic disk group■ Creating dynamic volumes■ Creating a SQL Server database and pointing the database and log paths to the SFW volumes

Reviewing the prerequisites

A Veritas Storage Foundation for Windows (SFW) Quick Recovery solution can be implemented on either a standalone system or a clustered system. Quick Recovery snapshots are supported in both Veritas Cluster Server (VCS) and Microsoft clusters.

This solution assumes that the required 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 for Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.

Supported software

Quick Recovery snapshots using Storage Foundation for Windows (SFW) require the following software:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW) with the FlashSnap option or Veritas Storage Foundation HA 5.1 Service Pack 2 for Windows (SFW HA) with the FlashSnap option.
- Any of the following Microsoft SQL 2008 servers and their operating systems

The following table lists the Microsoft SQL Server 2008 versions supported with SFW HA 5.1 Service Pack 2.

Table 4-2 Supported Microsoft SQL Server 2008 versions

SQL Server 2008	Windows Servers
Microsoft SQL Server 2008, 32-bit Standard Edition, Enterprise Edition, or Web Edition on Windows Server 2003 (SQL Server 2008 SP1)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)
Microsoft SQL Server 2008, 64-bit Standard Edition, Enterprise Edition, Enterprise IA64 Edition, Web Edition on Windows Server 2003 (SQL Server 2008 SP1)	<ul style="list-style-type: none"> ■ Windows Server 2003 (64-bit) Standard x64 Edition, Enterprise x64 Edition or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) ■ Windows Server 2003 (64-bit) for Itanium-based systems Enterprise Edition or Datacenter Edition (SP2 required for both)

Table 4-2 Supported Microsoft SQL Server 2008 versions (Continued)

SQL Server 2008	Windows Servers
<p>Microsoft SQL Server 2008, 32-bit Standard Edition, Enterprise Edition, or Web Edition on Windows Server 2008 (SQL Server 2008 SP1)</p>	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, Datacenter Edition, or Web Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008
<p>Microsoft SQL Server 2008, 64-bit Standard Edition, Enterprise Edition, Enterprise IA64 Edition, Web Edition on Windows Server 2008 (SQL Server 2008 SP1)</p>	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 64-bit Itanium (IA64) ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

The following table lists the Microsoft SQL Server 2008 R2 versions supported with SFW HA 5.1 Service Pack 2.

Table 4-3 Supported Microsoft SQL Server 2008 R2 versions

SQL Server 2008 R2	Windows Servers
Microsoft SQL Server 2008 R2, 32-bit Standard Edition, Enterprise Edition, Datacenter Edition on Windows Server 2003	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)
Microsoft SQL Server 2008 R2, 64-bit Standard Edition, Enterprise Edition, Datacenter Edition on Windows Server 2003	<ul style="list-style-type: none"> ■ Windows Server 2003 (64-bit) Standard x64 Edition, Enterprise x64 Edition or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) ■ Windows Server 2003 (64-bit) for Itanium-based systems Enterprise Edition or Datacenter Edition (SP2 required for both)

Table 4-3 Supported Microsoft SQL Server 2008 R2 versions (Continued)

SQL Server 2008 R2	Windows Servers
Microsoft SQL Server 2008 R2, 32-bit Standard Edition, Enterprise Edition, Datacenter Edition on Windows Server 2008	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, Datacenter Edition, or Web Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008
Microsoft SQL Server 2008 R2, 64-bit Standard Edition, Enterprise Edition, Datacenter Edition on Windows Server 2008	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 64-bit Itanium (IA64) ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

Storage requirements and best practices

Review the following best practices and requirements for snapshot storage:

- The hardware for Quick Recovery should include sufficient storage to be able to create dynamic volumes on separate disks or LUNs for the following purposes:
 - Databases
 - Transaction logs
 - FILESTREAM filegroups (for SQL Server 2008)
 - Split-mirror snapshots of the database and transaction log volumes
- For the split-mirror snapshots, make sure to select disks or LUNs that are not used for production data.
- You can create a snapshot of more than one volume on the same disk or LUN as long as there is sufficient space available.

Configuration requirements and best practices

Review the following configuration requirements and best practices:

- The system and boot volumes must reside on a separate disk (Harddisk0) from the dynamic volumes used for the SQL user-defined databases and split-mirror snapshots.
- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. Upgrade any disk groups created using an earlier version of Volume Manager for Windows before creating Quick Recovery snapshots.
- 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 on Microsoft Disk Management Disk Groups, see *Veritas Storage Foundation Administrator's Guide*.
- Database, transaction logs, and FILESTREAM filegroups (for SQL Server 2008) must be stored on disks within a single dynamic disk group.
- Database, transaction logs, and FILESTREAM filegroups (for SQL Server 2008) should be on separate disks so that disk failure does not affect anyone of these filegroups.
- User-defined database, transaction logs, and FILESTREAM filegroups (for SQL Server 2008) may not be stored in the same volume as the SQL Server program files or system data files.
- In order to perform a roll-forward recovery to the point of failure, database, transaction logs, and FILESTREAM filegroups (for SQL Server 2008) must be in separate volumes.

- Locate snapshot volumes on separate disks from any database, log, and FILESTREAM filegroup (for SQL Server 2008) volumes so that the snapshot process will not interfere with database operations.
- 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.

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

- 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 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.

Reviewing the configuration

A sample setup containing the following object names is used to illustrate how to create, refresh, and restore a snapshot set for a specified SQL Server database.

Table 4-5 Object names used in SQL Server Quick Recovery tasks

Name	Drive Letter	Object
billing		SQL Server instance name
billing_DB		component (SQL Server database name)
harddisk3		Disk where the mirrors will be created
billing_DG		disk group name for the Billing disk group
billing_data	L:	volume for Billing database user data
billing_log	M:	volume for Billing database log

Table 4-5 Object names used in SQL Server Quick Recovery tasks (Continued)

Name	Drive Letter	Object
billing_FS	N:	volume for Billing FILESTREAM filegroups (if FILESTREAM is implemented)
billing_datasnap	S:	volume for Billing database user data snapshot
billing_logsnap	T:	volume for Billing database log snapshot
billing_FSsnap	U:	volume for FILESTREAM filegroups database snapshot (if FILESTREAM is implemented)
billing_DB.xml		name for the metadata file generated when the snapshot is created
tlog1.bak, tlog2.bak		log files saved at c:\backup prior to restoring

Configuring SQL Server storage with Veritas Storage Foundation for Windows

To use the SFW Quick Recovery snapshot functionality with SQL database components, you must store the database components on Veritas Storage Foundation for Windows (SFW) dynamic volumes. Configuring your SQL Server storage with SFW includes the following tasks:

- Creating one or more dynamic disk groups
- Creating volumes for the databases and transaction logs
- Creating a new database and set the appropriate paths to point to the new SFW volumes.

If your SQL Server environment is already configured with SFW, skip this section.

Creating dynamic disk groups

Create one or more dynamic disk groups.

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.

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, `billing_DG`).
- 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:

```
> vxdbg -gbilling_DG init Harddisk1 Harddisk2 Harddisk3
```

where `-gbilling_DG` 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.

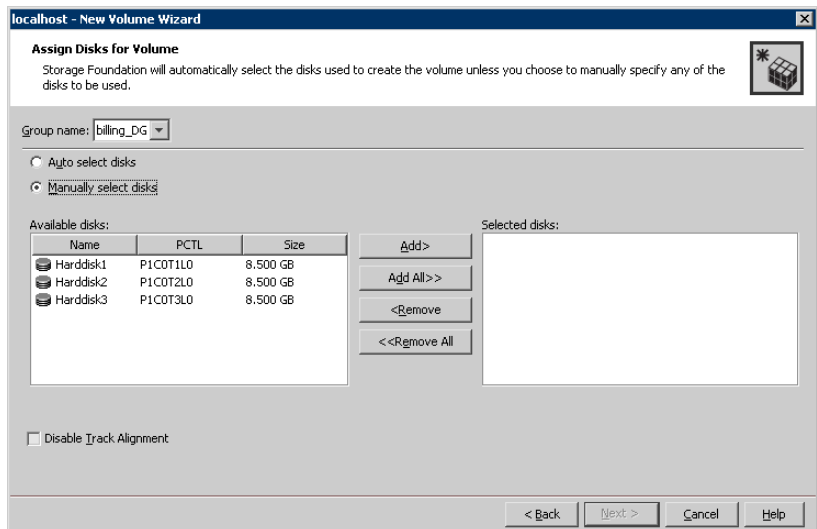
For the complete syntax for the `vxdg init` command, see *Veritas Storage Foundation Administrator's Guide*.

Creating dynamic volumes

Create volumes for the database and transaction log.


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, `billing_DG`), 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 (for example, `billing_DG`). 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, `billing_data`).
 - Enter the size.
 - Select the layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 7 Assign a drive letter to the volume (for example, L: for the `billing_data` volume) and click **Next**.
- 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 transaction log (for example, `billing_log`) and any other database in your configuration.

To create a volume from the command line

- 1 Type the CLI commands:

```
> vxassist [-b] -gbilling_DG make billing_data L:  
> vxassist [-b] -gbilling_DG make billing_log M:
```

This will create volumes in the `billing_DG` disk group named `billing_data` and `billing_log` on drive letters L and M respectively.

- 2 Modify and repeat this command as necessary to create additional volumes.

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

Pointing the databases and log paths to the SFW volumes

Create a new database and point the database and log paths to the SFW volumes.

Refer to your Microsoft SQL Server documentation for information about creating a new database and setting the database and log paths.

Implementing Quick Recovery for SQL Server 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 provides the most complete method of configuring snapshot sets and is therefore recommended for initial implementation. 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 in a SQL Server instance. Each database can have different filegroups, such as Database, Log, and FILESTREAM filegroups (for SQL Server 2008).

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 time scheduled for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again.

The snapshot creation process integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. The snapshot is done while the database is online and without disrupting the database operations.

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 database to snapshot
- Number of snapshot sets for each database
- 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 (Copy or Full)

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 or instances. Even if you do not save the settings to a template file, they are still stored for the databases to which they apply. You can access them later by selecting the same instance and database in the wizard.
- Implement the settings immediately.

Backup types for snapshot sets

When creating a snapshot set, you can specify that the snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

About snapshot templates

A snapshot template contains all the settings that you implemented with the Quick Recovery Configuration Wizard for selected databases of a particular instance. After configuring settings with the wizard, you can save the template settings to a template (.tpl) file (an XML file) for reuse on another database or instance. 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 database.

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 a database, you cannot select that database and import a template to overwrite those settings. You must instead modify the settings. However, you can import a template for another database for which settings do not exist.

Templates store the following information:

- The selected database
- The number of snapshot sets for the database
- When to prepare the snapshot mirrors
- The snapshot schedule and rules related to the schedule, including backup type (but not the names of optional scripts)
- The current date and time when the template is created

Templates and multiple components

When you apply a template that has multiple components, the wizard first attempts to match the names of components (databases) in the template to the databases you selected in the wizard.

If it matches a name, it applies the information stored under that name in the template to that database. For example, if you select the Billing database in the wizard and the template has settings for Billing, it applies those settings.

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

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

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.

See "[Scheduling and creating snapshot sets](#)" on page 52.

Tasks for implementing snapshot sets with the configuration wizard

Table 5-1 outlines the high-level objectives and the tasks to complete each objective.

Table 5-1 Tasks for implementing snapshot sets with the configuration wizard

Objective	Tasks
“Reviewing the prerequisites” on page 52	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Scheduling and creating snapshot sets” on page 52	<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 61	<ul style="list-style-type: none"> ■ Viewing the status of scheduled snapshots to determine whether they were successful
“Deleting or modifying schedules” on page 64	<ul style="list-style-type: none"> ■ Deleting schedules that are no longer needed or modifying schedules

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 and review the recommendations and best practices.
See [“Reviewing the prerequisites”](#) on page 34.
- Set up SQL Server for use with Storage Foundation for Windows (SFW).
See [“Configuring SQL Server storage with Veritas Storage Foundation for Windows”](#) on page 41.
- 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 scheduler service is configured with domain administrator privileges.
- 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.
- Although not required to run the wizard, the SQL Server VSS Writer service must be started for snapshot mirror preparation and scheduled snapshots to occur.

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. The wizard also creates the snapshot mirrors required for the snapshots.

See [“About the Quick Recovery Configuration Wizard”](#) on page 47.

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

See [“Reviewing the prerequisites”](#) on page 52.

In addition, if you plan to reuse settings, you should be familiar with the information about snapshot templates.

See “[About snapshot templates](#)” on page 49.

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	Select the SQL instance for which you want to configure the snapshot sets. The selection details are displayed in the panel.
---------------------------	---

Select objects or view details	Select the SQL databases for which you want to configure the snapshot sets. Databases with names longer than 50 characters are not supported by the wizard.
--------------------------------	--

The wizard checks for existing schedules for the selected database. 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	<p>Create or modify the Quick Recovery settings for the selected components.</p> <p>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.</p> <p>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.</p>
Import a template	<p>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.</p> <p>This option is available only if no settings exist for a selected component.</p>

- 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	<p>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.</p> <p>If you are modifying an existing schedule, the snapshot set name cannot be changed.</p>
--------------	--

XML Metadata File Name	<p>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.</p> <p>If you are modifying an existing schedule, the XML file name cannot be changed.</p>
Mirror Preparation Time	<p>Click Edit and in the Mirror Preparation dialog box, specify the date and time for mirror preparation and click OK.</p> <p>Preparing the snapshot mirrors slows down performance, so you may want to choose a time of low usage.</p>
Snapshot Disks	<p>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.</p>
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 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.

For a non-clustered environment, the default location on Windows Server 2003 is:

```
C:\Documents and Settings\All Users\  
Application Data\VERITAS\VSSXML\  
application name
```

and on Windows Server 2008:

```
SystemDrive:\ProgramData\VERITAS\VSSXML\  
application name
```

If a `redirect.txt` file has been created, the path in that file is shown in the field.

For new schedules for a clustered server environment, if there is no `redirect.txt` file, the field is empty. Enter the full path to a location on shared storage, for example: `G:\SnapshotSets`.


You must store the XML files on a separate volume from the volumes that are included in snapshots.

Note: If the snapshot XML files are not stored separately from the volumes that are included in the snapshots, a restore will fail.

Ensure that you use a consistent location. That way, you can easily find the XML metadata file, which is required for recovery.


- 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:

You can specify that snapshots be created as either a Full backup or Copy backup type. Any snapshots taken between the time a full backup is taken and the time a log backup is taken can be used as a base to restore and replay the logs. If the replay of logs is not necessary, then either type of snapshot can be used to restore the database to the time of the snapshot set. A Copy backup does not affect the sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

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 on other SQL databases, 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 database.

- 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.
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. For SQL, you must manually start the SQL Server VSS Writer Service. Then right-click the Storage Agent node in the VEA and click Refresh.

Situation

Resolution

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.

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:

- 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.

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.

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 2003:

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

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.

See “[Scheduling and creating snapshot sets](#)” on page 52.

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 for SQL Server

This chapter covers the following topics:

- [About scheduling or creating an individual snapshot set](#)
- [Tasks to schedule a new snapshot set](#)
- [Tasks to create a one-time snapshot set](#)
- [Reviewing the prerequisites](#)
- [Preparing the snapshot mirrors](#)
- [Scheduling a new snapshot set](#)
- [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 5, “Implementing Quick Recovery for SQL Server with the configuration wizard”](#) on page 47.

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

Note: Adding a snapshot set using the VSS 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.

At times you may want to create a one-time snapshot of a specific volume or volumes. You can do so using either the vxsnap command line utility or from the Veritas Enterprise Administrator (VEA) console using the VSS Snapshot and Snapback wizards.

If you want to snapshot only a single volume rather than multiple volumes in the database, you can use the VEA Snapshot Volume Wizard rather than the VSS SQL Snapshot Wizard. See the *Storage Foundation for Windows Administrator's Guide*.

Tasks to schedule a new snapshot set

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

Table 6-1 Tasks for scheduling a new snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 70	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 70	<ul style="list-style-type: none"> ■ Creating snapshot mirrors using the using the VEA Prepare command or the vxsnap utility
“Scheduling a new snapshot set” on page 72	<ul style="list-style-type: none"> ■ Using the VSS Snapshot Scheduler Wizard to create the initial snapshot set and set up the schedule for keeping it refreshed.

Tasks to create a one-time snapshot set

[Table 6-2](#) outlines the high-level objectives and the tasks to complete each objective.

Table 6-2 Tasks for creating a one-time snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 70	■ Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 70	■ Creating snapshot mirrors using the VEA Prepare command or the vxsnap utility
“Creating a one-time snapshot set” on page 82	■ Creating the one-time snapshot set using the VEA or the vxsnap utility.

Reviewing the prerequisites

Review the following prerequisites:

- Ensure that your system hardware and software meets the requirements. See [“Reviewing the prerequisites”](#) on page 34 in [Chapter 4, “Preparing to implement Quick Recovery for SQL Server”](#).
- Set up SQL Server for use with Storage Foundation for Windows (SFW). See [“Configuring SQL Server storage with Veritas Storage Foundation for Windows”](#) on page 41 in [Chapter 4, “Preparing to implement Quick Recovery for SQL Server”](#).
- Ensure that the SQL Server VSS Writer service is started.
- 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 database volume.
- The `vxsnap prepare` command from the CLI
Use the `vxsnap prepare` command to prepare a mirror for each of the volumes associated with the database.
- The `vxassist prepare` command from the CLI
The `vxassist prepare` command can be used to prepare the volume in a similar manner to the VEA GUI `snap prepare` option.

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

Also ensure that the SQL Server VSS Writer service is started.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap** and then **Prepare**.
- 2 In the wizard, click **Next**.
- 3 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 **Next**.

Or choose **Auto select disks** and any available disks will be selected.

Next, **Specify the attributes** and click **Next**.

Finally, **Click Finish** to prepare volume for flashsnap.

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

Select existing mirror for snap and then click **Next**.

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, select **Do not use existing mirror for snap** and click **Next**, then click **Select Disk**. Use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box, and click **Next**.
- To use an existing mirror, click **Select existing mirror for snap**, select the desired mirror, and click **Next**.

- 4 Repeat [step 1](#) and [step 3](#) to create a snapshot mirror for each volume associated with the database.

Note: 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 vxsnap prepare command

- ◆ Type the command, as in the following example:

```
> vxsnap prepare component=billing_DB/writer=SQLServerWriter
source=L:/harddisk=harddisk3
source=M:/harddisk=harddisk3
```

In this example, snapshot mirrors for the database and log volumes of the `billing_DB` database are created. The mirror for the volume mounted at L: is created on disk 3 and the mirror for the volume mounted at M: is also created on disk 3.

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

```
vxsnap prepare component=ComponentName/writer=WriterName  
[server=ServerName] [instance=InstanceName] [-b]  
[source=volume/harddisk=harddisk...]
```

Scheduling a new snapshot set

Before you run the VSS Snapshot Scheduler Wizard to schedule a snapshot set for a database, you must prepare a snapshot mirror for each of the volumes in the database.

See “[Preparing the snapshot mirrors](#)” on page 70.

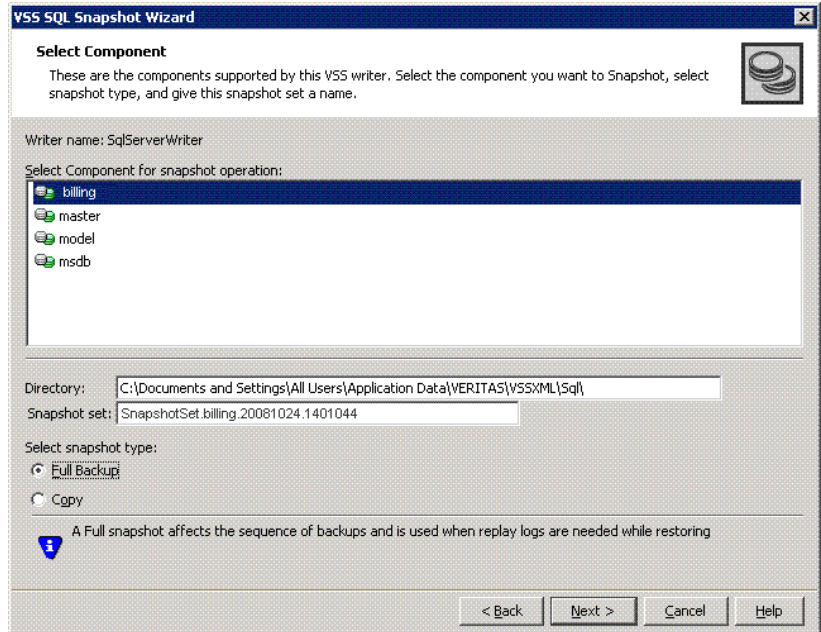
You can then use the VSS Snapshot Scheduler Wizard to schedule the initial snapshot set and to set up the schedule for keeping it refreshed.

When the scheduled snapshots occur, the snapshot mirrors are detached from the original volumes, creating separate on-host snapshot volumes as well as an XML file to store the SQL 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 database operations.

To schedule a snapshot for a selected component

- 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**. If the SQL node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQL** node.
- 4 Right-click the instance, and click **Schedule VSS Snapshot**.
- 5 In the Welcome panel, review the information and click **Next**.

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



Select component for snapshot operation	Select the database for the snapshot set.
Directory	<p>The XML file is stored by default in the directory shown on the screen.</p> <p>In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.</p> <p>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.</p> <p>There are two ways to change the XML file location.</p> <p>First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named <code>redirect.txt</code>. This text file should contain a single text line specifying the full path to the location of the XML file, for example, <code>G:\BackupSets</code>. Save the <code>redirect.txt</code> file in the default directory <code>C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML</code>.</p> <p>If SFW/SFW HA is installed on the D drive, then the path will be the same as above, but on the D drive</p> <p>Note: You must not use the volume name or volume path in the <code>redirect.txt</code> file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.</p>
Snapshot set	<p>Enter a name for the snapshot set, for example, <code>billing</code> or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time.</p> <p>The wizard creates the snapshot set metadata XML file with this name.</p>

Select snapshot type Select the snapshot type.

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. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

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

Volume Name Displays the name of the volume.

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.

- 8 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*.

- 9 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. If no interval is specified the snapshot occurs once.
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.
Post Command	Optionally, specify the full path of a command script to run after the snapshot is complete.

10 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.

- 11 Click **Next**.
- 12 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, the Storage Agent node, and the **Applications** node.
- 3 There are two ways to view the scheduled snapshot status:
 - Select the Scheduled Tasks tab. You can view the status of schedules and delete schedules from this pane. View the scheduled snapshots that are listed in the pane on the right
 - Expand the component for which you scheduled the snapshot. Click **Scheduled Tasks**. The scheduled snapshots are listed in the pane on the right.
- 4 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Scheduled Tasks** and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 5 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. For SQL, you must manually start the SQL Server VSS Writer Service. Then right-click the Storage Agent node in the VEA and click Refresh.

Situation	Resolution
Drive letters assignments for snapshot volumes do not occur.	<p>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.</p> <p>You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.</p>
A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.	<p>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.</p> <p>To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.</p>
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.	<p>You can choose from the following ways to avoid these problems:</p> <ul style="list-style-type: none">■ 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.
<p>The following problems can occur:</p> <ul style="list-style-type: none">■ 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.	

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 2003:

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

For Windows Server 2008:

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

Deleting or modifying schedules

You can delete a schedule from the VEA console, 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, the Storage Agent node, and the **Applications** node.
- 3 Select the Scheduled Tasks tab. You can view the status of schedules and delete schedules from this pane.
- 4 View the scheduled snapshots that are listed in the pane on the right.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Creating a one-time snapshot set

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

- The first step is to prepare the snapshot mirrors for the database volume or volumes. 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 70.
- The second step uses either the VSS SQL Snapshot Wizard or the `vxsnap create` command 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 a snapshot set XML file to store the application and snapshot volume metadata.


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

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

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. If the SQL node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQL** node.
- 4 Right-click the instance and click **VSS Snapshot**.
- 5 In the wizard, review the Welcome page and click **Next**.

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



Select Component for snapshot operation	Select the database for the snapshot set.
Directory	<p>The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory shown on the screen.</p> <p>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.</p> <p>There are two ways to change the XML file location.</p> <p>First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named <code>redirect.txt</code>. This text file should contain a single text line specifying the full path to the location of the XML file, for example, <code>G:\BackupSets</code>. Save the <code>redirect.txt</code> file in the default directory <code>C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML</code>.</p> <p>If SFW/SFW HA is installed on the D drive, then the path will be the same as above, but on the D drive.</p> <p>Note: You must not use the volume name or volume path in the <code>redirect.txt</code> file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore will fail.</p>
Snapshot set	Enter a name for the snapshot set, for example, <code>billing</code> or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time.
Select snapshot type	<p>Select the snapshot type.</p> <p>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. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.</p>

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

Volume Name	Displays the name of the volume.
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.

- 8 On 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*.

- 9 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 billing_DB.xml create
  source=L:/Newvol=billing_data
  source=M:/Newvol=billing_log
  writer=SQLServerWriter component=billing_DB
  backuptype=full
```

where `billing_DB.xml` is a name you assign the metadata file that is used to reattach the snapshot set or to recover using that snapshot set.

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

```
vxsnap -x filename create [-gDynamicDiskGroupName]
source=Volume [/DriveLetter=DriveLetter]
[/DrivePath=DrivePath] [/Newvol=NewVolName] [/Plex=PlexName] ...
writer=WriterName component=ComponentName
[server=ServerName] [instance=InstanceName]
[backuptype=FULL|COPY] [-O]
```

The *WriterName* and *ComponentName* are required. The component name is the name of the SQL database. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. You can use either a Full

backup or Copy backup for restoring from a snapshot (recovery). However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

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

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 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 86
- [“Creating the refreshed snapshot set”](#) on page 87

Note: The VSS Refresh option available in the VEA console from the VSS 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 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 database application 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 billing_DB.xml reattach  
writer=SQLServerWriter
```

where `billing_DB.xml` is the name of the XML metadata file created at the time the snapshot set was created.

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

```
vxsnap -x Filename [-f] [-b] reattach writer=WriterName
```

Creating the refreshed snapshot set

After you have reattached and resynchronized the snapshot set mirrors with the original volumes, using the VSS Snapback wizard or from the command line, create a new snapshot set of the database using either the VSS SQL Snapshot Wizard or the `vxsnap create` command.

See “[Creating a one-time snapshot set](#)” on page 82.

Recovering a SQL Server database

This chapter covers the following topics:

- [About recovering a SQL Server database](#)
- [Tasks for recovering a SQL Server database](#)
- [Prerequisites for recovering a SQL Server database](#)
- [Types of recovery](#)
- [Recovering using snapshots without log replay](#)
- [Recovering using snapshots and log replay](#)
- [Restoring snapshots and manually applying logs](#)
- [Recovering missing volumes](#)
- [Post-recovery steps](#)
- [Vxsnap restore command reference](#)

About recovering a SQL Server database

You can use the on-host Quick Recovery snapshot set to quickly recover a database after logical corruption.

You can also use a snapshot set to recover a database after production volumes are lost due to hardware failure. This recovery assumes that the failure does not affect the disk or disks where the snapshot set volumes are located.

You can use either the VSS SQL Restore Wizard from the Veritas Enterprise Administrator (VEA) console or the `vxsnap restore` command. Both of these methods integrate with VSS to notify the SQL 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.

Note: Refer to the Troubleshooting chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for troubleshooting information.

Tasks for recovering a SQL Server database

Table 7-1 outlines the high-level objectives and the tasks to complete each objective. The particular task required depends on the type of recovery you are using.

Table 7-1 Tasks for recovery using Quick Recovery snapshots

Objective	Tasks
“Prerequisites for recovering a SQL Server database” on page 91	■ Verifying the prerequisites for using VSS recovery
“Types of recovery” on page 91	■ Understanding the types of recovery
“Recovering using snapshots without log replay” on page 92	■ Recovering to a specified point in time using the VSS SQL Restore wizard or the vxsnap utility
“Recovering using snapshots and log replay” on page 94	■ Recovering to a point of failure using the SQL Restore wizard or the vxsnap utility
“Restoring snapshots and manually applying logs” on page 97	■ Recovering to the time of the snapshot set using the SQL Restore wizard or the vxsnap utility
“Recovering missing volumes” on page 99	■ Recovering missing production volumes from snapshot volumes, using the VSS SQL Restore wizard or the vxsnap utility
“Post-recovery steps” on page 105	■ Refreshing the snapshot set ■ Performing additional tasks in a VVR environment

Prerequisites for recovering a SQL Server database

You can use the VSS SQL Restore Wizard or vxsnap command line utility to recover a database from a snapshot set. Both the snapshot set and the snapshot set XML metadata file must be available.

If you are planning to recover missing volumes after hardware failure, additional prerequisites are described in the following topic:

“[Preparing for the recovery](#)” on page 100.

Types of recovery

[Table 7-2](#) gives an overview of the options you can select for recovery. The options are available from either the VSS SQL Restore Wizard or the vxsnap restore command.

Table 7-2 Recovery options

Selected option	Database state after recovery	Description
Recovery	online	Restore to the time of the snapshot set. Database and transaction log volumes are restored from the specified snapshot set database and log volumes. No additional transaction logs are applied.
Recovery, along with the option to restore missing volumes	offline	Restore to the time of the snapshot set. The specified snapshot set volumes are converted from read-only snapshot volumes to read/write volumes. No additional transaction logs are applied. You must then manually assign the drive letter/mount path of the missing production volumes to the previous snapshot volumes and bring the database online.

Table 7-2 Recovery options (Continued)

Selected option	Database state after recovery	Description
Recovery + log replay	online	<p>Restore the snapshot set and automatically roll-forward to the point of failure.</p> <p>Restores the database and transaction log volumes from the specified snapshot set and then automatically replays the specified logs to roll forward the recovery to the point of failure.</p> <p>Before using this option, you must back up your transaction logs within SQL Server. This log operation requires that at least one Full backup was created earlier.</p>
No Recovery	loading	<p>Restore the snapshot set and then manually apply logs in SQL.</p> <p>Restores the database and transaction log volumes from the specified snapshot set and leaves the database in a loading state. To bring the database back to an operational state, you must manually apply your backup transaction logs within SQL Server to the desired point in time.</p> <p>Before using this option, you must back up your transaction logs within SQL Server. This log operation requires that at least one Full backup was created earlier.</p>

Recovering using snapshots without log replay

The following procedure uses the Recovery option without the option to restore missing volumes. It restores the database from the snapshot set volumes to the time of the snapshot set. The database and transaction log volumes are restored but no additional transaction logs are applied.

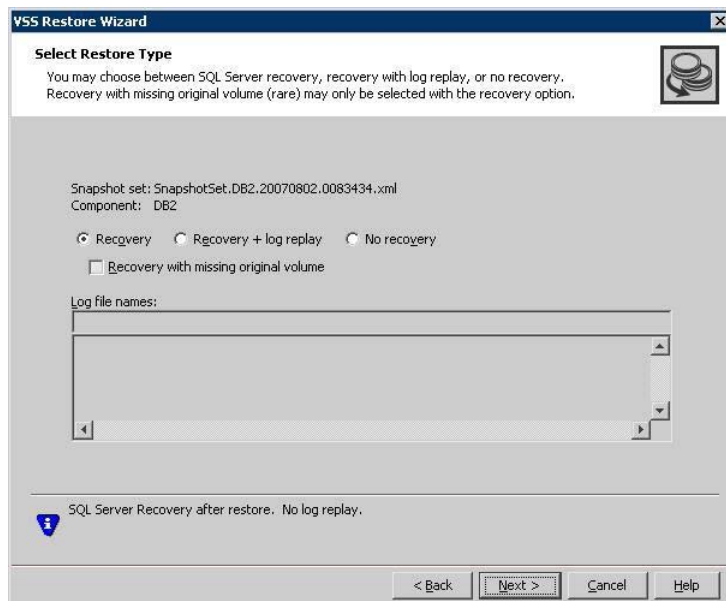
The Recovery option can also be used to recover a database for a missing volume. For the procedure to recover a missing volume, see the following topic: [“Recovering missing volumes”](#) on page 99

To recover a database without log replay using the VEA

- 1 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **Applications node**.
- 4 Right-click **SQL** and click **VSS Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation 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 logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the Select Restore Type panel, click **Recovery**.



- 8 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.
The database is restored to the time the snapshot set was created or last refreshed. If you took it offline earlier, bring it back online.
- 10 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To recover without log replay using the `vxsnap restore` command

- 1 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 Type the command as in the following example:

```
vxsnap -x billing_DB.xml restore RestoreType=RECOVERY noLogs
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command. The volumes in the snapshot set are restored and the database is left in an operational state. The database is restored to the time the snapshot set was created or last refreshed. If you took it offline earlier, bring it back online.
- 3 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Recovering using snapshots and log replay

The following procedure restores the database from the snapshot set volumes and applies the backed up transaction log files that you specify to automatically roll forward the recovery to the point of failure.

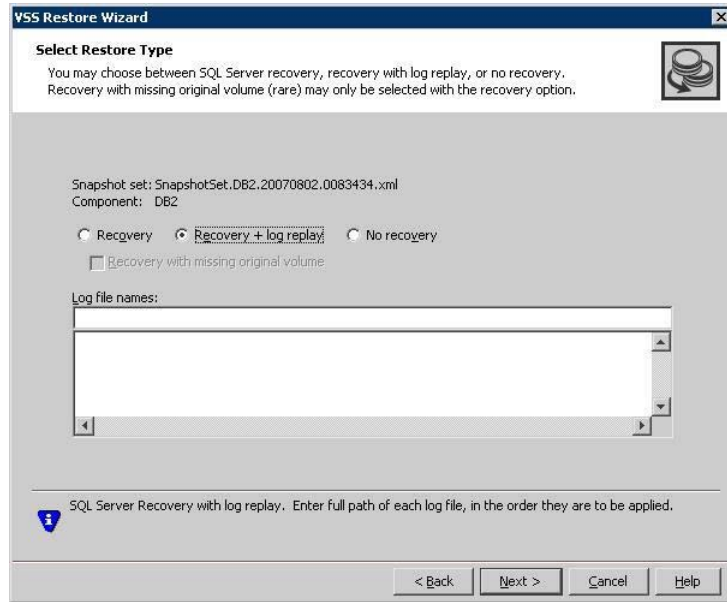
You can choose the procedure for the VSS SQL Restore Wizard or for the `vxsnap restore` command.

Warning: Before you begin, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

To use log replay for an automatic roll-forward recovery to a point of failure using the VEA

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **Applications** node
- 5 Right-click **SQL** and click **VSS Restore**.
- 6 Review the Welcome page and click **Next**.
- 7 Select the snapshot set XML metadata file to be used for this operation 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 logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 8 On the Select Restore Type panel, do the following and click Next:
 - Click **Recovery + Log replay**.

- Enter the full path of each log file, in the order they are to be applied.



- 9 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.
- 10 Verify the restore specifications and click **Finish**.
After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state. If you took it offline earlier, bring it back online.
- 11 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To use log replay for an automatic roll-forward recovery to a point of failure using the `vxsnap restore` command

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.

- 3 Type the command as in the following example:

```
vxsnap -x billing_DB.xml restore RestoreType=RECOVERY
      logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

 where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command and `c:\backup\tLog1.bak`, `c:\tLog2.bak` are the paths to the transaction log backup files. After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state. If you took it offline earlier, bring it back online.
- 4 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Restoring snapshots and manually applying logs

The following procedure uses the No Recovery option to restore the database from the database and log snapshot volumes. Selecting this option leaves the database in a loading state. You can then manually apply backed up transaction logs to recover the database to the desired point in time.

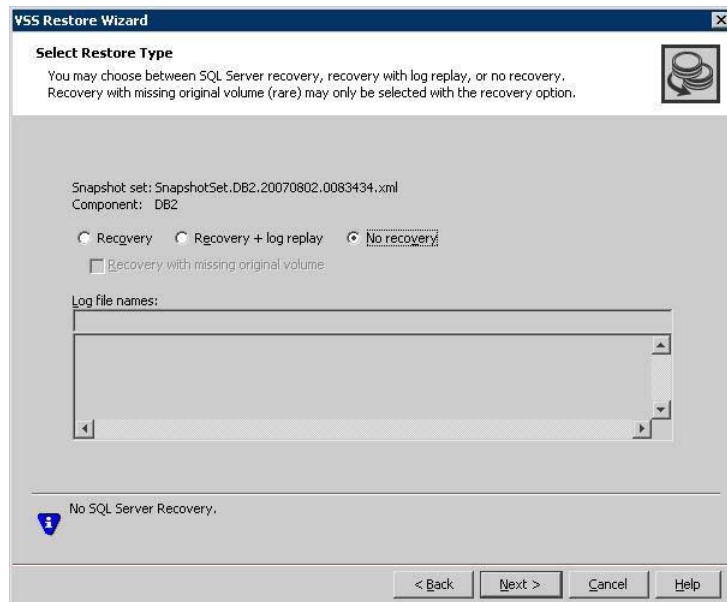
You can use either the VSS SQL Restore Wizard from the VEA or the `vxsnap restore` command.

Warning: Before you begin, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

To restore using the No Recovery option in the VEA

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the database volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **Applications** node
- 5 Right-click **SQL** and click **VSS Restore**.

- 6 Review the Welcome page and click **Next**.
- 7 Select the snapshot set XML metadata file to be used for this operation 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 logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 8 On the Select Restore Type panel, click **No Recovery** and click **Next**.



- 9 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 10 Verify the restore specifications and click **Finish**.
The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.
- 11 Use your preferred method to manually restore the backup transaction logs to the desired point in time and then bring the database back online.
- 12 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To restore using the No Recovery option with the `vxsnap restore` command

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles or third-party system management tools that may be accessing the database volumes. It is also recommended to bring the database offline.
- 3 Type the command as in the following example:

```
> vxsnap -x billing_DB.xml restore RestoreType=NO_RECOVERY
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command.
The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.
- 4 Use your preferred method to manually restore the backup transaction logs to the desired point-in-time and bring the database online.
- 5 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Recovering missing volumes

Typically missing volumes can result from a hardware failure. If the failure does not affect the disk or disks where the snapshot set resides, you can recover from the failure by restoring the missing volumes from the snapshot set.

Before performing the recovery, you complete several steps to prepare for it.

See “[Preparing for the recovery](#)” on page 100.

You can then recover the missing volumes using the VSS Restore Wizard or the `vxsnap` utility.

See “[Performing the recovery](#)” on page 103.

If the original volume is not present, the snapshot volume is changed from a read-only volume to a read-write volume. No additional transaction logs are applied.

Preparing for the recovery

Before using the VSS Restore Wizard or the vxsnap utility to recover missing production volumes from a snapshot set, you must prepare for the recovery. To prepare, you complete the following steps:

- Take any affected databases offline.
- Reassign the drive letters or mount points of the missing volumes to the snapshot volumes.
See [“Reassigning the drive letter or mount points of the missing volumes to the snapshot volumes”](#) on page 100.
- 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 102.

Reassigning the drive letter or mount points of the missing volumes to the snapshot volumes

If the production volume was 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.

There are two stages to the procedure:

- Note down the assigned drive letter or mount path of the production volume so that you can reassign the same drive letter or mount path to the snapshot volume. Then remove the existing drive letter or mount path from the production volume.
- Change the drive letter or mount path of the snapshot volume to the production volume drive letter or mount path.

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 remove the drive letter or mount path from the production volume

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

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

- 1 Right-click the snapshot volume, click **File System > Change Drive Letter and Path**.
- 2 On the Drive Letter and Paths screen, 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 On the Drive Letter and Paths screen, 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 On the Drive Letter and Paths screen, 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 Navigate to the folder where the production volume was mounted.
- 9 Click **OK**.

- 10 Click **OK** to assign the mount point.

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.

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.
- 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 storage group. 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.

Performing the recovery

Ensure that you have prepared for the recovery first.

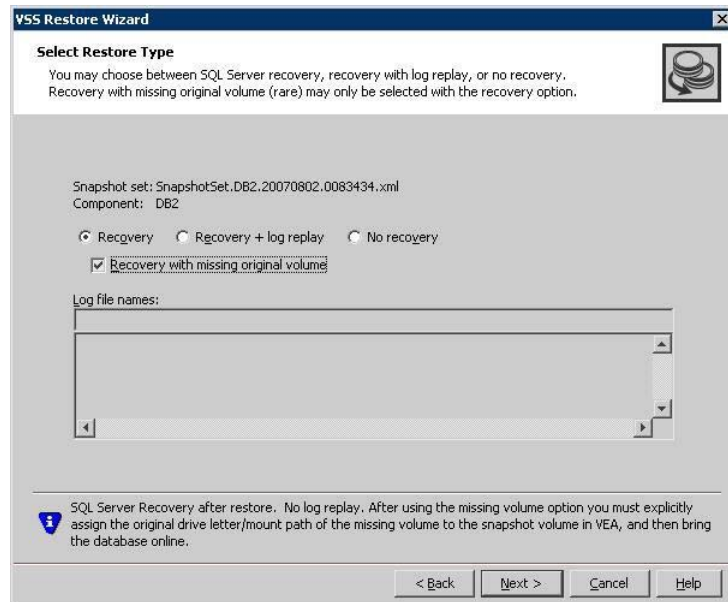
See “[Preparing for the recovery](#)” on page 100.

To recover a database with one or more missing volumes using the VEA

- 1 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click **SQL** and click **VSS Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation 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 logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the Select Restore Type panel, click **Recovery** and select **Recovery with missing original volume**.



- 8 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.
The snapshot of the missing volume is changed from a read-only volume to a read-write volume.
- 10 If you have not already done so, in the VEA, ensure that the drive letter or mount path of the missing production volume is assigned to the snapshot volume.
- 11 Bring the database online.
If the production volume was missing, the snapshot volume is now changed to the production volume. The database is restored to the time the snapshot set was created or last refreshed.
- 12 To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To recover a database with one or more missing volumes using the `vxsnap restore` command

- 1 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 Type the command as in the following example:

```
vxsnap -x billing_DB.xml -r restore RestoreType=RECOVERY noLogs
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command.
The snapshot of the missing volume is changed from a read-only volume to a read-write volume.
- 3 If you have not already done so, in the VEA, reassign the drive letter or mount path of the missing volume to the snapshot volume.
- 4 Bring the database online.
If the production volume was missing, the snapshot volume is now changed to the production volume. The database is restored to the time the snapshot set was created or last refreshed.
- 5 To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Post-recovery steps

After you have performed any of the recovery methods, whether point-in-time or roll forward, create a new snapshot set.

See “[Creating a one-time snapshot set](#)” on page 82 in [Chapter 6, “Scheduling or creating an individual snapshot set for SQL Server”](#)

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.

Vxsnap restore command reference

The following is a summary of the vxsnap restore command syntax and attribute descriptions.

The command syntax is as follows:

```
vxsnap -x Filename [-b] [-f] [-r] restore  
{RestoreType=[RECOVERY|NO_RECOVERY]}  
[noLogs|logFiles=tlog1,tlog2,...] writer=WriterName
```

The following attributes apply:

-x <i>Filename</i>	The metadata file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file.
-b	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
-f	Forces the operation. Make sure the volume is not in use before using this option.
-r	Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume. Use this option only with <i>Recovery noLogs</i> . After using this option you must explicitly assign the original drive letter/mount path of the missing volume to the snapshot volume in the VEA and then bring the database online.
RestoreType= [RECOVERY NO_RECOVERY]	Specifies the type of database recovery, either recovery or no recovery: <i>RECOVERY</i> can be used with either the <i>noLogs</i> or <i>logFiles=tlog1,tlog2,...</i> attributes. <i>NO_RECOVERY</i> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in an loading state so that you can manually replay backup logs to a specific point in time.

<i>noLogs</i>	Database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.
<i>logFiles=tlog1,tlog2,...</i>	Transaction log backup files to be applied with the RECOVERY option to achieve a point of failure recovery. Each transaction log must have a unique name and be created using the “overwrite existing media” option within SQL Server. The database is left in an online state.
<i>writer=WriterName</i>	The name for the SQL Server VSS Writer; used to located the default directory to search for the XML metadata file. Specify SQLServerWriter .

Vxsnap utility command line reference for SQL Server

This chapter covers 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 Windows Volume Shadow Copy Service (VSS) to create split-mirror snapshots of all the volumes in the specified component.

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.

- vxsnap is used for SQL 2008 and SQL 2005; SQL 2000 uses a different version called vxsnapsql.

Vxsnap keywords

The vxsnap utility has the following keywords:

prepare	Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes. Note: Either the <code>prepare</code> or <code>start</code> keyword may be used in the CLI; however <code>prepare</code> is recommended.
create	Creates simultaneous snapshots of all volumes in the specified component, providing a point-in-time snapshot set.
reattach	Reattaches and resynchronizes an existing snapshot set to the original database volumes.
restore	Restores a SQL database from a snapshot set.

You can view an online description of the command syntax by typing the following:

```
vxsnap keyword -?
```

vxsnap prepare

Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes.

Syntax

```
vxsnap prepare component=ComponentName/writer=WriterName  
[server=ServerName] [instance=InstanceName] [-b]  
source=volume/harddisk=harddisk...
```

Attributes

The following attributes apply:

<code>component=ComponentName</code>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<code>writer=WriterName</code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.
<code>server=ServerName</code>	SQL Server server name.
<code>instance=InstanceName</code>	SQL Server instance name.
<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, harddisk2.

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

Example

```
vxsnap prepare component=billing_DB/writer=SQLServerWriter  
instance=SQLINSTANCE server=SQLVirtualServer  
source=L:/harddisk=harddisk3 source=M:/harddisk=harddisk3
```

This command will create snapshot mirrors of all the volumes contained in the billing_DB database. The snapshot mirror of the volume mounted on L: will be created on harddisk 3 and the snapshot mirror of the volume mounted on M: will also be created on harddisk 3.

vxsnap create

Creates snapshot(s) of the specified volume(s) or SQL database (component). Allows volumes to be snapshotted simultaneously.

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.

Syntax

```
vxsnap -x Filename create [-gDynamicDiskGroupName]  
source=Volume [/DriveLetter=DriveLetter]  
[/DrivePath=DrivePath] [/Newvol=NewVolName]  
[/Plex=PlexName]...  
[writer=WriterName] [component=ComponentName]  
[server=ServerName] [instance=InstanceName]  
[backuptype=FULL|COPY] [-o]
```

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\SQL). If you wish to place the file in another directory, specify a full path before the file name, for example J:\XML\Image1.xml.

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

-g*DynamicDiskGroupName*

Indicates an SFW dynamic disk group name. If any volume is specified using the SFW volume name, this option is required.

<i>source=Volume</i>	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.
[/plex= <i>PlexName</i>]	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= <i>DriveLetter</i>]	The drive letter to be assigned to the new snapshot volume.
[/DrivePath= <i>DrivePath</i>]	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= <i>NewVolName</i>]	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
<i>writer=WriterName</i>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.
<i>component=ComponentName</i>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<i>server=ServerName</i>	SQL Server server name.
<i>instance=InstanceName</i>	SQL Server instance name.

<code>backuptype=FULL COPY</code>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default. To back up logs in SQL Server so that you can restore the database using SQL log replay, at least one Full backup must have been created earlier.
<code>-o</code>	Allows an existing XML file of the same name to be overwritten. If <code>-O</code> is not specified the <code>vxsnap create</code> command does not overwrite an existing XML file of the same name and the operation fails.

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

Examples

```
vxsnap -x billing_DB.xml create
      source=G:/Newvol=billing_data
      source=H:/Newvol=billing_log
      writer=SQLServerWriter component=billing_DB
      backuptype=full
```

This example creates a snapshot set based on the component `billing_DB` which contains volume G, the database volume, and volume H, the log volume. The snapshot volumes are named `billing_data` and `billing_log`, respectively. The XML file, `billing_DB.xml`, is used to store the VSS metadata that identifies the snapshot set. This file is used in the command to reattach the snapshot set to resynchronize it and in the command to recover the database using the snapshot set.

vxsnap reattach

Reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

Syntax

```
vxsnap -x Filename [-f] [-b] reattach writer=WriterName
```

Attributes

The following attributes apply:

<code>-x <i>Filename</i></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=<i>WriterName</i></code>	Unique ID of the VSS writer, for example, <code>SQLServerWriter</code> or the GUID for the writer.

Note: Make sure that the snapshot volumes are not in use before using this command.

Example

```
vxsnap -x billing_DB.xml reattach  
writer=SQLServerWriter
```

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.

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

vxsnap restore

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to recover a corrupted or missing SQL Server database.

Exclusive access to the SQL Server database is required for this operation.

Before using this command verify that the source volumes and the snapshot volumes are not in use.

Syntax

```
vxsnap -x Filename [-b] [-f] [-r] restore  
{RestoreType=[RECOVERY|NO_RECOVERY]}  
[noLogs|logFiles=tlog1,tlog2,...] writer=WriterName
```

Attributes

The following attributes apply:

-x <i>Filename</i>	The metadata file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file.
-b	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
-f	Forces the operation. Make sure the volume is not in use before using this option.
-r	Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume. Use this option only with Recovery noLogs. After using this option you must explicitly assign the original drive letter/mount path of the missing volume to the snapshot volume in the VEA and then bring the database online.

<code>RestoreType=</code> <code>[RECOVERY NO_RECOVERY]</code>	<p>Specifies the type of database recovery, either recovery or no recovery:</p> <p><code>RECOVERY</code> can be used with either the <code>noLogs</code> or <code>logFiles=tlog1,tlog2,...</code> attributes. <code>RECOVERY</code> leaves the database in an online state.</p> <p>To back up logs so that you can restore the database using SQL log replay, at least one Full backup must have been created earlier.</p> <p><code>NO_RECOVERY</code> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in an loading state so that you can manually replay backup logs to a specific point in time.</p>
<code>noLogs</code>	<p>Database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.</p>
<code>logFiles=tlog1,tlog2,...</code>	<p>Transaction log backup files to be applied with the <code>RECOVERY</code> option to achieve a point of failure recovery and leave the database in an online state. Each transaction log must have a unique name and be created using the “overwrite existing media” option within SQL Server.</p>
<code>writer=WriterName</code>	<p>The name for the SQL Server VSS Writer; used to located the default directory to search for the XML metadata file. Specify SQLServerWriter.</p>

Examples

Following are examples of the main types of restore operation:

- Recovering using snapshots without log replay

```
vxsnap -x TestDB.xml restore RestoreType=RECOVERY  
noLogs
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and brings the database online. The database is restored to the time the snapshot set was created or last refreshed.

You can use the `-r` option with the `RECOVERY noLogs` restore type if a production volume is missing due to hardware failure:

```
vxsnap -x TestDB.xml -r restore RestoreType=RECOVERY  
noLogs
```

This command uses the information in the `TestDB.xml` file to restore all the volumes in the snapshot set. Any missing volume is changed from a read-only volume to a read-write volume. After using the `-r` option you must explicitly assign the original drive letter/mount path of the missing production volume to the snapshot volume in the VEA. You then bring the database online.

- Recovering using snapshots and log replay

```
vxsnap -x TestDB.xml restore RestoreType=RECOVERY  
logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

This command uses the information in the `TestDB.xml` file to restore all the volumes in the snapshot set and then applies the specified transaction log backups (`c:\backup\tLog1.bak` and `c:\tLog2.bak`) and brings the database online.

- Restoring snapshots and manually applying logs

```
vxsnap -x TestDB.xml restore RestoreType=NO_RECOVERY
```

This command uses the information in the `TestDB.xml` file to restore all the volumes in the snapshot set and leaves the database in a loading state so that backup logs can be manually restored to a specific point in time.

Index

B

- backup types for snapshot sets 49
- backups, snapshot-assisted 22

C

- commands for snapshots
 - vxsnap 109
 - vxsnap create 111
 - vxsnap prepare 110
 - vxsnap reattach 114
 - vxsnap restore 115
- configuring
 - SQL Server storage with Veritas Storage Foundation for Windows 41

D

- database
 - preferred hardware layout 40
- drive letters 100
- dynamic disk groups
 - creating with the command line 42
 - creating with VEA 41
- dynamic volumes
 - creating with the command line 45
 - creating with VEA 43

F

- FastResync 27
- FlashSnap
 - integration with Microsoft Volume Shadow Copy Service 28
 - overview 27

H

- hardware recovery
 - adding disks to the dynamic disk group 102
 - replacing hardware 102

M

- mirrors for Quick Recovery snapshots (SQL Server) 52, 70, 71
- mount points
 - changing 100

P

- pre and post snapshot script file locations 30

Q

- Quick Recovery
 - file locations 30
 - Microsoft clustering considerations 29
 - VVR considerations 29
- Quick Recovery Configuration Wizard
 - overview 24, 47
 - prerequisites 52
 - running 52
- Quick Recovery for SQL Server
 - components 27
 - hardware requirements 38
 - methods 24
 - overview 21
 - prerequisites 34
 - software requirements 34
 - storage configuration recommendations 39

R

- recovery
 - SQL Server 89
- resynchronizing
 - snapshot (SQL Server) 86

S

- sample configurations
 - Quick Recovery 40
- schedules for snapshot sets
 - creating with Quick Recovery
 - Configuration Wizard 52
 - creating with VSS Snapshot
 - Scheduler Wizard 72
 - deleting 64, 65, 81
 - modifying 64, 81
 - synchronizing after adding a
 - cluster node 65
 - troubleshooting 62
- snapshot sets *See* snapshots
- snapshot volumes
 - changing drive letter 100
 - changing mount points 100
- snapshots
 - Microsoft cluster
 - considerations 29
 - VVR considerations 29
- snapshots for SQL Server
 - backup types 49
 - creating one-time snapshot
 - set 82
 - creating with Quick Recovery
 - Configuration Wizard 52
 - creating with VSS Snapshot
 - Scheduler Wizard 72
 - deleting 81
 - deleting schedules 64
 - modifying 81
 - modifying schedules 64
 - reattaching manually 86
 - reattaching split-mirror 86
 - refreshing 86
 - scheduling with Quick Recovery
 - Configuration Wizard 52
 - scheduling with VSS Snapshot
 - Scheduler Wizard 72
 - templates 49
 - troubleshooting 62

- using for recovery 89
 - viewing status 61, 78
- Solutions Configuration Center
 - context sensitivity 16
 - overview 15
 - running wizards remotely 18
 - starting 16
 - wizard descriptions 18
- SQL Server storage
 - configuration overview 41
 - creating dynamic disk groups 41
 - creating dynamic volumes 43
 - requirements and best
 - practices 39
- synchronized snapshots 75
- synchronizing schedules in a
 - cluster 65

T

- templates for snapshot sets
 - description 49
 - multiple components 50
 - schedule start dates 50

V

- Volume Shadow Copy Service 27
- VSS defined 27
- VSS Snapback Wizard
 - overview 26
 - using 86
- VSS Snapshot Scheduler Wizard
 - overview 25
 - using 72
- VSS Snapshot Wizard
 - overview 26
 - using 82
- vxsnap command reference 109
- vxsnap commands
 - create 111
 - reattach 114
 - restore 115
 - start 110