

Veritas™ Volume Replicator Web GUI Administrator's Guide

AIX

5.0 Maintenance Pack 3



Veritas™ Volume Replicator Web GUI Administrator's Guide

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: 5.0 MP3

Document version: 5.0MP3.0

Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas Storage Foundation and Veritas 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 Appendix to this Documentation or TPIP ReadMe 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
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

<http://www.symantec.com/techsupp/>

Contacting Technical Support

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

<http://www.symantec.com/business/support/index.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:

<http://www.symantec.com/techsupp/>

Customer service

Customer service information is available at the following URL:

<http://www.symantec.com/techsupp/>

Customer Service is available to assist with 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 maintenance 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

Maintenance agreement resources

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

Asia-Pacific and Japan	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

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

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

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

<http://www.symantec.com>

Select your country or language from the site index.

Contents

Technical Support	4
Chapter 1 Getting started	11
Introduction	11
Features of VRW	11
Web browser requirements	12
Veritas Volume Replicator on the Web	13
Related Veritas documents	13
Role-based administration	13
Features available with the operator role	14
Chapter 2 Setting up replication	15
About setting up replication	15
Overview of setting up replication	15
Example scenario	16
Creating a Replicated Data Set	16
Creating a Primary RVG of the RDS	17
Adding a Secondary to the RDS	20
Adding a bunker Secondary to the RDS	25
Add Bunker: Prerequisites panel	25
Add Bunker: Bunker Host, DiskGroup and Protocol panel	25
Add Bunker: Confirmation panel	26
Synchronizing the Secondary and starting replication	27
Setting up replication using automatic synchronization	27
Setting up replication using full synchronization	28
Setting up replication using block-level backup and checkpoint	29
Setting up replication using difference-based synchronization	30
Setting up replication when data volumes are initialized with zeroes	31

Chapter 3	Viewing configuration and status information	33
	About viewing configuration and status information	33
	Viewing summary information	33
	Managing Summary view	34
	Replicated Data Sets view	35
	Viewing detailed information about the RDS	37
	RDS detail view - Summary tab	37
	RDS detail view - Primary tab	40
	RDS detail view - Secondaries tab	43
	RDS detail view - Bunker tab	47
	RDS detail view - Clusters tab	49
	RDS detail view - Alerts tab	50
	Information displayed in RDS views	50
	Conventions for the RDS icons	50
	Conventions for the RVG icons	50
	Primary RVG states	51
	Secondary RVG data status	52
	Replication Status	52
Chapter 4	Administering Veritas Volume Replicator	55
	About administering Veritas Volume Replicator	55
	Administering data volumes	55
	Associating volumes to a Replicated Data Set	56
	Resizing a data volume in a Replicated Data Set	58
	Dissociating a data volume from its Replicated Data Set	59
	Administering replication	61
	Changing the replication settings	61
	Starting replication to a Secondary	63
	Pausing replication to a Secondary	65
	Resuming replication to a Secondary	66
	Stopping replication to a Secondary	67
	Changing the IP addresses used for replication	68
	Administering the Replicated Data Set	69
	Removing a Secondary from a Replicated Data Set	69
	Removing a Bunker from a Replicated Data Set	70
	Removing a Primary RVG	71
	Administering checkpoints	72
	Creating checkpoints	73
	Ending checkpoints	73
	Viewing checkpoints	74
	Deleting checkpoints	74
	Administering the SRL	75

	Protecting from SRL overflow	76
	Incrementally synchronizing the Secondary after SRL overflow	76
	Resizing the SRL	78
Chapter 5	Transferring the Primary role	81
	About transferring the Primary role	81
	Migrating the Primary role	81
	Migrate Role: Prerequisites panel	83
	Migrate Role: Summary panel	83
	Taking over from an original Primary	83
	Takeover dialog window	84
	Failing back to the original Primary	85
	Failing back using fast-failback synchronization	86
	Failing back using difference-based synchronization	88
	Using a bunker for disaster recovery	90
	Updating the Secondary from the bunker	91
	Activating a bunker in an RDS	92
	Deactivating a bunker	92
Appendix A	Troubleshooting VRW	95
	About troubleshooting VRW	95
	Application page is not displayed in the browser	95
	Message	95
	Problem	95
	Solution	96
	Missing buttons on certain wizards	96
	Problem	97
	Solution	97
	Log in redirects the user back to the login page	97
	Problem	97
	Solution	97
Index		99

Getting started

This chapter includes the following topics:

- [Introduction](#)
- [Features of VRW](#)
- [Web browser requirements](#)
- [Veritas Volume Replicator on the Web](#)
- [Related Veritas documents](#)
- [Role-based administration](#)

Introduction

Veritas Volume Replicator Web Console (VRW) is a Web-based graphical user interface that enables you to create, manage, and administer Veritas Volume Replicator (VVR) configurations using a Web browser. The Veritas Volume Replicator *Web GUI Administrator's Guide* (this guide) provides information on configuring, and administering VVR using VRW. It describes the features and options available in VRW and enables you to start using VRW. This guide also provides information on how to perform advanced VRW configuration tasks. The troubleshooting section enables you to recover from problems that you may encounter while using VRW.

This guide is supplemental to the Veritas Volume Replicator *Administrator's Guide*, which you should read before using VRW.

Features of VRW

VRW includes the following features:

- **VRW Views**
VRW provides views that display information about the selected VVR objects and the VVR environments. For example, the Summary view displays a list of Replicated Data Sets (RDSs) present on the selected hosts and the replication status for each RDS, whereas the RDS view displays details of the RDS and its status.
- **Web-Based Wizards**
VRW provides Web-based wizards to perform VVR administrative tasks, including creating the Primary, adding a Secondary, starting replication, and performing migration to change the Primary role. Each wizard provides instructions to enable you to perform its tasks from start to finish.
- **Centralized Management of VVR**
VVR can be centrally managed through the Storage Foundation Management Server. For more information, see the documentation for Storage Foundation Management Server.
- **Managing User Roles**
VRW supports the roles provided by SF Management Server. A user with Administrator access rights can assign to users and user groups one of the predefined roles: Guest, Operator, and Administrator.
See [“Role-based administration”](#) on page 13.

Web browser requirements

Depending on the platform, one of the following Web browsers is required to run VRW:

Table 1-1 Web browsers and their supported platforms

Operating System	Firefox 1.0.7 (or later)	Internet Explorer 5.5 (or later)	Mozilla 1.6 (or later)	Netscape Navigator 7.0 (or later)
AIX			✓	
HP-UX			✓	
Red Hat	✓		✓	✓
Solaris			✓	
SUSE	✓		✓	✓
Windows	✓	✓		✓

Your browser must support either JScript 5.5 (provided with Microsoft Internet Explorer 6 or later) or JavaScript 1.2 (provided with Netscape Navigator 7.x). Additionally, for Internet Explorer, Secure Sockets Layer (SSL) 3.0 ActiveX scripting must be enabled.

If you are using pop-up blockers (including Yahoo Toolbar or Google Toolbar), either disable them or configure them to accept pop-ups from the Web Server to which you will connect. Note that pop-up blocking is enabled by default in Netscape 7.2.

Note: For Internet Explorer 6.0 on Windows 2003 (Server and Advanced Server), set the default Intranet zone security level to MEDIUM or lower.

Veritas Volume Replicator on the Web

For comprehensive, up-to-date information about the Veritas Volume Replicator product offerings, visit the Veritas Web site:

<http://www.veritas.com/Products/www?c=product&refId=3>

Related Veritas documents

For more information on any of the topics presented in this guide, refer to the Veritas Volume Replicator (VVR) or the Veritas Volume Manager (VxVM) documentation sets.

See the *Veritas Storage Foundation Release Notes* for more information on these documentation sets.

Role-based administration

Authorization is access control. It determines which activities a user or group can perform on a particular managed host or disk group.

When a user is created, the user is assigned to a specific security role. The user may be assigned different roles on different hosts.

A user can be assigned one of following roles:

- **Administrator:** A user assigned to this role has complete access to the system. The root user on a host is automatically granted Administrator privileges for that host. These privileges cannot be changed.

- **Operator:** A user assigned to this role can perform some management operations, such as performing backups, administrating the database, and making some limited configuration changes.
A user with the operator role can also view the objects in the network and print reports.
See [“Features available with the operator role”](#) on page 14.
- **Guest:** A user assigned to this role cannot make any changes to the configuration. A Guest can view the objects in the network and print reports.
All authenticated users by default belong to this group.

Features available with the operator role

The following features are accessible to Operators:

- Start Replication
- Stop Replication
- Pause Replication
- Resume Replication
- Start Checkpoint
- Delete Checkpoint
- End Checkpoint
- Replay DCM
- Set Replication Attributes

Setting up replication

This chapter includes the following topics:

- [About setting up replication](#)
- [Overview of setting up replication](#)
- [Creating a Replicated Data Set](#)
- [Adding a bunker Secondary to the RDS](#)
- [Synchronizing the Secondary and starting replication](#)
- [Setting up replication when data volumes are initialized with zeroes](#)

About setting up replication

You can configure Veritas Volume Replicator (VVR) using Volume Replicator Web Console (VRW). VVR enables you to set up replication either when the data volumes are zero initialized or contain valid data, and when the application is running or stopped.

Before setting up a Replicated Data Set, lay out your VVR configuration.

Follow the best practices or recommendations listed to ensure successful configuration of VVR.

See the *Veritas Volume Replicator Administrator's Guide* for details.

Overview of setting up replication

To configure and set up replication, perform the following tasks in the order presented below:

- [Creating a Replicated Data Set](#)

■ [Synchronizing the Secondary and starting replication](#)

Note: The procedure to set up replication is the same either when the application is running or stopped, unless noted otherwise.

Example scenario

This chapter explains how to use Volume Replicator Web Console (VRW) to set up a simple VVR configuration under different situations with the help of examples. The examples explain the procedure to set up a VVR configuration containing one Primary and one Secondary. However, VRW enables you to create and set up VVR configurations containing multiple Secondary hosts. The example configuration looks like this:

Primary Host name: seattle

Disk_group	hrdg
Primary RVG	hr_rvg
Primary RLINK to Secondary london	rlk_london_hr_rvg
Primary data volume #1	hr_dv01
Primary data volume #2	hr_dv02
Primary SRL volume	hr_srl

Secondary Host name: london

Disk group	hrdg
Secondary RVG	hr_rvg
Secondary RLINK to Primary seattle	rlk_seattle_hr_rvg
Secondary data volume #1	hr_dv01
Secondary data volume #2	hr_dv02
Secondary SRL volume	hr_srl

Creating a Replicated Data Set

To create a Replicated Data Set (RDS) using VRW, perform the following tasks in the order listed below:

- [Creating a Primary RVG of the RDS](#)
- [Adding a Secondary to the RDS](#)
- [Adding a bunker Secondary to the RDS](#)

Creating a Primary RVG of the RDS

The first step in creating an RDS is creating its Primary RVG. VRW enables you to create a Primary RVG of an RDS using the Create Primary wizard. The Create Primary wizard enables you to associate existing data volumes and the Storage Replicator Log (SRL) with the Primary RVG. The Create Primary wizard performs the following operations:

- Creates the Primary RVG on the host to which you are connected.
In a shared disk group environment, you should be connected to the Master in the cluster.
- Associates the specified data volumes and SRL with the RVG.
- Associates Data Change Maps (DCMs) to the data volumes in the RVG.
- Enables the Primary RVG by starting it.
- Associates the specified volume sets (if any) to the RVG.

To add more volumes after creating the RVG, use the Associate Volume wizard.

Observe the following prerequisites:

Prerequisites	The data volumes and the SRL must exist in a disk group on the Primary host.
	All the data volumes used by the application MUST be in the same RVG.
	The data volumes and the SRL must be started.
	All the data volumes to be replicated and the SRL must be the same type. They must be all VxVM ISP volumes or VxVM non-ISP volumes.
	The SRL cannot be a volume set or a component volume of a volume set.

For instructions on creating the data volumes and SRL, see the example in the *Veritas Volume Replicator Administrator's Guide*.

To create a Primary RVG of an RDS

- 1 From the Configuration task pane of any view, select **Create Primary**.
In a shared disk group environment, you must be connected to the Master node of the cluster in order to create a Primary.
- 2 Select the volume type. If the data volumes to be replicated were created using Intelligent Storage Provisioning, select the checkbox **The volumes are VxVM ISP volumes**. Click **Next**.
- 3 Complete the RVG Name and Disk Group panel, and then click **Next**.
See [“Create Primary Wizard: Specify RVG Name and Disk Group panel”](#) on page 19.
- 4 Complete the Data Volume Selection page, and then click **Next**.
See [“Create Primary Wizard: Data Volume Selection panel”](#) on page 19.
- 5 Complete the SRL Volume Selection page.
See [“Create Primary Wizard: SRL Volume Selection panel”](#) on page 20.
- 6 Click **Finish** to proceed with the creation of the Primary RVG.
 - To make any changes to the RVG name or to select a different disk group, click **Back**.
 - To exit the Create Primary wizard and cancel the Create Primary operation, click **Cancel**.

After the Primary RVG is created, the following message displays in a window:

```
The Primary RVG is created on seattle.
```

- 7 Click **OK** in the Result page.

Create Primary Wizard: Prerequisites panel

Use the Create Primary wizard to create a Primary RVG, which is the first step in creating an RDS. The Create Primary wizard enables you to associate existing data volumes and the Storage Replicator Log (SRL) with the Primary RVG. Review the Prerequisites panel and indicate the type of volumes.

See [“Creating a Primary RVG of the RDS”](#) on page 17.

The volumes are VxVM ISP volumes	An RVG cannot contain both VxVM ISP volumes and non-VxVM ISP volumes. Select the checkbox if the RVG will include VxVM ISP volumes. If this option is selected, the subsequent Create Primary wizard panels only display VxVM ISP volumes.
---	---

Create Primary Wizard: Specify RVG Name and Disk Group panel

Use this panel to specify the name for the Primary RVG that you are creating, and the name of the disk group.

See [“Creating a Primary RVG of the RDS”](#) on page 17.

Host Name	Displays the name of the host. This field is not editable.
RVG (RDS) Name	Enter a name for the Primary RVG. The RDS name is the same as the RVG name. A valid RDS name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period. For example, hr_rvg.
Disk Group	From the drop-down list, select the name of the disk group, that is, the disk group that contains the SRL and the data volumes to be associated with the Primary RVG. For example, hrdg. The disk group is either shared or private. Note: If the checkbox The volumes are VxVM ISP volumes was selected on the previous page, the list only displays disk groups that contain VxVM ISP data volumes.

Create Primary Wizard: Data Volume Selection panel

Use the Data Volume Selection panel to select the data volumes to be associated to the Primary RVG.

Host Name	Displays the name of the host. This field is not editable.
RVG (RDS) Name	Displays the name of the RDS. This field is not editable.

Data Volumes

Select one or more data volumes or volume sets to be replicated. The selected data volumes are associated to the Primary RVG. All of the data volumes used by the application **MUST** be included in the same RVG.

To select multiple data volumes from the list, hold down the CTRL key and click the names of the required data volumes.

To select a range of data volumes, select the first data volume in the range, hold down the SHIFT key, and then click the last name in the range. For example, hr_dv01 and hr_dv02.

For volume sets, the list displays the volume set name. The list does not display the names of the component volumes, because a component volume cannot be associated to an RVG. Select the volume set name to associate all of its component volumes to the RVG.

Note: If the checkbox The volumes are VxVM ISP volumes was selected on the first page, the list only displays VxVM ISP data volumes.

Create Primary Wizard: SRL Volume Selection panel

Use the SRL Volume Selection panel to select the volume to be associated as the Storage Replication Log (SRL) for the Primary RVG.

Replication Log

Select the volume that is to be configured as the SRL. For example, hr_srl. The list does not display volume sets or component volumes of a volume set, because they cannot be used for the SRL volume. The SRL size must be at least 110MB.

Note: If the checkbox The volumes are VxVM ISP volumes was selected on the first panel, the list only displays VxVM ISP volumes.

Create Primary Wizard: Confirmation panel

Use the Confirmation panel to review the configuration for the Primary RVG to be created. Click Finish to create the Primary.

Adding a Secondary to the RDS

To add a Secondary RVG to the RDS, use the Add Secondary wizard from any host that is already a part of the RDS.

Note: If the RDS consists of only the Primary host, use the Add Secondary wizard that is available from the Primary. Do not use the Add Secondary wizard from the Secondary host that you want to add to the RDS.

The host to which you are connected is referred to as the local host, and the name of the RVG on the local host is referred to as the local RVG. The local RVG represents the RDS to which it belongs, and the RDS is referred by its local RVG name.

When adding the Secondary, the Add Secondary wizard performs the following operations:

- Creates and adds a Secondary RVG with the same name as the Primary RVG, to the specified RDS on the Secondary host. By default, the Secondary RVG is added to the disk group with the same name as the Primary disk group. You can specify a different disk group by using the options provided in the Add Secondary wizard.
- Adds DCMs to the Primary and Secondary data volumes if they do not have DCMs. If any of the data volumes or the SRL on the Secondary has a DRL, the DRL is removed before the data volume or SRL is associated with the RVG.
- Associates to the Secondary RVG existing data volumes of the same names and sizes as the Primary data volumes; it also associates an existing volume with the same name as the Primary SRL, as the Secondary SRL.
- Creates and associates Primary and Secondary RLINKs with default RLINK names `rlk_remotehost_rvgname`. The wizard enables you to specify different names for the Primary and Secondary RLINKs.

Before creating and adding the Secondary RVG to the specified RDS, do the following:

- The Secondary data volumes and the SRL must be in the same disk group on the Secondary. Use a disk group with the same name as the Primary disk group, if possible.
- Create data volumes with the same names and sizes as the Primary data volumes.
- Create an SRL with the same name and size as the Primary SRL.
- Make sure the `/etc/vx/vras/.rdg` file contains the Primary disk group ID. For details, see the *Veritas Volume Replicator Administrator's Guide*. To display the Primary disk group ID, log on to the Primary host and issue the following command:

```
# vxprint -l diskgroup
```

To add a Secondary to an RDS

- 1 Navigate to the detail view of the RDS to which you want to add the Secondary RVG.
- 2 From the Configuration task pane, select **Add Secondary**.
- 3 Review the prerequisites. Click **Next**.

See [“Add Secondary Wizard: Prerequisites panel”](#) on page 23.

- 4 Complete the Primary and Secondary Host Names page, and then click **Next**.
See [“Add Secondary Wizard: Primary and Secondary Host Names panel”](#) on page 23.

- 5 In the RLINKs and Secondary Disk Group Names panel, the default values for the Primary RLINK, Secondary RLINK, and the Secondary Disk Group display. To accept the default values, click **Next**.

If required, change the values for the Primary RLINK, Secondary RLINK, and the Secondary Disk Group by completing the RLINKs and Secondary Disk Group Names page, and then click **Next**.

See [“Add Secondary Wizard: RLINKs and Secondary Disk Group Names panel”](#) on page 24.

- 6 Click Finish to add the Secondary RVG.

■ To make any changes, click **Back**.

The Add Secondary operation succeeds if the SRL and data volumes with the required properties are available in the specified Secondary disk group.

- 7 Click **OK** on the Result page. The Secondary RVG view for the newly added Secondary RVG is displayed.

Warning: If the data volumes on the Primary contain valid data, you must synchronize the Secondary data volumes with the Primary data volumes before replicating to the Secondary.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

If the data volumes are initialized with zeroes, set up replication as described.

See [“Setting up replication when data volumes are initialized with zeroes”](#) on page 31.

If required, add a bunker Secondary.

See [“Adding a bunker Secondary to the RDS”](#) on page 25.

Add Secondary Wizard: Prerequisites panel

Use the Add Secondary wizard to create a Secondary RVG on the Secondary host and add it to the specified RDS. Use the Add Secondary wizard from any host that is already a part of the RDS.

Review the Prerequisites panel and proceed to the next step only after the prerequisites are met.

See [“Adding a Secondary to the RDS”](#) on page 20.

Add Secondary Wizard: Primary and Secondary Host Names panel

Use the Primary and Secondary Host Names panel to specify the host names or IP addresses for the Primary and Secondary hosts in the RDS.

RDS	Displays the name of the RDS. This field cannot be edited.
Primary Host	Enter a host name or specify the IP address that is to be used for replication. The Primary host name must be resolvable and reachable from the Secondary host. For example, seattle. If you entered the IP address, it must be reachable from the Secondary host. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
Secondary Host	Enter a host name or specify the IP address that can be used for replication. The Secondary host name must be resolvable and reachable from the Primary host. For example, london. If you entered the IP address, it must be reachable from the Primary host. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.

See [“Adding a Secondary to the RDS”](#) on page 20.

Add Secondary Wizard: RLINKs and Secondary Disk Group Names panel

Use the RLINKs and Secondary Disk Group Names panel to change the values for the Primary RLINK, Secondary RLINK, and the Secondary Disk Group names. If no changes are required, click Next to accept the default values.

RDS	Displays the name of the RDS. This field cannot be edited.
Primary RLINK	The default RLINK name is in the form <code>rlk_remotehost_rvgname</code> . Enter another Primary RLINK name if the default name is not appropriate. A valid RLINK name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
Secondary RLINK	The default RLINK name is in the form <code>rlk_remotehost_rvgname</code>. Enter another Secondary RLINK name if the default name is not appropriate. A valid RLINK name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
Secondary Disk Group	By default, the Secondary Disk Group box displays the name of the Primary disk group. This will be used as the name of the disk group on the Secondary that is to contain the Secondary RVG. If the default name is not appropriate, enter another Secondary disk group name. This disk group must contain the SRL and data volumes that are to be associated with the Secondary RVG that is being created. A valid disk group name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.

See [“Adding a Secondary to the RDS”](#) on page 20.

Add Secondary Wizard: Confirmation panel

Use the Confirmation panel to review the configuration for the Secondary RVG to be added. Click Finish to add the Secondary RVG.

See [“Adding a Secondary to the RDS”](#) on page 20.

Adding a bunker Secondary to the RDS

Use the Add Bunker wizard to add a bunker Secondary to the RDS. After adding the bunker Secondary, you can proceed with synchronizing the Secondary and starting replication.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

To add a bunker Secondary to the RDS

- 1 Select the name of the Primary for which you want to create the bunker Secondary. For example, hr_rvg.
- 2 Choose **Replication > Add Bunker**. The Add Bunker wizard is displayed.
- 3 Review the information on the Prerequisites panel, then click **Next**.
- 4 Complete the Add Bunker wizard page.
See [“Add Bunker: Bunker Host, DiskGroup and Protocol panel”](#) on page 25.
- 5 In the next page, the message indicates the result of the add bunker operation. If the operation of adding the bunker Secondary was successful the page displays a success message.

If the operation fails, the page displays the failure message and lists the possible reasons of failure in the Result Details area. Read the details carefully and take appropriate corrective action. For more information about the error, note down the Unique Message Identifier (UMI) that is displayed with the error message, and then refer to the Veritas Technical Services knowledge base.
- 6 Click **Finish**. The Add Bunker wizard sets the same default values that are set by the Add Secondary wizard.

Add Bunker: Prerequisites panel

Use the Add Bunker dialog window to add a bunker Secondary to the RDS.

Review the information on the Prerequisites panel, and proceed to the next step only after the prerequisites are met.

See [“Adding a bunker Secondary to the RDS”](#) on page 25.

Add Bunker: Bunker Host, DiskGroup and Protocol panel

Use the Bunker Host, DiskGroup and Protocol page to specify the host names, disk group, and protocol to use when adding a bunker Secondary to the RDS.

RDS	Displays the name of the RDS for which you are adding the bunker. This field cannot be edited. Example: <code>hr_rvg</code>
Primary Host	Enter a host name or specify the IP address that can be used for replication. The Primary host name must be resolvable and reachable from the Bunker host. For example, <code>seattle</code>. If you entered the IP address, it must be reachable from the bunker host. The Replication IP column in the Secondary RVG view displays the value of the Primary host.
Bunker Host	Enter a host name or specify the IP address that can be used for replication. The Bunker host name must be resolvable and reachable from the Primary host. For example, <code>portland</code>. If you entered the IP address, it must be reachable from the Primary host. The Replication IP column in the Primary RVG view displays the value of the Bunker host.
Bunker DiskGroup	If the name of the bunker disk group is other than the default, enter the correct name. If the protocol is UDP or TCP, the disk group usually is the same as the disk group of the Primary RVG. For STORAGE protocol, the bunker disk group has a different name than the disk group for the Primary RVG, because both disk groups will be imported on the Primary host. The bunker disk group must contain the SRL that is to be associated with the Secondary bunker RVG that is being added.
Protocol	Indicates the network connectivity between the bunker Secondary and the Primary. UDP is selected by default. If the bunker is replicating over IP, set the protocol to UDP or TCP. If the storage is directly accessible by the Primary, for example, DAS or NAS, set the protocol to STORAGE. Select UDP, TCP, or STORAGE.

Add Bunker: Confirmation panel

Confirm the attributes for the operation. Click Finish to add the bunker Secondary to the RDS.

See [“Adding a bunker Secondary to the RDS”](#) on page 25.

Synchronizing the Secondary and starting replication

Before starting replication, the Secondary data volumes must be synchronized with the Primary data volumes. You can use VRW to synchronize the Secondary with the Primary with one of the following methods:

- Automatic synchronization

See [“Setting up replication using automatic synchronization”](#) on page 27.

Note: For component volumes of a volume set, use Start Replication with the Using Autosync option to synchronize the RVG.

- Full synchronization

See [“Setting up replication using full synchronization”](#) on page 28.

- Block-level tape backup

See [“Setting up replication using block-level backup and checkpoint”](#) on page 29.

- Difference-based synchronization

See [“Setting up replication using difference-based synchronization”](#) on page 30.

For more information on the methods to synchronize the Secondary and choosing the appropriate method to synchronize the Secondary, see the *Veritas Volume Replicator Administrator's Guide*.

Setting up replication using automatic synchronization

You can synchronize the Secondary using automatic synchronization when the data volumes contain data and when the application is active or inactive.

The instructions in this section assume that the RDS has been created.

See [“Creating a Replicated Data Set”](#) on page 16.

Using the automatic synchronization option enables you to automatically synchronize the Secondary data volumes with the Primary data volumes in an RDS and start replication; automatic synchronization makes the Secondary data volumes up to date with the Primary data volumes.

Replication to another Secondary can be started only after this automatic synchronization completes. Automatic synchronization copies the data on the Primary to the Secondary over the network using the Data Change Map (DCM). To use automatic synchronization, each data volume in the RVG must have an associated Data Change Map (DCM).

To set up replication using automatic synchronization

- 1 Navigate to Replication Status or Replication Settings sections of one of the following views:
 - the main Summary view
 - the Secondaries tab of the RDS view
 - the Summary tab of the RDS view
- 2 Select the Secondary to be synchronized, select **Start Replication** from the drop-down list, and click **Go**.
- 3 In the Start Replication dialog box, select **Using Automatic synchronization** to synchronize the Secondary and start replication using automatic synchronization.
- 4 Click **Finish**. A message displays the status of your request to start replication.
- 5 Click **OK**. The RDS view for `hr_rvg on seattle` displays the status of replication at this time. At any time, refresh the view to display the current status of replication.

Setting up replication using full synchronization

You can synchronize the Secondary using full synchronization with checkpoint when the application is active or inactive.

The instructions in this section assume that the RDS has been created.

See [“Creating a Replicated Data Set”](#) on page 16.

To set up replication using full synchronization

- 1 Synchronize the Secondary using the following command:

```
# vradmin -g diskgroup -c checkpoint_name -full syncrvg \  
local_rvgname sec_hostname...
```
- 2 After the synchronization completes, use the Start Replication wizard to start replication to the Secondary with the checkpoint.
- 3 In the Start Replication dialog box, select **Using checkpoint**.
From the drop-down menu, select the checkpoint name which you entered in step 1.

- 4 Click **Finish**. A message displays the status of your request to start replication.
- 5 Click **Ok**. The Secondary RVG view for hr_rvg on seattle displays the status of replication at this time. At any time, refresh the view to get the current status of replication.

If you specified multiple Secondary hosts in step 1, repeat step 2 through step 5 to start replication for each Secondary RVG.

Setting up replication using block-level backup and checkpoint

You can synchronize the Secondary using block-level backup and checkpoint when the application is active or inactive. The instructions in this section assume that the RDS has been created.

See [“Creating a Replicated Data Set”](#) on page 16.

To set up replication using block-level backup and checkpoint

- 1 In the Primary tab for the RDS view, in the Primary RVG section, select **Start Checkpoint** to start a checkpoint on the Primary.
- 2 In the Start Checkpoint dialog box, enter a name for the checkpoint and click **Finish**. Note down the checkpoint name you entered, for example, `checkpt_ddmmyyyy`.
- 3 Perform a block-level backup of the data volumes in the Primary RVG.
- 4 Select **End Checkpoint** to end the Primary checkpoint when the backup is complete.
- 5 To confirm that you want to end the checkpoint, click **Finish**.
- 6 Restore the backup to the Secondary data volumes.
- 7 When step 6 is completed, check whether the checkpoint you created is still valid in the Primary detail view. Refresh the view: if the checkpoint status shows Completed, the checkpoint is valid. If it shows another status, it may have overflowed or otherwise not be valid. Repeat step 1 to step 6.
- 8 Use the Start Replication wizard to start replication.
- 9 In the Start Replication dialog box, select **Using Checkpoint**. From the drop-down menu, select the checkpoint created above. This uses the selected checkpoint to synchronize the Secondary and start replication.
- 10 Click **Finish**. A message displays the result of your request to start replication.

- 11 Click **Ok**. The Secondaries tab of the RDS view for `hr_rvg on seattle` displays the current replication status, for the replication that has just been started. At any time, refresh the view to display the current status of replication.
- 12 In the Secondaries tab of the RDS view, check whether the `consistent` flag is set in the Data Status column for the Primary RLINK. The RLINK becomes `consistent` only after the data contained in the checkpoint is sent to the Secondary.

If the Secondary is consistent, the synchronization was successful. If the checkpoint overflows before the Secondary becomes consistent, the synchronization failed. Increase the size of the SRL, and then repeat step 1 to step 11.

For instructions on resizing the SRL, see the *Veritas Volume Replicator Administrator's Guide*.

It is likely that there might be writes beyond the checkpoint that are yet to be sent to the Secondary after the `consistent` flag is set on the RLINK. In the RDS view, check **Status** to find out whether the RLINK is up to date.

The same backup and the corresponding checkpoint can be used to set up additional Secondary hosts while the checkpoint is still valid. If a checkpoint has overflowed, its corresponding backup cannot be used to resynchronize the Secondary using that checkpoint. Eventually, any checkpoint becomes STALE and unusable. VRW does not display a warning to indicate if a checkpoint becomes unusable. However, the status of the checkpoint is displayed in the Primary tab of the RDS view.

See “[Viewing detailed information about the RDS](#)” on page 37.

Setting up replication using difference-based synchronization

You can synchronize the Secondary using difference-based synchronization with checkpoint when the application is active or inactive.

The instructions in this section assume that the RDS has been created.

See “[Creating a Replicated Data Set](#)” on page 16.

To set up replication using difference-based synchronization

- 1 Synchronize the Secondary using the following command:

```
# vradmin -g diskgroup -c checkpoint_name syncrvg \
local_rvgname sec_hostname...
```

- 2 After the synchronization completes, use the Start Replication wizard to start replication to the Secondary with the checkpoint.

- 3 In the Start Replication dialog box, select **Using checkpoint** and select the checkpoint name from the drop-down menu. This uses the selected checkpoint to synchronize the Secondary and start replication.
- 4 Click **Finish**. A message displays the result of your request to start replication.
- 5 Click **Close**. The Secondary RVG view for `hr_rvg` on `seattle` displays the status of replication at this time. At any time, refresh the view to get the current status of replication.

If you specified multiple Secondary hosts in step 1, repeat step 2 through step 5 to start replication for each Secondary RVG.

Setting up replication when data volumes are initialized with zeroes

Because the Primary data volumes are initialized with zeroes, the data on the Secondary node need not be synchronized with the Primary. However, we recommend that you zero initialize the Secondary data volumes. The instructions in this section assume that the RDS has been created.

See [“Creating a Replicated Data Set”](#) on page 16.

Use the Resynchronization is not required option to start replication only when the Primary data volumes are zero initialized or when the contents of the volumes on the Primary and Secondary are identical and the application is inactive.

Note: Do not start the application or mount the file system before completing step 5.

To set up replication when data volumes are initialized with zeroes

- 1 Navigate to the Managing Summary view, the Secondaries tab of the RDS view or the Summary tab of the RDS view.
- 2 In the Replication Settings section, select the Secondary to be synchronized, and select **Start Replication** from the drop-down list.
- 3 In the Start Replication dialog box, select **Resynchronization is not required**.
- 4 Click **Finish**. A message displays the result of your request to start replication.
- 5 Click **OK**. The RDS view for `hr_rvg` on `seattle` is automatically refreshed to display the replication status at a particular point in time.

After replication is started, start the application on the Primary.

Viewing configuration and status information

This chapter includes the following topics:

- [About viewing configuration and status information](#)
- [Viewing summary information](#)
- [Viewing detailed information about the RDS](#)
- [Information displayed in RDS views](#)

About viewing configuration and status information

Volume Replicator Web Console (VRW) enables you to get information about the VVR environment using the views. The Replicated Data Sets view displays consolidated information about the Replicated Data Sets (RDSs) in a Veritas Volume Replicator (VVR) environment; the other views provide detailed information about the selected object.

Viewing summary information

The Volume Replicator Web Console provides summary information about VVR objects in the following views:


- [Managing Summary view](#)
- [Replicated Data Sets view](#)
- [RDS detail view - Summary tab](#)

Managing Summary view

The Volume Replicator section of the Managing Summary view displays information about the RDSs present on the connected host. Use the Managing Summary view to view summary information about all of the RDSs present on the connected host.

Figure 3-1 shows the Managing Summary view.

Figure 3-1 Managing Summary view



RDS Replication Status

This section shows the count of RDSs based on the current state of the replication (replication status) between the Primary RVG and the Secondary RVG for the RDS.

See “Replication Status” on page 52.

The RDS replication status section shows the following information:

- **Replicating:** The count of Replicated Data Sets where the replication status for the Secondary is replicating, resync in progress or logging to DCM. If the RDS contains multiple Secondaries, the RDS is counted as replicating if at least one Secondary is replicating.
- **Not Replicating:** The count of Replicated Data Sets where the Secondary is not connected or replication is not occurring. For example, the replication status

is not replicating or paused by user. If the RDS contains multiple Secondaries, the RDS is counted as not replicating only if none of its Secondaries is replicating. RDSs that have no Secondary, or that are missing a Primary, are counted as not replicating.

RDS Behind By Status

This section shows the count of Replicated Data Sets based on the amount of time the Secondary is behind the Primary. If the RDS contains multiple Secondaries, the Behind-By status reflects the status of the most up-to-date Secondary for that RDS.

- (0-15] minutes: The Secondary is up-to-date, or the last write that the Secondary received from the Primary was less than fifteen minutes ago.
- (15 - 120] minutes: The last write that the Secondary received from the Primary more than fifteen minutes (inclusive) ago but less than two hours ago.
- Greater than 2 hours: The last write that the Secondary received from the Primary was more than two hours ago.
- Not Available: indicates the Behind-By status is not available. The status may not be available if the RDS has no Secondary, or if the Secondary is paused or stopped.

Alerts section

Displays configuration problems in an RDS, if any. The alerts section displays the RDS name together with a brief description about the configuration error.

Replicated Data Sets view

Select Replicated Data Sets in the Managing Summary pane to display the Replicated Data Sets view. The Replicated Data Sets view includes the following tabs:


- [Replicated Data Sets view - Summary tab](#)
- [Replicated Data Sets view - Alerts tab](#)

Replicated Data Sets view - Summary tab

The Summary tab of the RDSs view provides summary information about each RDS on the connected host, including the replication status, the logging status, and the names of the Primary and Secondary hosts.

[Figure 3-2](#) shows the Summary tab of the RDSs view

Figure 3-2 Summary tab of the RDSs view



Replicated Data Sets

The Replicated Data Sets section displays the following information:

Name	The RDS icon and the name of the RDS. The RDS name is a navigation link to the RDS view. See Table 3-1 on page 50.
Primary	The name of the Primary host. If the Primary has not been configured correctly, this field is blank and the word <code>unknown</code> is displayed in the replication status column.
Secondary/Bunker	The name of the Secondary host or bunker Secondary host.
Replication Status	The replication status to the Secondary in text format. See Table 3-5 on page 53. The replication status <code>n/a</code> means that the RDS is configured incorrectly and the Primary is not known. A text message also displays if a Secondary does not exist.
Behind By (Size)	Indicates the size of the writes remaining to bring the Secondary up-to-date with the Primary.

Behind By (Time)	Indicates the elapsed time since the last write that the Secondary received from the Primary.
Current Mode	Shows the replication mode, either <code>synchronous</code> or <code>asynchronous</code> .
Current Logging	Indicates whether the RDS is currently using the SRL or the DCM.
% Log Usage	Shows either the percentage of the SRL that is full or the percentage of the DCM that is dirty.
Alert	Displays configuration problems in the RDS, if any.
Data Status	Shows the Data Status of the Secondary RVG. See Table 3-4 on page 52.
Secondary Disk Group	The name of the Secondary disk group. For a shared disk group, (Shared) displays after the name.

Replicated Data Sets view - Alerts tab

The Alerts tab of the RDSs view displays information about any configuration errors for the RDSs that reside on the connected host. The Alerts tab displays the RDS name together with a brief description about the configuration error.

Viewing detailed information about the RDS

From the Replicated Data Sets Summary view, you can navigate to the following detailed views:

- [RDS detail view - Summary tab](#)
- [RDS detail view - Primary tab](#)
- [RDS detail view - Bunker tab](#)
- [RDS detail view - Clusters tab](#)
- [RDS detail view - Alerts tab](#)


RDS detail view - Summary tab

The Summary tab of the RDS view displays information about the Primary RVG, replication status of the Secondary RVG, and the data volumes in the RDS. Use this view to administer replication and data volumes. From the Managing Summary view, click on the RDS name to display the Summary tab of the RDS view. For

example, click the RDS name `hr_rvg`, to display the Summary tab of the RDS view for `hr_rvg`.

Figure 3-3 shows the Summary tab of the RDS view.

Figure 3-3 Summary tab of the RDS view



Primary Replicated Volume Group

This section shows the details for the RVG on the Primary.

Host	Displays Primary RVG icon and the name of the Primary host. For a shared disk group, the name of the logowner displays as the Primary host. If the Primary has not been configured correctly, this field is blank and the word <code>unknown</code> is displayed in the replication status column. The Primary RVG icon indicates the I/O state. See Table 3-2 on page 51.
Name	The name of the RVG (RDS name). The RVG state displays in parenthesis after the name. See “Primary RVG states” on page 51.

Disk Group	The name of the Primary disk group. For a shared disk group, (Shared) displays after the name.
Replicated Data Size	The total size of the replicated data volumes in the RDS. If a volume set is associated to the RDS, the size includes the total size of each of the component volumes of the volume set in addition to the total size of any independent volumes.
Storage Replication Log	The name of the Primary SRL.
Storage Replication Log Size	The size of the SRL.
Bunker	The name of the bunker host. Displays N/A if the RDS does not contain a bunker.
Cluster	The name of the cluster if the Primary RVG uses a shared disk group. Select the cluster name to view details about the cluster.

Replication status

Displays the status-related information of each Secondary in the selected RDS. This section displays the following information:

Secondary	Displays the Secondary RVG icon and the name of the Secondary host. For a shared disk group, the name of the logowner displays as the Secondary host. The Secondary RVG icon indicates the configuration state. See Table 3-2 on page 51.
Data status	Shows the data status. See Table 3-4 on page 52.
Replication status	Shows the replication status. See “Replication Status” on page 52.
Current mode	Shows the replication mode, either <code>synchronous</code> or <code>asynchronous</code> .
% Log Usage	Shows either the percentage of the SRL that is full or the percentage of the DCM that is dirty.

Behind By (Size)	Indicates the size of the writes remaining to bring the Secondary up-to-date with the Primary.
Behind By (Time)	Indicates the elapsed time since the last write that the Secondary received from the Primary.
Current Logging	Indicates whether the SRL or DCM is currently in use (with respect to the Secondary),
Alert	Alert messages display if there is a problem with the configuration. For example, an alert message displays when there is some configuration error in the RDS.
Disk Group	The name of the Secondary disk group. For a shared disk group, (Shared) displays after the name.

Data volumes

Displays the data volumes associated with the Primary RVG or the Secondary RVG, if the Primary RVG is not present in the selected RDS. This section lists the following information for each data volume:


Name	Name of the data volume and the status icon.
Vset	Displays the name of the volume set, if the volume is a component volume of a volume set. Displays N/A, if the volume is an independent volume.
Size	Size of the data volume.
DCM	Yes or No depending on whether a DCM is associated to the data volume.

RDS detail view - Primary tab

To display the Primary tab of the RDS view, select the host name for the Primary from any VRW view. For example, click the host name `seattle`, to display the Primary tab of the RDS view for `hr_rvg`.

Figure 3-4 shows the Primary tab of the RDS view

Figure 3-4 Primary tab of the RDS view



Primary Replicated Volume Group

This section shows the details for the RVG on the Primary.

- Host

Displays Primary RVG icon and the name of the Primary host.

For a shared disk group, the name of the logowner displays as the Primary host.

If the Primary has not been configured correctly, this field is blank and the word `unknown` is displayed in the replication status column.

The Primary RVG icon indicates the I/O state.

See [Table 3-2](#) on page 51.
- Operating System

The operating system and version for the Primary host.
- Name

The name of the RVG (RDS name). The RVG state displays in parenthesis after the name.

See [“Primary RVG states”](#) on page 51.
- RVG Version

The RVG version.

Disk Group	The name of the Primary disk group. For a shared disk group, (Shared) displays after the name.
Disk Group Version	The disk group version for the RVG.
Data Volume Count	The number of data volumes in the RVG. If the RVG includes a volume set, the data volume count includes the number of the component volumes in the volume set in addition to the number of independent volumes.
Replicated Data Size	The total size of the replicated data volumes in the RDS. If a volume set is associated to the RDS, the size includes the total size of each of the component volumes of the RDS in addition to the total size of any independent volumes.
Checkpoint count	The number of Primary checkpoints on the RVG.
Checkpoints	The names and information about each Primary checkpoint. The information includes usage, % of log used, and the status of the checkpoint (started or complete).
Bunker	The name of the bunker host. Displays N/A if the RDS does not contain a bunker.
Cluster	The name of the cluster if the Primary RVG uses a shared disk group. Select the cluster name to view details about the cluster.

Storage Replication Log

This section displays the following information about the Storage Replication Log:

Name	The name of the SRL volume.
Size	The size of the SRL volume.

Replication Settings to Secondaries

This section shows the details about the replication settings for each Secondary RVG.

Secondary	Displays Secondary RVG icon and the name of the Secondary host. For a shared disk group, the name of the logowner displays as the Secondary host.
-----------	--

Replication IP	Name of the remote host and the IP address of the host in parenthesis.
Protocol	Transport protocol being used for replication.
Bandwidth Limit	Shows the bandwidth limit for VVR if specified, or none.
Replication Mode	Displays synchronous or asynchronous.
SRL Protection	Displays autodcm, dcm, override, off, or fail.
Latency Protection	Displays override, off, or fail depending on the chosen option.
Latency Low Mark	Number of updates in the SRL before the protection becomes inactive and updates succeed.
Latency High Mark	Specifies the maximum number of waiting updates in the SRL before the protection becomes active and updates stall or fail.

Data volumes

Displays the data volumes associated with the Primary RVG or the Secondary RVG, if the Primary RVG is not present in the selected RDS.

This section lists the following information for each data volume:


Name	Name of the data volume.
Vset	Displays the name of the volume set, if the volume is a component volume of a volume set. Displays N/A, if the volume is an independent volume.
Size	Size of the data volume.
DCM	Yes or No depending on whether a DCM is associated to the data volume.

RDS detail view - Secondaries tab

The Secondaries tab of the RDS detail view shows detailed information about the Secondaries within an RVG. To get to the Secondaries tab of the RDS detail view, click on a Secondary name from the Summary tab of the RDSs view or the Primary tab of the RDS view.

[Figure 3-5](#) shows the Secondaries tab of the RDS detail view.

Figure 3-5 Secondaries tab of the RDS detail view



The Secondaries tab of the RDS view displays the following information:

Replication status

Displays the status-related information of each Secondary in the selected RDS. This section displays the following information:

- Secondary
- Displays the Secondary RVG icon and the name of the Secondary host.

For a shared disk group, the name of the logowner displays as the Secondary host.

The Secondary RVG icon indicates the configuration state.
See [Table 3-2](#) on page 51.
- Data status
- Shows the data status.
See [Table 3-4](#) on page 52.
- Replication status
- Shows the replication status.
See [“Replication Status”](#) on page 52.
- Current mode
- Shows the replication mode, either synchronous or asynchronous.

% Log Usage	Shows either the percentage of the SRL that is full or the percentage of the DCM that is dirty.
Behind By (Size)	Indicates the size of the writes remaining to bring the Secondary up-to-date with the Primary.
Behind By (Time)	Indicates the elapsed time since the last write that the Secondary received from the Primary.
Current Logging	Indicates whether the SRL or DCM is currently in use (with respect to the Secondary),
Alert	Alert messages display if there is a problem with the configuration. For example, an alert message displays when there is some configuration error in the RDS.
Disk Group	The name of the Secondary disk group. For a shared disk group, (Shared) displays after the name.

Secondary Replicated Volume Group

This section shows the details for the RVG on the Secondary.

Host	Displays the name of the Secondary host. For a shared disk group, the name of the logowner displays as the Secondary host. The Secondary RVG icon indicates the configuration state. See Table 3-2 on page 51.
Name	The name of the RVG (RDS name). The RVG state displays in parenthesis after the name. See “Primary RVG states” on page 51.
Disk Group	The name of the Secondary disk group. For a shared disk group, (Shared) displays after the name.
Disk Group Version	The disk group version for the RVG.
RVG Version	The RVG version.
Operating System	The operating system and version for the Secondary host.

Data Volume Count	The number of data volumes in the Secondary RVG. If the RVG includes a volume set, the data volume count includes the number of the component volumes in the volume set in addition to the number of independent volumes.
Replicated Data Size	The total size of the replicated data volumes in the RDS. If a volume set is associated to the RDS, the size includes the total size of each of the component volumes of the RDS in addition to the total size of any independent volumes.
Bunker	The name of the bunker host. Displays N/A if the RDS does not contain a bunker.
Cluster	The name of the cluster if the Secondary RVG uses a shared disk group. Select the cluster name to view details about the cluster.

Storage Replication Log

This section displays the following information about the Storage Replication Log:

Name	The name of the SRL volume.
Size	The size of the SRL volume.

Replication settings to Primary

This section shows the details about the replication settings to the Primary for the specified Secondary RVG.

Primary	Name of the Primary
Replication IP	The IP address of the host.
Protocol	Transport protocol being used for replication.
Bandwidth Limit	Shows the bandwidth limit for VVR if specified, or none.
Replication Mode	Displays synchronous or asynchronous.
SRL Protection	Displays autodcm, dcm, override, off, or fail.
Latency Protection	Displays override, off, or fail depending on the chosen option.
Latency Low Mark	Number of updates in the SRL before the protection becomes inactive and updates succeed.

Latency High Mark	Specifies the maximum number of waiting updates in the SRL before the protection becomes active and updates stall or fail.
-------------------	--

Data volumes

Displays the data volumes associated with the Secondary RVG. This section lists the following information for each data volume:

Name	Name of the data volume and volume status icon.
Vset	Displays the name of the volume set, if the volume is a component volume of a volume set. Displays N/A, if the volume is an independent volume.
Size	Size of the data volume.
DCM	Yes or No depending on whether a DCM is associated to the data volume.

RDS detail view - Bunker tab

The bunker tab of the RDS view shows detailed information about the bunker for the RDS. To get to the bunker tab, click on a Bunker Secondary name from the Summary RDSs view or the Primary RDS view.

The bunker tab displays the following information:

Bunker replication status

Displays the status-related information of each Bunker in the selected RDS. This section displays the following information:

Bunker	Displays the Secondary RVG icon and the name of the Secondary host. For a shared disk group, the name of the logowner displays as the Secondary host. The Secondary RVG icon indicates the configuration state. See Table 3-2 on page 51.
Data status	Shows the data status. See Table 3-4 on page 52.
Replication status	Shows the replication status. See “Replication Status” on page 52.

Current mode	Shows the replication mode, either synchronous or asynchronous.
% Log Usage	Shows either the percentage of the SRL that is full or the percentage of the DCM that is dirty.
Behind By (Size)	Indicates the size of the writes remaining to bring the Secondary up-to-date with the Primary.
Behind By (Time)	Indicates the elapsed time since the last write that the Secondary received from the Primary.
Current Logging	Indicates whether the SRL or DCM is currently in use (with respect to the Secondary),
Alert	Alert messages display if there is a problem with the configuration. For example, an alert message displays when there is some configuration error in the RDS.
Disk Group	The name of the Bunker disk group. For a shared disk group, (Shared) displays after the name.

Bunker Replicated Volume Group

This section shows the details for the RVG on the bunker.

Host	Displays Secondary RVG icon and the name of the Secondary host. For a shared disk group, the name of the logowner displays as the Secondary host. The Secondary RVG icon indicates the configuration state. See Table 3-2 on page 51.
Name	The name of the RVG (RDS name). The RVG state displays in parenthesis after the name. See “Primary RVG states” on page 51.
Disk Group	The name of the Bunker disk group. For a shared disk group, (Shared) displays after the name.
Disk Group Version	The disk group version for the RVG.
RVG Version	The RVG version.
Operating System	The operating system and version for the Secondary host.
Cluster	The name of the cluster if the Secondary RVG uses a shared disk group. Select the cluster name to view details about the cluster.

Storage Replication Log

This section displays the following information about the Storage Replication Log:

Name	The name of the SRL volume.
Size	The size of the SRL volume.

Replication settings to Primary

This section shows the details about the replication settings to the Primary for the specified Secondary RVG.

Primary	Name of the Primary
Replication IP	The IP address of the host.
Protocol	Transport protocol being used for replication.
Bandwidth Limit	Shows the bandwidth limit for VVR if specified, or none.
Replication Mode	Displays synchronous or asynchronous.
SRL Protection	Displays autodcm, dcm, override, off, or fail.
Latency Protection	Displays override, off, or fail depending on the chosen option.
Latency Low Mark	Number of updates in the SRL before the protection becomes inactive and updates succeed.
Latency High Mark	Specifies the maximum number of waiting updates in the SRL before the protection becomes active and updates stall or fail.

RDS detail view - Clusters tab

The Clusters tab shows the cluster information if the Primary RVG or the Secondary RVG has a shared disk group. From the Summary RDSs view or RDS detail view, click on the name of a Primary RVG or Secondary RVG using a shared disk group to see the Clusters tab of the RDS detail view.

Cluster information

This section shows the details about the cluster for the specified RDS.

Name	Name of the cluster and status icon.
------	--------------------------------------

Log owner	Displays the logowner for the cluster.
Master	Displays the Master for the cluster.
Members	Displays the host names for the cluster.
Role	Indicates the role of the current host, for example, whether the host is a Primary or a Secondary.

RDS detail view - Alerts tab

The Alerts tab of the RDS view displays configuration problems for the selected RDS, if any. The alerts section displays the RDS name together with a brief description about the configuration error.

Information displayed in RDS views

The following sections show the information displayed in RDS views.

Conventions for the RDS icons

The icon that overlays the RDS icon indicates whether the Primary and the corresponding Secondary are configured correctly. The following table shows the icons for the two cases. If either the Primary or any of the Secondaries are not configured correctly, the error icon will appear.

Table 3-1 RDS_Configuration_State_Icons

Icon	Description
	Displays when the RDS does not have any configuration errors.
	Displays when the RDS has some configuration errors. See the <i>Veritas Volume Replicator Administrator's Guide</i> for details about possible mis-configurations in an RDS.

Conventions for the RVG icons

Different icons indicate the different RVG states that are displayed under different conditions. The following table explains the meaning of the icons that overlay the host icon to indicate the different states.

Table 3-2 Conventions for Primary I/O States or Secondary RVG Configuration States

Icon	Description
	Primary RVG. Displays when the Primary RVG is enabled for I/O and there are no configuration errors in the Primary RVG.
	Primary RVG. Displays when the Primary RVG is not enabled for I/O or there are configuration errors in the Primary RVG.
	Secondary RVG. Configuration is correct. Displays when replication to this Secondary is possible, that is, if the Secondary is configured correctly. In the case where a Primary RVG is an acting Secondary (that is, a Primary-Primary configuration), a Primary RVG icon displays with its I/O state overlaid.
	Secondary RVG. Configuration has an error. Displays when replication to this Secondary is not possible, that is, if the Secondary is not configured correctly.

Primary RVG states

The Primary RVG states field shows the current state of the Primary RVG.

Table 3-3 Primary RVG States

Value	Meaning
acting_secondary	This Primary RVG is currently the acting Secondary as part of the fast failback process. Writes to the data volumes in this RVG are disabled independent of whether the RVG is started or stopped.
disabled for I/O	Primary RVG is disabled for I/O, that is, the RVG is stopped.
enabled for I/O	Primary RVG is enabled for I/O, that is, RVG is started.
needs recovery	State after an import or reboot. The <code>vxrvrg recover rvg</code> command clears this state.
passthru	The Primary RVG is in passthru mode because the Primary SRL is detached, missing, or unavailable. For more information on RVG PASSTHRU mode, see the <i>Veritas Volume Replicator Administrator's Guide</i> .

Secondary RVG data status

The following table lists the values for the data status field and their meanings:

Table 3-4 Secondary RVG Data Status	
Value	Meaning
consistent, behind	Secondary data is consistent but not up-to-date with the Primary data. If the data status is "behind", the time stamp information is also shown. The time stamp information shows the hours, minutes, and seconds which the data writes are behind.
consistent, stale	The data on this Secondary is consistent. Replication to this Secondary has been stopped; the Primary RLINK is detached.
consistent, up-to-date	The Secondary data is consistent and is current or up-to-date with the Primary data. The Primary role can be migrated to this Secondary.
inconsistent	The data on the Secondary volumes is not consistent and the Secondary cannot take over.
needs recovery	State after an import or reboot. The <code>vxrlink recover</code> command clears this state.
N/A	Current state of the Secondary data cannot be determined. This may occur because of a configuration error on this Secondary. More information about the state can be found in the RLINK section of the Primary or Secondary RVG view.

Replication Status

The Replication Status field shows the current state of the replication between the Primary RVG and the Secondary RVG. The following table shows the values for this field and their meanings.

Table 3-5 Replication Status

Value	Meaning
logging to DCM	DCM is active for this Secondary, that is, new updates on Primary are tracked using DCM for this Secondary. The following information may be displayed: needs dcm resynchronization—To continue replication, resynchronize the Secondary using DCM resynchronization. See “Incrementally synchronizing the Secondary after SRL overflow” on page 76. needs failback synchronization—To continue replication, start failback synchronization to this Secondary. See “Failing back using fast-failback synchronization” on page 86. See the <i>Veritas Volume Replicator Administrator’s Guide</i>.
needs failback synchronization	This Primary RVG is acting as Secondary as part of the fast failback process. To continue replication, start failback resynchronization on the new Primary.
not replicating	Data is not being replicated to Secondary because Primary RLINK is in needs recovery state. Primary RLINK needs to be recovered before replication can resume.
paused by user	Replication to Secondary paused because of some administrative action that results in the following states: primary paused—Primary RLINK is paused. secondary paused—Secondary RLINK is paused.
paused due to error	Replication to Secondary paused because of the following errors: secondary config error—Secondary has some configuration error. secondary log error—Secondary SRL has an I/O error. See the <i>Veritas Volume Replicator Administrator’s Guide</i> for more information on interpreting RLINK flags.
paused due to network disconnection	Replication to Secondary is paused because of some network problem.
replicating	connected—Replication can take place if there are updates on the Primary data volumes.

Table 3-5 Replication Status (*continued*)

Value	Meaning
resync in progress	Resynchronization to the Secondary is in progress. autosync—Resynchronization type is autosync. dcm resynchronization—Resynchronization after an SRL overflow. failback synchronization—Resynchronization using failback logging.
resync paused by user	Resynchronization to Secondary paused because of some administrative action that results in the following states: primary paused—Primary RLINK is paused. secondary paused—Secondary RLINK is paused.
resync paused due to error	Resynchronization to Secondary paused because of the following errors: secondary config error—Secondary has some configuration error. secondary log error—Secondary SRL has an I/O error. See the <i>Veritas Volume Replicator Administrator's Guide</i> for more information on interpreting RLINK flags.
resync paused due to network disconnection	Resynchronization to Secondary is paused because of some network problem.
stopped	Replication to Secondary is stopped because of the following: primary detached—Primary RLINK is detached. secondary detached—Secondary RLINK is detached.
N/A	The replication status cannot be determined. More information about the replication status can be found in the Primary or Secondary RVG view.

Administering Veritas Volume Replicator

This chapter includes the following topics:

- [About administering Veritas Volume Replicator](#)
- [Administering data volumes](#)
- [Administering replication](#)
- [Administering the Replicated Data Set](#)
- [Administering checkpoints](#)
- [Administering the SRL](#)

About administering Veritas Volume Replicator

Volume Replicator Web Console (VRW) enables you to administer Veritas Volume Replicator (VVR) Secondary hosts, data volumes, and volumes from any host in a Replicated Data Set (RDS). VRW provides Web-based wizards that enable you to perform various tasks on an RDS. The specified tasks are performed on all the hosts.

Administering data volumes

VRW enables you to administer data volumes from any host in an RDS. You can perform the following volume administration tasks using VRW:

- [Associating volumes to a Replicated Data Set](#)
- [Resizing a data volume in a Replicated Data Set](#)

■ **Dissociating a data volume from its Replicated Data Set**

The VRW wizards that manage volume operations perform the operations on all RVGs in the RDS. The host from which the wizard is accessed is referred to as the local host. The name of the RVG on the local host represents the name of the RDS to which it belongs.

Note: If the associating or dissociating volume operation fails on any of the hosts in the RDS, the original configuration remains unchanged.

Associating volumes to a Replicated Data Set

Use the Associate Volume wizard to associate a volume or a volume set to a Replicated Data Set (RDS) even when replication is in progress. This command associates the selected volumes to all the RVGs of the RDS. For each volume you are associating, ensure that volumes of the same name and same length exist on all the hosts in the RDS. You must create volumes of the required layout on the Primary and the Secondary.

Before associating volumes, the Associate Volume wizard displays a warning and prompts the user to confirm whether or not the Primary and Secondary volumes contain the same data.

To verify that the Primary and Secondary volumes contain the same data, see the information provided in the *Veritas Volume Replicator Administrator's Guide*.

Observe the following prerequisites:

- | | |
|---------------|---|
| Prerequisites | Create volumes of the same name and length as the Primary volume on all the hosts in the RDS. |
| | Verify that the volumes to be associated are inactive. |
| | Synchronize the volumes. |

To associate a volume to an RDS

- 1 Select the RDS detail view for the RDS for which the data volume is to be associated.
- 2 From the Configuration menu in the task pane, select **Associate Volume**.
- 3 In the Volume Selection page, select the name of a data volume or a volume set to be associated to the RDS. Click **Next**.

See “[Associate Volume: Select Volume panel](#)” on page 57.

- 4 In the confirmation page, check the box to indicate that the prerequisites have been met and click **Next**. If the operation is successful, a message displays to indicate that each data volume has been successfully associated to each corresponding RVG in the RDS.
See [“Associate Volume: Confirmation panel”](#) on page 57.
- 5 In the summary page, click **Finish**.
- 6 VRW attempts the operation and displays the Result panel. Click **OK**.

Associate Volume: Select Volume panel

Use the Associate Volume wizard to associate a data volume or a volume set to an RDS, even when replication is in progress. This operation associates the selected volumes to all the RVGs of the RDS.

For each volume you are associating, ensure that volumes of the same name and same length exist on all the hosts in the RDS. You must create volumes of the required layout on the Primary and the Secondary.

RDS	Displays the name of the RDS to which you are adding volumes. This field cannot be edited.
Data Volumes	Select the name of a data volume or volume set to be associated to the RDS. For a volume set, you cannot associate an individual component volume of the volume set to the RDS, unless it is a component of a volume set that is already associated to the RDS. You can associate the entire volume set to the RDS, which associates all of the volumes in the volume set to the RDS. However, if there are component volumes that do not need to be replicated, you can remove them from the RDS later with the Dissociate Volume task.

Associate Volume: Confirmation panel

The Primary and Secondary volumes should contain the same data before the volumes are associated to the RDS. Before associating volumes, the Associate Volume wizard displays a warning and prompts the user to confirm whether or not the Primary and Secondary volumes contain the same data.

Observe the following prerequisites:

Prerequisites	Displays the prerequisites for associating volumes to an RDS. This field cannot be edited.
---------------	--

Prerequisites are met Check the box to indicate that the prerequisites have been met. Do not check the checkbox until you are certain that the prerequisites are actually met. Otherwise, associating the volumes may result in the Secondary having unusable data.

Associate Volume: Summary panel

Confirm the attributes for the operation. Click Finish to associate the volumes.

Resizing a data volume in a Replicated Data Set

Use the Resize Volume task to resize a data volume associated to a Replicated Data Set (RDS) even when replication is in progress. The Resize Volume task resizes the specified data volume in all the RVGs in an RDS.

Important notes for resizing a data volume:

- If the Primary data volume contains a file system, then resizing the data volume also resizes the file system.
- If replication to any Secondary RVG is in progress, resizing the volume pauses the Secondary, resizes the Primary data volume, and then resumes replication.
- If resizing the volume fails on any of the hosts in the RDS, original volume sizes are not restored. The volumes sizes won't match on the Primary and the corresponding Secondary hosts. To correct the mismatch, correct the error condition, run the Resize Volume wizard again to resize the volume, and resume replication.
- When you increase the size of a data volume, the newly added portions on the Primary and Secondary volumes are not synchronized.

Observe the following prerequisites:

Prerequisites: The data volume must exist in the disk group and be associated with the RVGs for all hosts in the RDS.

Make sure there is enough space in the disk group on the Primary and the Secondary to increase the size of a data volume.

To resize a data volume in an RDS

- 1 Select the RDS detail view for the RDS for which the data volume is to be resized.
- 2 Select the data volume to be resized.
- 3 Select **Resize Volume** in the drop-down list and click **Go**.

- 4 View the Warnings page and click **Next**.
- 5 Complete the Resize Volume Options page.
See “[Resize Volume: Volume Selection panel](#)” on page 59.
- 6 Click **Finish** to process the Resize Volume request.
- 7 Click **Ok** in the Result panel.

Resize Volume: Warnings panel

Use the Resize Volume wizard to resize a data volume associated to an RDS, even when replication is in progress. The Resize Volume task resizes the specified data volume in all the RVGs in an RDS.

Review the warnings and click Next to continue.

Resize Volume: Volume Selection panel

Use this panel to specify the data volume to be resized.

RDS	Displays the name of the RDS for which the data volume is to be resized. This field cannot be edited.
Data Volume	Select the name of the data volume to be resized. For example, <code>hr_dv01</code> .
New Size	Specify the size for the data volume. The default size unit is megabytes. To specify a size unit, select either kilobytes, or gigabytes from the drop-down list.
Force resize volume even if it may result in loss of data	Select the checkbox to force the resize volume operation.

Resize Volume: Summary panel

Confirm the attributes for the operation. Click Finish to resize the volumes.

Dissociating a data volume from its Replicated Data Set

Perform the Dissociate Volume task to remove data volumes from an RDS. The Dissociate Volume task only dissociates the data volumes from the RVGs in an RDS and does not delete them.

The Dissociate Volume task removes the specified data volumes from the RDS in either of the following conditions:

- When replication is in progress that is, the RLINKs are in the CONNECT state.
- When replication is stopped, that is, the RLINKs are detached. To stop replication, use Stop Replication.

Note: If you want the data volumes on the Primary and the Secondary to be consistent after they are dissociated, stop the application and make sure that the RLINKs are up to date before dissociating a data volume from its RDS.

To dissociate a data volume from its RDS even when the Primary RVG has not been stopped, select the Dissociate the data volume even if the Primary RVG has not been stopped option from the Dissociate Volume dialog box.

To dissociate a data volume from its RDS

- 1 Select any view for the RDS for which the data volume is to be dissociated.
- 2 In the Data Volumes section, select the volume.
- 3 Select Dissociate Volume from the drop-down list and click **Go**.
- 4 Complete the Dissociate Volume dialog, and click **Next**.
See “[Dissociate Volume dialog](#) ” on page 60.
- 5 Click **Next**.
- 6 On the Confirmation page, click **Finish** to dissociate the volume.
- 7 On the Result page, Click **OK**. VRW updates and displays the RDS or RVG view from which you accessed the Dissociate Volume task.

Dissociate Volume dialog

Use this dialog window to remove data volumes from an RDS. The Dissociate Volume task only dissociates the data volumes from the RVGs in an RDS and does not delete them.

RDS	Displays the name of the RDS for which you are dissociating the data volumes. This field cannot be edited.
Data Volume	Select the name of the data volume or volume set to be dissociated from the RDS. For a volume set, you can dissociate the entire volume set or you can dissociate a component volume from the RDS. Dissociating a component volume removes the volume from the RDS but does not remove it from the volume set.

Option: Select this option to force removal of a data volume. We recommend that the applications and the Primary RVG be stopped before dissociating any data volumes.

Dissociate the data volume even if the Primary RVG has not been stopped

Dissociate Volume: Summary panel

Use this dialog window to remove data volumes from an RDS. Review the summary information and click Finish to remove the data volumes.

Administering replication

VRW enables you to administer Secondary hosts from any host in the RDS. You can perform the following replication tasks using VRW:

- [Changing the replication settings](#)
- [Pausing replication to a Secondary](#)
- [Resuming replication to a Secondary](#)
- [Stopping replication to a Secondary](#)
- [Changing the IP addresses used for replication](#)

The replication tasks are independent of the host from which they are accessed and can be accessed from any host in an RDS. The host on which the task is launched is referred to as the local host. The name of the RVG on the local host represents the RDS name to which it belongs.

Changing the replication settings

This section describes how to change the VVR replication attributes according to your requirements. The Change Replication Settings task enables you to set the following VVR replication attributes:

- Replication mode
- Latency protection
- SRL protection
- Replication protocol
- Bandwidth limit
- Packet size

For details on selecting the appropriate option for each replication attribute, see the *Veritas Volume Replicator Administrator's Guide*.

Note that the replication settings are applied to both the Primary and the Secondary.

To change replication settings to a Secondary in an RDS

- 1

In the Secondaries detail view of the RDS, in the Replication Settings to Primary section, select the Primary.
- OR
- In the Primary detail view of the RDS, in the Replication Settings to the Secondaries section, select the Secondary.
- 2

Select **Change Replication Settings** and click **Go**.
- 3

Complete the Change Replication Settings page.
See [“Change Replication Settings”](#) on page 62.
- 4

Click **Finish** to change the replication settings for the selected Secondary. A message indicates whether the operation succeeded or failed.

In the Result page, click **OK**. The Secondary RVG view displays the new replication settings.

Change Replication Settings

Use this dialog window to change replication settings. To use the replication settings effectively in your environment, it is important to understand how each replication setting effects replication when the Primary and Secondary are connected and disconnected.

Secondary	Select the appropriate Secondary from the drop-down list.
Replication Mode	The current replication mode is displayed by default. From the drop-down list, select the required mode of replication.
Asynchronous,	
Synchronous-override	To set the replication to synchronous mode, select Synchronous-override. To set the replication to asynchronous mode, select Asynchronous.

SRL Protection autodcm, dcm, off, override	From the drop-down list, select the required value to change the SRL protection mode. To enable SRL protection, select autodcm, dcm, or override. To disable SRL protection, select off. If you disable SRL protection, the SRL is allowed to overflow. For details on selecting the fail option, see <i>Veritas Volume Replicator Administrator's Guide</i>.
Latency Protection override, off	To enable latency protection and limit the number of waiting updates in the SRL, select override. When you select override the wizard displays the fields Latency Low Mark and Latency High Mark with default values of 9950 and 10000 respectively. To disable latency protection, select off. This does not limit the number of waiting updates in the SRL. To enable latency protection and limit the number of waiting updates in the SRL, select override. For details on selecting the override or fail option, see <i>Veritas Volume Replicator Administrator's Guide</i>.
Latency Low Mark	Displays only if Latency Protection is set to override. Enter the number of writes in the SRL at which the protection must become inactive and writes must succeed.
Latency High Mark	Displays only if Latency Protection is set to override. Enter the numerical limit above which you do not want the number of updates waiting in the SRL to grow. Must be specified when Latency Protection is set to override.
Protocol	Displays the current network protocol used for replication. Select the desired protocol from the drop-down list.
Bandwidth Limit	Enter a value, or the word none. Select the desired units from the drop-down list. The minimum value is 56 kbps.
Packet Size	The default is 8400 bytes. The minimum value is 1300 bytes and the maximum value is 65464 bytes.

Starting replication to a Secondary

Use Start Replication to start replication to a Secondary. When you start replication to a Secondary, the data volumes on the Secondary must be synchronized with the data volumes on the Primary. For this reason, the Start Replication task is described in conjunction with the synchronization procedures.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

Start Replication dialog

Use the Start Replication dialog to start replication to a Secondary. When you start replication to a Secondary, the data volumes on the Secondary must be synchronized with the data volumes on the Primary. For this reason, the Start Replication task is described in conjunction with the synchronization procedures. Understanding the method used to start the replication is critical in order to correctly and efficiently start the replication.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

Tip: Resynchronization is not required when the Primary data volumes are empty or when the contents of the volumes on the Primary and Secondary are identical and the application is inactive.

Complete the Start Replication dialog according to the method of synchronization:

RDS	Displays the name of the RDS for which you are starting replication. This field cannot be edited. For example, hr_rvg
Secondary	Select the name of the Secondary to which you are starting replication.
Using automatic synchronization	Select this option to automatically synchronize the Secondary data volumes with the Primary data volumes in an RDS and start replication; automatic synchronization makes the Secondary data volumes up to date with the Primary data volumes. Replication to another Secondary can be started only after this automatic synchronization completes. Automatic synchronization copies the data on the Primary to the Secondary over the network using the Data Change Map (DCM). To use automatic synchronization, each data volume in the RVG must have an associated Data Change Map (DCM). Note: For component volumes of a volume set, use Start Replication with the Using Autosync option to synchronize the RVG.
Resynchronization is not required	Use this option to start replication only when the Primary data volumes are zero initialized or when the contents of the volumes on the Primary and Secondary are identical and the application is inactive.

Using checkpoint	This option enables you to start replication with a checkpoint. This option is used in conjunction with full synchronization, block-level tape backup, or difference-based synchronization of Primary data volumes with the Secondary data volumes. Select the checkpoint name that was created during the synchronization. This uses the selected checkpoint to start replication. If no checkpoints are available, this field displays No usable checkpoints.
------------------	---

See [“Setting up replication using automatic synchronization”](#) on page 27.

See [“Setting up replication using full synchronization”](#) on page 28.

See [“Setting up replication using block-level backup and checkpoint”](#) on page 29.

See [“Setting up replication using difference-based synchronization”](#) on page 30.

Pausing replication to a Secondary

Use Pause Replication to pause replication to a specific Secondary in an RDS. The Pause Replication task temporarily pauses replication between the Primary and Secondary. The pause is initiated from the Primary and is independent of the host from which the command is entered. Pausing replication from the Primary enables you to perform administrative tasks, such as making changes to the network connecting two nodes.

During a pause, the Primary continues to log writes to the SRL but does not send the writes to the Secondary. The network session between the Primary and Secondary on behalf of the RLINK is broken. A resume re-establishes the Primary-Secondary network session and enables writes to continue from the point of the pause. If there is volume activity during the pause, synchronous RLINKs are forced to become asynchronous. A previously synchronous RLINK remains asynchronous after the resume until it catches up.

To pause replication to a Secondary

- 1 Navigate to the Secondaries detail view or the Summary view of the required RDS.
- 2 In the Replication Status section, select the Secondary.
- 3 Select **Pause Replication**, and click **Go**.
- 4 In the Pause Replication dialog box, click **Finish** to pause replication. A message indicates whether the operation succeeded or failed.
- 5 In the Result page, click **OK**. The Secondary RVG view displays the new replication status.

Pause Replication dialog

Use Pause Replication to pause replication to a specific Secondary in an RDS.

RDS	Displays the name of the RDS for which you are pausing replication. This field cannot be edited. For example, hr_rvg
Secondary	Select the name of the Secondary to which you are pausing replication.

Resuming replication to a Secondary

Use Resume Replication to resume replication between the Primary and a specific paused Secondary in an RDS. After replication resumes, the state of the Primary and Secondary RLINKs changes to CONNECT to indicate that replication has resumed. To check the state of RLINKs, check the RLINK section in the Primary RVG view.

If the Secondary has been paused for a long period of time, the SRL may have overflowed and the Secondary may need to be resynchronized by replaying the DCM or by another method.

To resume replication to a Secondary

- 1 Navigate to the Secondaries detail view or the Summary view of the required RDS.
- 2 In the Replication Status section, select the Secondary.
- 3 Select **Resume Replication**, and click **Go**.
- 4 In the Resume Replication Summary page, click **Finish** to resume replication. A message indicates whether the operation succeeded or failed.
- 5 In the Result page, click **OK**. The Secondary RVG view displays the new replication status.

Resume Replication dialog

Use Resume Replication to resume replication between the Primary and a specific paused Secondary in an RDS.

RDS	Displays the name of the RDS for which you are resuming replication. This field cannot be edited. For example, hr_rvg.
Secondary	Select the name of the Secondary to which you are resuming replication.

Stopping replication to a Secondary

Use Stop Replication from any host in an RDS to stop replication between the Primary and a specified Secondary.

The Stop Replication task fails if the Primary and Secondary RLINKs are not up to date. The dialog box provides an option to stop replication to a Secondary when the RLINKs are not up to date.

Do not use the Stop Replication task to temporarily stop replication; instead, use Pause Replication.

Note: Stopping replication when the RLINKs are not up to date requires a complete resynchronization of the data volumes.

To stop replication to a Secondary

- 1 Navigate to the Secondaries detail view or the Summary view of the required RDS.
- 2 In the Replication Status section, select the Secondary.
- 3 Select **Stop Replication**, and click **Go**.
- 4 Complete the Stop Replication dialog box.
 See “[Stop Replication dialog](#)” on page 67.
- 5 Click **Finish** to stop replication. A message indicates whether the stop replication operation succeeded or failed.
- 6 In the Result page, click **OK**. The Secondaries detail view displays the new replication status.

Stop Replication dialog

Use Stop Replication from any host in an RDS to stop replication between the Primary and a specified Secondary.

RDS	Displays the name of the RDS for which you are stopping replication. This field cannot be edited. For example, hr_rvg
Secondary	Select the Secondary to which you are stopping replication.
Stop replication even when the Secondary is not up-to-date	Select the checkbox to stop replication to the Secondary when the Secondary is not up-to-date. Stopping replication when the Secondary is not up-to-date also stops tracking of writes to the SRL with respect to the Secondary. Starting replication again later may require a complete resynchronization of the data volumes.

Changing the IP addresses used for replication

Use Change IP to change the IP addresses used for replication to a specific Secondary in an RDS.

Observe the following prerequisites:

Prerequisites

The new host names must be configured for proper resolution on both the Primary and Secondary sites.

The Secondary must be reachable from the Primary either through the previous network, the new network, or both the networks.

If the previous network is no longer available, the change IP task will fail if you are not currently connected to the Primary host.

To change the replication IP

- 1 In the Secondaries tab of the RDS detail view of the RDS, in the Replication Settings to Primary section, select the Secondary for which you want to change the replication IP.

OR

In the Primary tab of the RDS detail view of the RDS, in the Replication Settings to the Secondaries section, select the Secondary for which you want to change the replication IP.

- 2 Select **Change IP** and click **Go**.
- 3 In the Change IP dialog box, enter a value for a new Primary or a new Secondary, or both. The value can be either a host name or an IP address.
See [“Change IP”](#) on page 68.
- 4 Click **Finish** to change the replication IP for the selected Secondary. A message indicates whether the operation succeeded or failed.
- 5 In the Result page, click **OK**. The RVG view displays the new replication settings.

Change IP

Use this dialog window to change the IP address used for replication. Enter a value for a new Primary or a new Secondary, or both. The value can be either a host name or an IP address.

RDS	Displays the name of the RDS for which you are changing the IP address. This field cannot be edited. For example, <code>hr_rvg</code>
Secondary	Displays the name of the Secondary RVG for the selected RDS. To change the IP address for another Secondary, select it from the drop-down list.
New Primary	Enter a value for a new Primary. The value can be either a host name or an IP address. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
New Secondary	Enter a value for a new Secondary. The value can be either a host name or an IP address. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.

Administering the Replicated Data Set

This section describes the following tasks:

- [Removing a Secondary from a Replicated Data Set](#)
- [Removing a Bunker from a Replicated Data Set](#)
- [Removing a Primary RVG](#)

Removing a Secondary from a Replicated Data Set

Use Remove Secondary to remove a Secondary RVG from the RDS to which it belongs. Before performing this operation, you must stop replication to the specified Secondary.

The Remove Secondary operation is irreversible.

The Remove Secondary task performs the following operations:

- Dissociates the data volumes and SRL from the Secondary RVG.

Note: The Remove Secondary task does not delete data volumes and the SRL; it only dissociates them from the Secondary RVG.

- Removes the Secondary RVG from its RDS.
- Deletes the Secondary RVG.

- Deletes the Primary and Secondary RLINKs.

Observe the following prerequisites:

Prerequisites	Stop replication to the Secondary that is to be removed. See “Stopping replication to a Secondary” on page 67.
---------------	---

To remove a Secondary from an RDS

- 1 Navigate to the Secondaries tab of the RDS view of the RDS for which the Secondary is to be removed.
- 2 In the task pane, under **Configuration**, select **Remove Secondary**.
- 3 In the Remove Secondary dialog box, click **Finish** to remove the Secondary. A message indicates whether the Remove Secondary operation succeeded or failed.
- 4 In the Result page, click **Ok**. The RDSs view or Summary tab of the RDS view displays the new RDS configuration. If the removed RVG was on the local host, VRW displays the Summary tab, which no longer shows the RDS in which the removed RVG belonged. Otherwise, VRW displays the current RDSs view.

Remove Secondary dialog window

Use this dialog window to remove a Secondary RVG from the RDS to which it belongs. Before performing this operation, you must stop replication to the specified Secondary.

RDS	Displays the name of the RDS from which you are removing the Secondary RVG. This field cannot be edited.
Secondary	Select the Secondary RVG to be removed.

Removing a Bunker from a Replicated Data Set

Use Remove Bunker to remove a Bunker Secondary RVG from the RDS to which it belongs. You must stop replication to the bunker before removing it.

The Remove Bunker operation is irreversible.

Observe the following prerequisites:

Prerequisites	Stop replication to the Bunker Secondary that is to be removed.
---------------	---

See [“Stopping replication to a Secondary”](#) on page 67.

To remove a Bunker Secondary from an RDS

- 1 Navigate to the Bunker tab of the RDS view of the RDS for which the Bunker is to be removed.
- 2 In the task pane, under **Configuration**, select **Remove Bunker**.
- 3 In the Remove Secondary dialog box, click **Finish** to remove the Secondary. A message indicates whether the Remove Secondary operation succeeded or failed.
- 4 In the Result page, click **Ok**. The RDSs view or Summary tab of the RDS view displays the new RDS configuration. If the removed RVG was on the local host, VRW displays the Summary tab, which no longer shows the RDS in which the removed RVG belonged. Otherwise, VRW displays the current RDSs view.

Remove Bunker dialog window

Use Remove Bunker wizard to remove a bunker Secondary from the RDS. You must stop replication to the bunker before removing it.

The Remove Bunker task is irreversible.

RDS	Displays the name of the RDS from which you are removing the bunker Secondary. This field cannot be edited.
Bunker Host	Displays the name of the bunker to be removed. This field cannot be edited.

Removing a Primary RVG

Use Remove Primary to remove a Primary RVG from an RDS and thus delete the corresponding RDS.

Observe the following prerequisite for deleting a Primary RVG:

- All Secondaries in the RDS must be removed.
See [“Removing a Secondary from a Replicated Data Set”](#) on page 69.

The Remove Primary task performs the following by default:

- Dissociates the data volumes and SRL from the Primary RVG.
- Removes the Primary RVG.

The Remove Primary task does not delete data volumes or the SRL from the Veritas Volume Manager configuration.

To remove a Primary RVG

- 1 Navigate to the Summary tab or the Primary tab of the RDS view of the RDS from which to remove the Primary RVG.
- 2 In the task pane, under **Configuration**, select **Remove Primary**.
- 3 In the Remove Primary dialog box, if required, select the "Force removal of the Primary RVG even if it is not yet stopped" checkbox.
- 4 Click **Finish**. A message indicates whether the Remove Primary operation succeeded or failed.
- 5 In the Result page, click **OK**. The Summary tab or Primary tab of the RDS view displays the new RDS configuration. If the removed RVG was on the local host, VRW displays the Summary tab, which no longer shows the RDS in which the removed RVG belonged. Otherwise, VRW displays the current RDSs view.

Remove Primary dialog window

Use this dialog window to remove a Primary RVG from the RDS to which it belongs and then delete the RDS. Before performing this operation, remove all Secondaries from the RDS. The Remove Primary operation also dissociates all data volumes and the SRL before removing the Primary.

RDS	Displays the name of the RDS from which you are removing the Primary RVG. This field cannot be edited.
Force removal of the Primary RVG	Check this option to force the removal of the Primary RVG even if the application is still active or even if the Primary RVG is not yet stopped.

Administering checkpoints

This section describes the following tasks:

- [Creating checkpoints](#)
- [Ending checkpoints](#)
- [Viewing checkpoints](#)
- [Deleting checkpoints](#)

Creating checkpoints

VRW enables you to create Primary checkpoints. Primary checkpoints are associated with a Primary RVG. You can create a maximum of 46 checkpoints.

Checkpoints cannot be nested. A started checkpoint must be ended before a new one can be started.

To start a checkpoint

- 1 Select the Primary tab of the RDS detail view.
- 2 Select **Start Checkpoint** and click **Go**.
- 3 In the Start Checkpoint dialog box, enter a checkpoint name and click **Finish**. A message displays the status of the checkpoint.
- 4 Click **OK**.

Create Checkpoint

Use this dialog window to create a Primary checkpoint. Primary checkpoints are associated with a Primary RVG. You can create a maximum of 46 checkpoints.

RDS	Displays the name of the RDS for which you are creating the checkpoint. This field cannot be edited. For example, <code>hr_rvg</code>
Checkpoint Name	Enter the name to be used for the Primary checkpoint. A valid checkpoint name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 19 characters long. It cannot begin with a dash or a period, or end with a period. For example, <code>my_checkpoint</code>

Ending checkpoints

VRW enables you to end a checkpoint. For example, if a checkpoint was started before a tape backup on the Primary, end the checkpoint when the tape backup completes on the Primary. The checkpoint can then be used to restore the tape on the Secondary.

You must end any checkpoint that is started before creating a new checkpoint.

To end a checkpoint

- 1 Select the Primary tab of the RDS detail view.
- 2 Select **End Checkpoint** and click **Go**.

- 3
- In the End Checkpoint dialog box, click **Finish** to confirm that you want to end the checkpoint which is currently started. A message displays the status of the checkpoint.
- 4
- Click **Ok**.

End Checkpoint

Use this dialog window to end a checkpoint.

RDS	Displays the name of the RDS for which the checkpoint is to be ended. This field cannot be edited. For example, my_rvg.
Checkpoint	Displays the name of the checkpoint that has been started. This field cannot be edited. For example: <code>my_checkpt (Started, Usage=0, % Log=0)</code>

Viewing checkpoints

Checkpoints are displayed in the Primary tab of the RDS detail view for the RDS.

Checkpoint Count	Indicates the number of checkpoints available.
Checkpoints	Lists each checkpoint, and displays usage, % of log used, and the status of the checkpoint (started or complete).

Deleting checkpoints

VRW enables you to delete an existing checkpoint. Checkpoints can be deleted after they are no longer needed. If you reach the maximum number of checkpoints, you must delete a checkpoint before adding a new one. You can create a maximum of 46 checkpoints.

To delete a checkpoint

- 1
- Select the Primary tab of the RDS detail view.
- 2
- Select **Delete Checkpoint**.
- 3
- In the Delete Checkpoint dialog box, select the checkpoint from the drop-down box. Select the option box if the checkpoint has not been checkpoint ended and you still want to delete it.
- See [“Delete Checkpoint”](#) on page 75.

- 4 Click **Finish**. A message displays the status of the checkpoint delete operation.
- 5 Click **OK**.

Delete Checkpoint

Use this dialog window to delete an existing checkpoint.

RDS	Displays the name of the RDS for which you are deleting the checkpoint. This field cannot be edited. For example, hr_rvg
Checkpoint Name	Select the checkpoint to be deleted from the drop-down list of existing checkpoints. For example, hr_rvg If no checkpoint exists, the following message displays: There is no checkpoint to be deleted.
Delete checkpoint even if it is not yet checkended or not yet over flowed	Select this option to force removal of a checkpoint.

Administering the SRL

The size of the SRL is critical to the performance of replication. When the SRL overflows for a particular Secondary, the Secondary becomes out of date until a complete resynchronization with the Primary is performed. Because resynchronization is a time-consuming process and during this time the data on the Secondary cannot be used, it is important to prevent the SRL from overflowing. Hence, when initially configuring VVR, determine an appropriate size for the SRL. The maximum size of the SRL can be derived from various criteria, however, the size of the SRL volume cannot be less than 110 MB. If the size that you have specified for the SRL is less than 110MB, VVR displays a message that prompts you to specify a value that is equal to or greater than 110 MB.

For more information, see "Sizing the SRL" in the *Veritas Volume Replicator Planning and Tuning Guide*.

It is possible that an SRL of an appropriate size overflows because of changes in the environment. This section describes how to protect from SRL overflows and administer VVR if the SRL overflows.

Protecting from SRL overflow

To avoid complete synchronization of Secondary in the case of an SRL overflow, VVR provides autodcm or dcm mode of SRL protection.

For more information, see "Protection Against SRL Overflow—srlprot attribute" in the *Veritas Volume Replicator Administrator's Guide*.

To enable SRL protection, each data volume in the RDS must have an associated DCM.

For more information, see "Associating a Data Change Map to a Data Volume" in the *Veritas Volume Replicator Administrator's Guide*.

See ["Changing the replication settings"](#) on page 61.

Incrementally synchronizing the Secondary after SRL overflow

The default protection mode for the SRL is autodcm and every data volume in the RVG must have a DCM. When the SRL fills up, whether the RLINK is connected or not, DCM logging is activated and a bit corresponding to the region of the update is turned on for every incoming update. When you are ready to replay the DCM, start the DCM resynchronization process. To start the resynchronization, use the Replay DCM wizard.

Note that you can also use the cache or cache size parameters with the Replay DCM wizard. If you specify these attributes, the wizard first creates a space-optimized snapshot of the Secondary data volumes before starting the resynchronization. Data is transmitted to the Secondaries only after all the RLINKs taking part in the resynchronization have connected. All the Secondaries taking part in the resynchronization must remain connected for the resynchronization to continue. The resynchronization will pause if any of the Secondary RLINKs are paused. During DCM resynchronization, VVR does not maintain the order of updates to the Secondary. As a result, the Secondary remains inconsistent until the resynchronization operation is complete. Note that if the Primary becomes unavailable during the time the resynchronization is taking place, the applications cannot be restarted on the Secondary.

If the Secondary data volumes are mirrored, you can break off mirrors to retain consistent (though out-of-date) copies of data until the resynchronization is complete. However, to overcome this problem, create snapshots of the Secondary data volumes before the resynchronization starts.

The cache attribute specifies a name for the precreated cache object, on which the snapshots for the volumes in the specified RVG will be created. The cachesize attribute specifies a default size for the cache object with respect to the source

volume. You can specify only one of these attributes at one time with the Replay DCM wizard to create one cache object for each snapshot.

If you do not specify either of these options then the Replay DCM wizard resynchronizes the Secondary data volumes without creating the snapshots.

Replay DCM starts replay of a DCM that is active because of an SRL overflow or failback logging. Replay occurs on all RLINKs on which the `dcm_logging` flag is set. If any of these RLINKs are disconnected or paused, replay DCM is delayed until this condition clears on all RLINKs.

Observe the following prerequisite for incrementally resynchronizing the Secondary:

Prerequisites	The RVG must have the <code>dcm_logging</code> flag set.
---------------	--

To incrementally resynchronize the Secondary

- 1 From any view of the RDS for which the resynchronization is to be started, in the task bar, under Replication, select **Replay DCM**.
- 2 Select one of the options on the Summary page.
See [“Replay DCM Wizard”](#) on page 77.
- 3 Click **Finish** to start the incremental synchronization.
Refresh the RDS view to monitor the progress of the incremental synchronization. The RDS view shows how much data is left to send to each Secondary that is being resynchronized.

Replay DCM Wizard

Use Replay DCM to start replay of a DCM that is active because of an SRL overflow or failback logging. Replay occurs on all RLINKs on which the `dcm_logging` flag is set. If any of these RLINKs are disconnected or paused, replay DCM is delayed until this condition clears on all RLINKs.

Cache Name	Select Cache Name and specify the name for a pre-created cache object on which the snapshots for the data volumes in the Secondary RVG will be created.
Cache Size	Select Cache Size and specify a default size for the cache object with respect to the source volume. The default is Megabytes. Select another unit from the drop-down menu, if necessary.

none	Use this option if you have a consistent mirror of the Secondary data volumes and do not need to use snapshots, in which case break off the mirror first so that a consistent copy of the data is available in case something went wrong during the resynchronization.
------	--

Resizing the SRL

The size of the SRL is critical to the performance of replication. When initially configuring VVR, carefully determine an appropriate size for the SRL. If the SRL overflows for a particular Secondary, the Secondary becomes out of date and inconsistent until a complete resynchronization with the Primary is performed. Because resynchronization is a time-consuming process and during this time the data on the Secondary cannot be used, the SRL needs to be an appropriate size to prevent the SRL from overflowing.

If the SRL needs to be resized, use the Resize SRL dialog window to increase the size of the SRL. The SRL can be increased while the application is active or while replication is in progress.

Note: This operation is irreversible because the SRL size can only be resized to a larger size.

Observe the following prerequisites:

Prerequisites	Make sure there is enough space in the disk group on the Primary to increase the size of the SRL.
---------------	---

To resize the SRL

- 1 From any view of the RDS for which the resynchronization is to be started, in the task bar, under Configuration, select **Resize SRL**.
- 2 Complete the Resize SRL dialog box, and click **Next**.
See “[Resize SRL dialog window](#)” on page 78.
- 3 Click **Finish**.
- 4 Click **Ok**.

Resize SRL dialog window

Use the Resize SRL dialog window to increase the size of the SRL. The SRL can be increased while the application is active or while replication is in progress.

Note: This operation is irreversible because the SRL size can only be resized to a larger size.

SRL	Displays the name of the SRL volume to be resized. This field cannot be edited.
RDS	Displays the name of the RDS which contains the SRL. This field cannot be edited.
New Size	Enter the required size. The default size unit is megabytes. To specify a size unit, select kilobytes or gigabytes from the drop-down list. The minimum size of the SRL volume is 110 MB. The maximum size of the SRL can be derived from various criteria. Note: This operation is irreversible because the SRL size can only be resized to a larger size.

Resize SRL dialog window: confirmation panel

Confirm the attributes for the operation. Click Finish to resize the SRL.

Transferring the Primary role

This chapter includes the following topics:

- [About transferring the Primary role](#)
- [Migrating the Primary role](#)
- [Taking over from an original Primary](#)
- [Failing back to the original Primary](#)
- [Using a bunker for disaster recovery](#)

About transferring the Primary role

The Volume Replicator Web Console (VRW) wizards enable you to transfer the Primary role from a healthy or failed Primary to another in an RDS. VRW provides the following methods to transfer the Primary role:

- [Migrating the Primary role](#)
- [Taking over from an original Primary](#)
- [Failing back to the original Primary](#)

For details about each method, see the *Veritas Volume Replicator Administrator's Guide*.

Migrating the Primary role

Migration involves transferring a healthy Primary of a Replicated Data Set (RDS) to a Secondary when the application involved in replication is inactive. Migration

of a Primary to a Secondary is useful when the Primary must be brought down for reasons such as maintenance or to make the application active on another node.

Use the Migrate task to migrate the role of a healthy Primary to a specified Secondary. The Migrate task performs the following functions:

- Migrates the Primary role of an RDS to the Secondary RVG.
- Converts the Secondary RVG to the Primary RVG.
- Converts the original Primary of the RDS to Secondary in the RDS.
- Reconfigures only the original Primary and the new Primary.

Note: If an RDS contains multiple Secondaries, VRW does not automatically reconfigure the additional Secondary hosts as Secondary hosts of the new Primary. If required, you can reconfigure the additional Secondary hosts as Secondary hosts of the new Primary using the VVR command-line interface.

If the Migrate task fails on any of the hosts, the original configuration is restored.

Before migrating the Primary, be sure the following prerequisites are met:

Prerequisites	The data volumes in the RDS must be inactive; that is, any applications that use the Primary data volumes must be stopped. All active Secondaries (attached RLINKs) must be up-to-date. All RLINKs to the Secondary must be in the CONNECT state.
---------------	--

To migrate a healthy Primary

- 1 Stop all the applications that are involved in replication. For example, if the application is a file system, unmount it.
- 2 Verify that the RLINKs are up to date by checking the replication status in the RDS view. The Replication Status of the Secondary must be UP-TO-DATE. The Migrate operation fails if:
 - the Primary attached RLINKs are not in the UP TO DATE and CONNECT state.
 - the data volumes are active.
- 3 Navigate to the Summary view, Primary tab of the RDS view or Secondaries tab of the RDS view for the RDS.
- 4 In the Role Change section of the task pane, select **Migrate**.

- 5 In the Migration dialog box, select a Secondary to which you want to migrate the Primary role, and then click **Next**.
- 6 On the Summary page, click **Finish**. A message indicates whether the migration succeeded or failed.
- 7 In the Result page, click **Close**. The Primary RVG view for the new Primary is displayed. By default, the migration operation starts replication from the new Primary.
- 8 Restart the application. Because the application was stopped before the migration, an application recovery is not required. To start the application before enabling replication, perform the following steps:
 - Pause replication.
 - Start the application.
 - Resume replication.

Migrate Role: Prerequisites panel

Use the Migrate dialog window to migrate the Primary role from the Primary host to a Secondary host.

RDS	Displays the name of the RDS for which you are migrating the Primary role. This field cannot be edited. Example hr_rvg
Secondary	Select a Secondary to which you want to migrate the Primary role. The selected Secondary must be up-to-date. After the migration, this Secondary becomes the new Primary, and the original Primary becomes the Secondary. Example london

Migrate Role: Summary panel

View the configuration. Click Finish to migrate the Primary.

Taking over from an original Primary

The takeover procedure involves transferring the Primary role from an original Primary to a Secondary. When the original Primary fails or is destroyed because of a disaster, the takeover procedure enables you to convert a consistent Secondary to a Primary. The takeover of a Primary role by a Secondary is useful when the Primary experiences unscheduled downtimes or is destroyed because of a disaster.

Use the takeover procedure to transfer the Primary role from an original Primary to a Secondary. The takeover task must be performed on the Secondary host. The takeover task fails when performed on a Primary host.

The takeover procedure performs the following functions on the RDS to which the original Primary belongs:

- Converts the Secondary RVG to a Primary RVG.
- Removes the original Primary from the RDS.
- Enables the fast-failback feature on the new Primary depending on the failback option that you select from the Takeover wizard. Each data volume in the Secondary RVG must have an associated Data Change Map (DCM) to enable fast failback with or without the autofailback.

See [“Takeover dialog window”](#) on page 84.

To take over from the original Primary to the Secondary

- 1 Make sure that the data volumes on the Secondary host have associated DCMs if you want to use the fast failback logging feature.
- 2 Navigate to the Summary tab, Primary tab, or Secondaries tab of the RDS view of the RDS.
- 3 In the Role Change section of the task pane, select **Takeover**. The Takeover dialog box enables you to convert the selected Secondary to be the new Primary.
- 4 Complete the Takeover Options page.
See [“Takeover dialog window”](#) on page 84.
- 5 Click **Finish** to proceed with the Takeover operation. A message displays to indicate whether the takeover succeeded or failed.
- 6 Click **Ok**.
- 7 Verify whether fast-failback is enabled. If it is, the `dcm_logging` flag is displayed in the **Flags** column of the Primary RLINK in Primary RVG view.
- 8 If the takeover succeeds, start the application on the new Primary. Starting the applications on the new Primary after a takeover may require an application recovery to be run.

Takeover dialog window

Use this dialog window to transfer the Primary role from an original Primary to a Secondary. The Takeover dialog window enables you to specify the method you plan to use for failing back to the original Primary. We recommend that you use

the fast failback synchronization method to synchronize the original Primary with the new Primary.

See [“Failing back using fast-failback synchronization”](#) on page 86.

The Takeover dialog box provides the following options to fail back to the original Primary:

RDS	Displays the name of the RDS. This field cannot be edited.
Fast failback	This is the default. Use this option when you want to incrementally synchronize the new Primary with the original Primary. This option enables fast failback; however, the original Primary is not automatically converted to a Secondary after reboot. To synchronize the data volumes on the original Primary and convert the original Primary to a Secondary, use the Replay Failback Log task on the new Primary.
Auto fast failback	The Auto fast failback option enables you to automatically synchronize the original Primary when it becomes available. This option converts the original Primary to a Secondary when the original Primary becomes available after an unexpected failure. It also uses the DCM to automatically synchronize the data volumes on the original Primary, which is now the new Secondary, using fast failback. Each Secondary data volume must have an associated DCM.
No fast failback	Use the No fast failback option to change the role from Secondary to Primary without enabling fast failback. Select this option if you are sure that the original Primary will not recover or if most of the data on the new Primary is going to change while the original Primary is unavailable. This option does not automatically synchronize the original Primary. You must synchronize the original Primary after it becomes available either using difference-based synchronization or full synchronization. Use the <code>vradmin syncrvg</code> command to synchronize the original Primary using difference-based synchronization or full synchronization. See the <i>Veritas Volume Replicator Administrator's Guide</i> for details.

Failing back to the original Primary

After an unexpected failure, a failed Primary host might become available and find that one of its Secondary hosts has been promoted to a Primary as a result of a takeover. A takeover happens when a Secondary of the original Primary takes

over the Primary role because of an unexpected outage of the original Primary. The process of transferring the Primary role back to the original Primary is called failback.

VRW provides the following methods to fail back to the original Primary:

- [Failing back using fast-failback synchronization](#)
- [Failing back using difference-based synchronization](#)

Note: We recommend that you use the fast failback synchronization method.

For a comparison of fast failback versus difference-based synchronization, see the *Veritas Volume Replicator Administrator's Guide*.

Failing back using fast-failback synchronization

Before using the fast-failback synchronization method to fail back to the original Primary, verify whether the fast-failback feature was enabled when taking over the Primary role.

For more information about the fast-failback feature, see the chapter "Transferring the Primary Role" in the *Veritas Volume Replicator Administrator's Guide*.

To fail back to the original Primary using fast-failback synchronization

- 1 Synchronize the data volumes on the original Primary, which is now the new Secondary, with the data volumes on the new Primary using the Replay Failback Log dialog box. Do not perform this step if **Auto-fast failback** is selected in the Takeover dialog box during the takeover.
- 2 Stop the application on the new Primary.
- 3 Migrate the Primary Role to the original or failed Primary.

See ["Migrating the Primary role"](#) on page 81.

Example—failing back to the original Primary using fast failback

In this example, the Primary host `seattle` has restarted after an unexpected failure. After the failure, the Primary role was taken over by the Secondary host `london`. Each data volume on the Secondary `london` has a Data Change Map (DCM) associated with it and fast failback is enabled on `london` by selecting the Auto-fast failback option in the Takeover dialog box.

An application is running on `london` and incoming writes are logged to its DCM. This example shows how to fail back to the original Primary `seattle` using the fast-failback feature.

To fail back to the original Primary `london` using fast failback

- 1 Navigate to the Summary, Primary, or Secondaries tab of the RDS view of the RDS.
- 2 In the Role Change section of the task pane, select **Replay Failback Log**.
- 3 Complete the Replay Failback Log dialog box.
See [“Replay Failback Log dialog box”](#) on page 87.
- 4 Click **Finish**.

The Replay Failback Log wizard synchronizes the data volumes in the new Secondary RVG `hr_rvg` on `seattle` with the data volumes in the new Primary RVG `hr_rvg` on `london` using the fast-failback feature. This step is not required if **Auto fast failback** was selected in the Takeover dialog box when taking over from the original Primary.

- 5 Click **OK**. You can check the status of the synchronization using the RDS Summary view. Refresh the view to ensure the latest status is displayed. When the synchronization completes, go to the next step.
- 6 At a convenient time, stop the application on the new Primary `london`.
- 7 Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle` using the Migration task.

See [“Migrating the Primary role”](#) on page 81.

Replication from the original Primary `seattle` to the original Secondary `london` is started by default after the migration is completed.

Replay Failback Log dialog box

Use the Replay Failback Log wizard to synchronize the data volumes in the new Secondary RVG with the data volumes in the new Primary RVG using the fast-failback feature.

Cache Size	Select Cache Size and specify a default size for the cache object. The default is Megabytes. Select another unit from the drop-down menu, if necessary.
Cache Name	Select Cache Name and specify the name for a pre-created cache object.

None

Select this option if you are not using a cache.

Failing back using difference-based synchronization

This section describes failing back using difference-based synchronization.

To fail back to the original Primary using difference-based synchronization

- 1 Convert the original Primary to a Secondary using the Make Secondary task.
- 2 See [“Converting a Primary to a Secondary”](#) on page 88.
- 3 Synchronize the data volumes on the original Primary with the data volumes on the new Primary using the checkpoint option of the `vradmin syncrvg` command to synchronize the original Primary using difference-based synchronization.

See the *Veritas Volume Replicator Administrator's Guide* for details.

- 4 Start replication to the new Secondary with the checkpoint by completing the Start Replication task.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

- 5 Stop the application on the new Primary `london`.
- 6 Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle`. Replication from the original Primary `seattle` to the original Secondary `london` is started by default.

See [“Migrating the Primary role”](#) on page 81.

Converting a Primary to a Secondary

Use the Make Secondary task to convert an original Primary to a Secondary. The Make Secondary task can be launched from the original Primary only when one of its Secondary hosts has taken over the Primary role.

The Make Secondary task is used in the failback procedure to fail back to the original Primary. When the original Primary restarts, use the Make Secondary task to convert the original Primary to a new Secondary. Stop the application if the application restarts automatically when the Primary restarts.

Note: Use the Make Secondary task to fail back to the original Primary only if No fast failback was selected during the takeover from the original Primary.

To convert an original Primary to a Secondary

- 1 Navigate to the Summary, Primary, or Secondaries tab of the RDS view of the RDS on the Primary host, which is the original Primary that failed and is now restarted.
- 2 In the Role Change section of the task pane, select **Make Secondary**.
- 3 Complete the Make Secondary page.
See [“Make Secondary dialog window”](#) on page 89.
- 4 Click **Finish** to convert the original Primary to a Secondary. A message displays to indicate whether the takeover succeeded or failed.
- 5 Click **Ok**.

Make Secondary dialog window

Use this dialog window to convert an original Primary to a Secondary. The Make Secondary task is part of the procedure to fail back to the original Primary. After the original Primary recovers, the RDS is in a Primary-Primary configuration. The original Primary is converted to be a Secondary to the new Primary.

The Make Secondary task is only used if the No fast failback option was selected during the takeover of the original Primary.

RDS	Displays the name of the RDS which contains the original Primary. This field cannot be edited. Example: <code>hr_rvg</code>
Secondary	From the drop-down list, select the name of the original Primary to be converted to a Secondary. If the name of the original Primary is not available in the drop-down list, type it into the field. A value in the entry field takes precedence over the selected value in the selection box.

Example: Failing back using difference-based synchronization

In this example, the Primary host `london` has restarted after an unexpected failure. After the failure, the Secondary host `seattle` has manually taken over the Primary role. This example shows how to fail back to the original Primary `london` using difference-based synchronization.

To fail back using difference-based synchronization

- 1 From the RDS or RVG view, choose **Make Secondary**, and complete the Make Secondary page. Note: In the **New Primary** field, select or enter **london**.

See [“Converting a Primary to a Secondary”](#) on page 88.

The Make Secondary task converts the original Primary RVG `hr_rvg` on `seattle` to the Secondary RVG of the new Primary `london`.

- 2 Synchronize the data volumes on the original Primary with the data volumes on the new Primary using the checkpoint option of the `vradmin syncrvg` command to synchronize the original Primary using difference-based synchronization.

See the *Veritas Volume Replicator Administrator's Guide* for details.

The Synchronize RVG task synchronizes the data volumes in the RVG `hr_rvg` on `seattle` with the data volumes on the new Primary RVG `hr_rvg` on `london` using the difference-based synchronization and checkpoint.

- 3 Start replication to the Secondary RVG (original Primary) `hr_rvg` on `seattle` from the new Primary RVG `hr_rvg` on `london`.
- 4 In the Start Replication page, select the **Using checkpoint** option, and select the checkpoint `ckpt_presync` from the drop-down list.
- 5 Stop the applications on the new Primary `london`.
- 6 Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle` using **Migrate**.

See [“Migrating the Primary role”](#) on page 81.

Replication from the original Primary `seattle` to the original Secondary `london` is started by default after the migration is completed.

Using a bunker for disaster recovery

If the Primary fails, update the Secondary from the bunker.

See [Updating the Secondary from the bunker](#)

After the Secondary is up-to-date, the Secondary can take over the role of the Primary.

See [“Taking over from an original Primary”](#) on page 83.

When the Primary recovers, fail back to the original Primary.

See [“Failing back to the original Primary”](#) on page 85.

Updating the Secondary from the bunker

When disaster strikes and the Primary goes down, the Secondary can be updated using the bunker.

Note: If the Primary SRL has overflowed for a Secondary, or if the Secondary is inconsistent because it is resynchronizing, you cannot use the corresponding bunker SRL to recover the Secondary. Because the bunker does not have data volumes, it cannot use DCM to track overflows.

To update the Secondary from the bunker

- 1 If replication to the bunker SRL uses the STORAGE protocol, import the disk group containing the bunker SRL onto the bunker node and then recover it using either the Veritas Volume Manager CLI or the Veritas Volume Manager GUI.

For instructions, see the *Veritas Volume Manager Administrator's Guide* or the *Veritas Volume Manager Online Help*.

- 2 Activate the bunker using the Activate Bunker task.

See [“Activating a bunker in an RDS”](#) on page 92.

- 3 Start replication to the Secondary from the bunker using the Start Replication task.

See [“Synchronizing the Secondary and starting replication”](#) on page 27.

If you have more than one Secondary using this bunker, repeat the Start Replication task for each Secondary.

- 4 Monitor the status of the replication from bunker to Secondary using the RDS view.

- 5 When the Secondary is up-to-date, stop replication to the Secondary.

See [“Stopping replication to a Secondary”](#) on page 67.

All the replays from the bunker must have been stopped before you deactivate the bunker.

Do not proceed to step 6 until the replication is stopped.

- 6 After the bunker has been used to transfer data to the secondary and is not needed any more, deactivate the bunker using the Deactivate Bunker task.

See [“Deactivating a bunker”](#) on page 92.

The Secondary is now up-to-date and can take over as Primary.

See [“Taking over from an original Primary”](#) on page 83.

Activating a bunker in an RDS

Use the Activate Bunker task to convert the bunker Secondary from receiving mode (Secondary) to replicating mode (Primary).

To activate a bunker in an RDS

- 1 Make sure you have logged on to the Bunker host. To log on to a host from VRW, choose **File > Connect** and specify the name of the host to which you want to connect.
- 2 In any view, select the Bunker host name from the view, and then choose **Activate Bunker** from the **Role Change** menu in the task pane.
- 3 Click **Finish**. A message indicates whether the activate bunker task succeeded or failed.

See [“Using a bunker for disaster recovery”](#) on page 90.

See [“Deactivating a bunker”](#) on page 92.

Activate Bunker

Use the Activate Bunker task to convert the bunker Secondary from receiving mode (Secondary) to replicating mode (Primary).

RDS	Displays the name of the RDS. This field cannot be edited. Example: my_rvg.
Bunker Host	Displays the name of the bunker to be activated. This field cannot be edited. Example: my_bunker.

Deactivating a bunker

Use the Deactivate Bunker task to deactivate the bunker after the bunker has been used to transfer data to the Secondary and it is not needed for any more replays.

Note: Deactivate the bunker only after all the replays from the bunker have been stopped.

To deactivate a bunker in an RDS

- 1 Make sure you have logged on to the Bunker host. To log on to a host from VRW, choose **File > Connect** and specify the name of the host to which you want to connect.
- 2 In any view, select the Bunker host name from the view, and then choose **Deactivate Bunker** from the **Role Change** menu in the task pane.
- 3 Click **Finish**. A message indicates whether the deactivate bunker task succeeded or failed.

Deactivate Bunker

Use the Deactivate Bunker task to deactivate the bunker after the bunker has been used to transfer data to the Secondary and it is not needed for any more replays.

Note: Deactivate the bunker only after all the replays from the bunker have been stopped.

RDS	Displays the name of the RDS. This field cannot be edited. Example my_rvg
Bunker Host	Displays the name of the bunker to be deactivated. This field cannot be edited.

Troubleshooting VRW

This appendix includes the following topics:

- [About troubleshooting VRW](#)
- [Application page is not displayed in the browser](#)
- [Missing buttons on certain wizards](#)
- [Log in redirects the user back to the login page](#)

About troubleshooting VRW

This chapter describes the errors that may occur when using VRW and the solutions to recover from these errors.

Application page is not displayed in the browser

This section describes an error that may occur when using VRW and the solutions to recover from this error.

Message

The following message displays:

The page cannot be displayed.

Problem

You may get this error message after you have specified the URL `http://london:8443/vvr` in the address field of your browser.

Solution

This section provides a workaround.

To recover from the page display error

- 1 Determine if the client on which the browser is running has `http` and `https` connection to the server running VRTSweb.

```
C:\> ping london  
Unknown host london.
```

The message `Unknown host london.` displays if there is no connection between the client and the server running VRTSweb.

Establish the connection and reconfigure the client so that the client can ping the VRTSweb server.

- 2 Determine the port numbers configured for use by VRTSweb.

```
# /opt/VRTSweb/bin/webgui listports  
8181 http  
8443 https
```

By default, VRTSweb is configured to use the port number 8181 (for HTTP) and 8443 (for HTTPS) as shown in the example.

- 3 Make sure that VRTSweb is running by issuing the following command:

```
# /opt/VRTSweb/bin/monitorApp VEAWeb  
Web Application "VEAWeb" is OFFLINE
```

The message `Web Application "VEAWeb" is OFFLINE` displays if the VRTSweb is not running. If the VRTSweb is offline, bring it online by starting VRW.

```
# /opt/VRTSweb/bin/webgui/veaw start
```

Wait for VRTSweb to start and for the VVR application to be loaded.

- 4 Enter the URL using the port number that is displayed in step 2.

Missing buttons on certain wizards

This section describes an error that may occur when using VRW and the solutions to recover from this error.

Problem

Certain buttons may be not be visible on some wizards and dialog windows.

Solution

Change font size for your browser, or use the scroll bars until the buttons are visible.

Log in redirects the user back to the login page

This section describes an error that may occur when using VRW and the solutions to recover from this error.

Problem

Attempting to log in redirects the user back to the login page, although user ID and password are correct. This is due to a limitation with Internet Explorer (IE), if the host name contains any underscores.

Solution

Use one of the following workarounds, listed in order of preference:

- **Create an alternate host name.**
 Edit the `/etc/hosts` file on the UNIX machine to have an alternate host name without any underscores. Also, on the Windows client running IE, edit the `hosts` file to contain the alternate host name. Use this alternate host name in the URL.
- **Use Netscape instead of IE.**
- **In IE, use the IP address in the URL, instead of the host name.**

Index

A

- activating a bunker 92
- Add Secondary wizard 20
- adding a bunker
 - using VVR VEA 25
- Associate Volume wizard 56
- automatic synchronization 27
- AutoSync 27

B

- bandwidth limit
 - changing 61
- Bunker Secondary
 - removing 70
- bunker Secondary
 - activating 92
 - adding
 - using VVR VEA 25
 - deactivating 92

C

- checkpoint attach
 - using 29
- checkpoint option 28, 31
- checkpoints
 - creating 73
 - deleting 74
 - ending 29
 - viewing 74
- Create Primary wizard 17

D

- data volumes
 - dissociating 59
- deactivating a bunker 92
- detailed views 37
- dissociate data volumes from an RDS 59
- Dissociate Volume wizard 59

E

- ending a checkpoint 29
- errors, page display problem 95
- example
 - scenario 16
- examples
 - failing back using difference-based synchronization 89
 - failing back using fast failback 86

F

- failback
 - fast failback 85
 - recommended method 85
 - using Make Secondary wizard 88
- failback options 84
- failing back
 - difference-based synchronization 90
 - fast-failback synchronization 86
 - overview 85
- fast failback 85
- force attach 31

I

- IP addresses, changing 68

L

- latency high mark
 - changing 63
 - viewing 43, 47, 49
- latency low mark
 - changing 63
 - viewing 43, 46, 49
- latency protection
 - changing 61

M

- Make Secondary wizard 88
- Migrate wizard 82

N

No fast-failback option 85

P

packet size

changing 61

page display problems 95

Pause Replication wizard 65

pausing replication 65

Primary

deleting 71

Primary RVG

prerequisites for creating 17

transferring role 81

R

RDS

administering volumes 55

associating volumes 56

creating 16

icons 51

pausing replication 65

removing a bunker 70

removing a Secondary 69

removing data volumes 59

resizing data volumes 58

status information 36

stopping replication 67

transferring RVG roles 81

RDS states

Primary RVGs 51

Remove Bunker wizard 70

Remove Primary wizard 71

Remove Secondary wizard 69

removing a Secondary RVG 69

replication

overview of setting up 15

pausing 65

resuming 66

stopping 67

replication mode

changing 61

replication protocol

changing 61

replication settings

bandwidth limit 61

packet size 61

replication mode 61

replication settings *(continued)*

replication protocol 61

SRL protection 61

Resize Volume wizard 58

resizing data volumes 58

Resume Replication wizard 66

RVG

icons 50

RVG states

Primary RVGs 50

S

Secondary

adding to RDS 20

removing 69

removing bunker 70

Secondary RVG

pausing replication 65

removing a Secondary 69

stopping replication 67

SRL protection

changing 61

Stop Replication wizard 67

Summary view 34

synchronization

difference-based 88

synchronize Secondary

methods 27

synchronizing Secondary

automatic synchronization 27

T

Takeover wizard 84

Taking over from a failed Primary 84

tasks

Deactivate Bunker 92

U

using checkpoint option 28, 31

V

views

Summary view 34

volumes

adding to Secondary RVG 21

administering 55

associating to an RDS 56

removing from an RDS 59

volumes (*continued*)

 resizing 58

VRW tasks

 Activate Bunker 92

VVR VEA wizards

 Add Bunker 25

W

Web wizards

 Add Secondary 20

 Associate Volume 56

 Create Primary 17

 Dissociate Volume 59

 Make Secondary 88

 Migrate 82

 Pause Replication 65

 Remove Bunker 70

 Remove Secondary 69

 Resize Volume 58

 Resume Replication 66

 Stop Replication 67

 Takeover 84