

VERITAS Volume Replicator Web Console 4.1

Administrator's Guide

HP-UX

N13014G

June 2005

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo, and VERITAS Storage Foundation are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000 Fax 650-527-2908
www.veritas.com

Third-Party Legal Notices

Data Encryption Standard (DES) Copyright

Copyright © 1990 Dennis Ferguson. All rights reserved.

Commercial use is permitted only if products that are derived from or include this software are made available for purchase and/or use in Canada. Otherwise, redistribution and use in source and binary forms are permitted.

Copyright 1985, 1986, 1987, 1988, 1990 by the Massachusetts Institute of Technology. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided as is without express or implied warranty.

Apache Software

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work.

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.



8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.



Contents

Preface	xi
Audience	xi
How This Guide Is Organized	xii
Related VERITAS Documents	xii
Conventions	xiii
Getting Help	xiv
Documentation Feedback	xiv
Chapter 1. Introducing VERITAS Volume Replicator Web Console	1
Features of VRW	1
Supported Browsers	2
Installing and Uninstalling VRW	2
Getting Started	3
Starting VRW	3
Accessing and Logging On to VRW	4
About	4
Logging Out of VRW	5
Using Help	5
Selecting Hosts	5
VRW Views	6
Summary View	6
Detailed Views	7
RDS View	8
Primary and Secondary RVG Views	9



Navigating VRW	11
VRW Conventions	11
Using Links in the VRW Views	11
Using the VRW Menu Bar	11
Menus and Menu Options	12
Menus Available from the Summary View	12
Menus Available from the RDS and RVG Views	12
Menu Options for Views	12
Chapter 2. Setting Up Replication	15
Overview of Setting Up Replication	15
Creating a Replicated Data Set	17
Creating a Primary RVG of the RDS	17
Adding a Secondary to the RDS	20
Synchronizing the Secondary and Starting Replication	25
Setting Up Replication Using Automatic Synchronization	26
Setting Up Replication Using Full Synchronization	26
Setting Up Replication Using Block-level Backup and Checkpoint	28
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
Viewing a Summary of the RDSs	33
Using the Summary View	34
Selecting Hosts	34
Information Displayed in the Summary View	35
Conventions for the RDS Icons	36
Conventions for the RVG Icons	36
RVG Status	37
Replication Status	39
Viewing RDS Information	41



Using the RDS View	41
Information Displayed in the RDS View	41
Viewing RVG Information	43
Using the RVG Views	44
Primary RVG View	45
Secondary RVG View	45
Information Displayed in the RVG Views	46
Viewing RLINK Information	49
Viewing Detailed RLINK Information	50
Chapter 4. Administering VERITAS Volume Replicator	53
Administering Users	53
User Roles	53
User Profiles for Central Managing of Hosts	54
Configuring VRW Host Access	54
Configuring Users	55
Configuring User Profiles	55
Configuring Hosts	56
Editing Users	57
Editing Profiles	60
Administering Data Volumes	63
Associating Volumes to a Replicated Data Set	63
Verifying the Data on the Primary and Secondary Volumes	64
Synchronizing Volumes on the Local Host and Remote Hosts	65
Resizing a Data Volume in a Replicated Data Set	66
Dissociating a Data Volume from its Replicated Data Set	67
Administering Replication	69
Changing the Replication Settings	70
Starting Replication to a Secondary	71
Pausing Replication to a Secondary	72



Resuming Replication to a Secondary	72
Stopping Replication to a Secondary	73
Changing the IP Addresses Used for Replication	74
Administering the Replicated Data Set	75
Removing a Secondary from a Replicated Data Set	75
Removing a Primary RVG	76
Administering Checkpoints	77
Creating Checkpoints	77
Ending Checkpoints	77
Viewing Checkpoints	78
Deleting Checkpoints	78
Creating RVG Snapshots	78
The Instant Snapshot Feature	79
Instant Full Snapshots	79
Instant Space-Optimized Snapshots	82
Instant Plex-Breakoff Snapshots	83
Traditional Snapshots	85
Displaying Snapshots	86
Administering the SRL	87
Protecting from SRL Overflow	87
Incrementally Synchronizing the Secondary After SRL Overflow	87
Resizing the Primary SRL	89
Performing Data Verification	90
Verifying RVG Offline	90
Verifying RVG Online	92
Using the Command Window	93
Chapter 5. Transferring the Primary Role	95
Migrating the Primary Role	95
Example—Migrating From a Healthy Primary	97



Taking Over From an Original Primary	99
Failback Options	99
Failing Back to the Original Primary	102
Failing Back Using Fast-Failback Synchronization	102
Example—Failing Back to the Original Primary Using Fast Failback	102
Failing Back Using Difference-Based Synchronization	104
Converting a Primary to a Secondary	105
Synchronize the Secondary RVG with the Primary RVG	106
Example: Failing Back Using Difference-Based Synchronization	107
Chapter 6. Additional Settings for VRW	109
Configuring the VRW Application	109
Configuring VRW Using the Configuration File	110
Editing the Configuration File	110
Starting and Stopping the VRW Application	111
Appendix A. Administering VERITAS Web Server	113
Getting Started 114	
Reviewing the Web Server Configuration	115
Configuring Ports for VRTSweb	116
Retrieving the List of Ports	116
Adding Ports	117
Deleting Ports	119
Changing the Administrative Port	120
Managing VRTSweb SSL Certificates	121
Viewing SSL Certificate Information	121
Creating a Self-Signed SSL Certificate	121
Exporting SSL Certificate to a File	122
Configuring a CA-Signed SSL Certificate	122
Cloning the VRTSweb SSL Certificate	123



Configuring SMTP Notification for VRTSweb	124
Retrieving the Name of the Configured SMTP Server	124
Setting the SMTP Server	125
Retrieving Configured SMTP Recipients	126
Adding an SMTP Recipient	127
Deleting an SMTP Recipient	129
Configuring VRTSweb Logging	130
Retrieving Log Levels	130
Modifying Log Levels	131
Modifying Maximum Size Limit and Rollover Count for Logs	132
Modifying the Maximum Heap Size for VRTSweb	134
Appendix B. Troubleshooting VRW	135
Problem After Successfully Logging on to VVR	135
Application Page is Not Displayed in the Browser	136
Missing Close Button on Certain Wizards	137
Correct RDS Name is not Displayed	137
Log in redirects the user back to the login page	138
Appendix C. Accessibility and VRW	139
Navigation and Keyboard Shortcuts	139
Navigation in the Web Console	139
Support for Accessibility Settings	140
Support for Assistive Technologies	140
Index	141



Preface

VERITAS™ Volume Replicator Web Console (VRW) is the Web-based Graphical User Interface for VVR. The *VERITAS Volume Replicator Web Console Administrator's 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.

Audience

This guide is for system administrators who are responsible for configuring and maintaining VVR. It assumes that the user has:

- ◆ A basic understanding of system administration.
- ◆ A working knowledge of the UNIX operating system.
- ◆ A basic working knowledge of VERITAS Volume Manager (VxVM) and VERITAS VERITAS Volume Replicator (VVR).



How This Guide Is Organized

[Chapter 1. “Introducing VERITAS Volume Replicator Web Console” on page 1](#) gives an overview of VRW, explains the features of VRW, and helps you get started with VRW.

[Chapter 2. “Setting Up Replication” on page 15](#) describes the procedure to set up an RDS, synchronizing, and starting replication to a Secondary.

[Chapter 3. “Viewing Configuration and Status Information” on page 33](#) provides information on using VRW to find consolidated information about an RDS and its replication status.

[Chapter 4. “Administering VERITAS Volume Replicator” on page 53](#) describes the procedure of using VRW’s wizards to administer VERITAS Volume Replicator (VVR) components from any host in a Replicated Data Set (RDS).

[Chapter 5. “Transferring the Primary Role” on page 95](#) describes the procedure to transfer and fail back the Primary role using specific VRW features.

[Chapter 6. “Additional Settings for VRW” on page 109](#) describes the advanced VVR configuration tasks that can be performed using VRW.

[Appendix A. “Administering VERITAS Web Server” on page 113](#) describes the configuration settings for VERITAS Web Server.

[Appendix B. “Troubleshooting VRW” on page 135](#) describes the errors that may occur when using VRW and provides solutions to recover from these errors.

[Appendix C. “Accessibility and VRW” on page 139](#) provides information about the accessibility features and compliance of VRW.

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. Refer to the *VERITAS Volume Replicator Release Notes* for more information on these documentation sets.



Conventions

Convention	Usage	Example
monospace	Used for path names, commands, output, directory and file names, functions, and parameters.	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
monospace (bold)	Indicates user input.	# ls pubs C:\> dir pubs
<i>italic</i>	Identifies book titles, new terms, emphasized text, and variables replaced with a name or value.	See the <i>User's Guide</i> for details. The variable <i>system_name</i> indicates the system on which to enter the command.
bold	Depicts GUI objects, such as fields, list boxes, menu selections, etc. Also depicts GUI commands.	Enter your password in the Password field. Press Return .
blue text	Indicates hypertext links.	See " Getting Help " on page xiv.
#	Unix superuser prompt (all shells).	# cp /pubs/4.0/user_book /release_mgnt/4.0/archive
C:\>	Windows user prompt.	C:\> copy \pubs\4.0\user_book c:\release_mgnt\4.0\archive



Getting Help

For technical assistance, visit <http://support.veritas.com> and select phone or email support. This site also provides access to resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service. Use the Knowledge Base Search feature to access additional product information, including current and past releases of VERITAS documentation.

For license information, software updates and sales contacts, visit <https://my.veritas.com/productcenter/ContactVeritas.jsp>. For information on purchasing product documentation, visit <http://webstore.veritas.com>.

Documentation Feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to vvrdocs@veritas.com. Include the title and part number of the document (located in the lower left corner of the title page), and chapter and section titles of the text on which you are reporting. Our goal is to ensure customer satisfaction by providing effective, quality documentation. For assistance with topics other than documentation, visit <http://support.veritas.com>.



Introducing VERITAS Volume Replicator Web Console

1

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.

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.

- ◆ Show Commands

VRW is a Web-based interface for executing VVR or VxVM commands, as an alternative to using the command line. When you complete a wizard or dialog box and then click **OK**, VRW executes the command. Clicking on **Show Commands** displays the command or commands corresponding to your selections in the wizard or dialog box.

The Command Window, Edit User List and Edit User Profile operations do not use VVR or VxVM commands; hence they do not have the Show Commands button.

- ◆ Managing User Roles

VRW provides the ability to define user roles. This enables you to create non-root users and to set read-only access for some users.



- ◆ Centralized Management of configuration on multiple hosts

VRW provides a centralized console to administer RDSs on multiple hosts from one host, through user profiles. You can customize a list of hosts for which you want to administer RDSs. VRW must be installed on each host in the list.

Supported Browsers

VRW supports the following browsers:

Platform	Supported Browser
Solaris 8, Solaris 9, Solaris 10	Netscape 6.2 or later
Windows NT 4.0, Windows 2000	Internet Explorer 5.0 or later and Netscape 6.2 or later
AIX 5.1 and AIX 5.2	Netscape 6.2 or later
September 2004 HP-UX 11i version 2.0	Netscape 6.2 or later
SUSE Linux Enterprise Server 9 (SLES 9) SP1 and Red Hat Enterprise Linux 4 (RHEL 4) Update 1	Mozilla 1.4 and Netscape 6.2 or later

Installing and Uninstalling VRW

For instructions on installing and uninstalling VRW, see the *VERITAS Volume Replicator Installation Guide*.



Getting Started

This section provides information about starting and accessing VRW.

Starting VRW

After VRW is installed successfully on all the VVR hosts, start VRW on all the hosts.

- ❖ Start VRW using the following command:

```
# /opt/VRTSvrw/bin/vrw start
```

You can now use the browser to manage VVR.



Accessing and Logging On to VRW

Perform the following steps to access VERITAS Volume Replicator Web Console:

1. Launch the Web browser that is supported for your system.
2. In the Location or address edit box of the browser, enter the URL of the server.

`https://hostname:8443/vvr`

where *hostname* is the name or the IP address of the host on which VRW is installed and 8443 is the default Web server port for VRW.

For example, if VRW is installed on the server `seattle`, enter the following URL:

`https://seattle:8443/vvr`

Note The client must have an `http` connection to the host (for example, `seattle`).

3. Accept the self-signed certificate (issued by VERITAS) to proceed.

You can prevent this certificate from appearing every time you connect to the console by installing a CA-signed certificate. See “[Configuring a CA-Signed SSL Certificate](#)” on page 122 for instructions.

The browser displays the VRW login page.

4. Enter the user name and password. By default, `root` is the only user with VRW access after installation. For information about configuring additional users, see “[Administering Users](#)” on page 53.

The VRW Summary view is displayed. For more information, see “[Summary View](#)” on page 6.

About

The **About** item provides information about VVR, including copyright details, the version of VRW, and the version of VVR that was shipped with VRW.

▼ To view information about VVR

1. Click **About** in the top right corner of any view.
2. Click **OK** to close the About window.



Logging Out of VRW

▼ To log out of VRW

Click **Logout** in the top right corner of any view.

Using Help

VRW provides online help. Use the Contents, Index, and Search features in the help system to locate the information you need.

▼ To access the online help

Click **Help** in the top right corner of any view, or

Click on the **Help** button on a wizard page to get help on the wizard.

Selecting Hosts

The upper right corner of each view in VRW displays two drop-down lists: Profiles and Hosts. The Profiles list displays the *profiles* defined for that user. Profiles are used to define a customized list of hosts for which that user can manage and view RDS and RVG data.

When you select a profile, the host list is populated with the hosts defined in that profile and the keyword **All**. When **All** is selected (the default), VRW displays RDS information for all the hosts defined in the current profile. Select a host from the drop-down host list to filter the view, showing only RDSs on that host. Selecting a host also redirects the connection to that host. For more information about profiles and the host list, see “[User Profiles for Central Managing of Hosts](#)” on page 54.

Note If any of the RVGs in an RDS are on a host in the profile, VRW displays information about all RVGs in that RDS. Although some of the RVGs may not reside on any host in the host list, they can be viewed and modified as needed.



VRW Views

VRW provides different views that provide information about the VVR objects and their relationships. The Summary view is displayed by default when you log on to VRW. If you log on after a session time-out using the same browser session, VRW displays the last selected view. You can navigate to different views from the Summary view. VRW provides the following views:

- ◆ Summary view—Displays the list of Replicated Data Sets (RDSs) present on the specified hosts, including any remote RVGs belonging to these RDSs.
- ◆ RDS view—Displays details of an RDS and its status.
- ◆ Primary RVG view—Displays details of the Primary RVG and its status.
- ◆ Secondary RVG view—Displays details of the Secondary RVG and its status.

Summary View

When you log on to VRW, the Summary view displays by default.

The screenshot shows the VERITAS Volume Replicator Web Console interface. The title bar indicates the user is logged in as 'root' on the 'seattle' host. The main content area displays the 'Summary of RDSs' view, which includes a table with the following data:

RDS	RVG	RVG Status	Replication Status
hr_rvg	Primary: seattle Secondary: london	enabled for I/O consistent, up-to-date	replicating (connected)

Below the table, there is an 'Alerts' section showing 'No Alerts'.

The Summary view displays information about the RVGs present on the host list, such as the replication status, RVG state, and the names of the Primary and Secondary hosts. By default, the Summary view shows information for all the hosts in the host list. If a specific host is selected, the list shows RDSs present on that host and any remote RVGs belonging to these RDSs.

The Alert section in the lower region of the Summary view displays configuration problems in an RDS, if any. The RDS name along with a brief description about the mis-configuration is displayed.

Detailed Views

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

- ✓ The RDS view—To display the RDS view, click the RDS name.
- ✓ Primary RVG view—To display the Primary RVG view, click the Primary host name of the RDS to which the Primary RVG belongs.
- ✓ Secondary RVG view—To display the Secondary RVG view, click the Secondary host name of the RDS to which the Secondary RVG belongs.

Each view has a similar layout.



RDS View

- ✓ To display the RDS view, click the RDS name.

The screenshot displays the VERITAS Volume Replicator Web Console interface. The left pane shows a tree view with 'hr_rvg' selected, containing sub-items for 'seattle' and 'london'. The right pane shows the configuration for the 'hr_rvg' RDS. The 'Primary' section shows the 'seattle' host is 'enabled for I/O' with a 'hrdg' disk group and 0 checkpoints. The 'Replication Log' shows a log named 'hr_srl' with a size of 4.00 GB and a layout of 'CONCAT'. The 'Replication Status' section shows the 'london' host is in a 'consistent, up-to-date replicating (connected)' state with an 'SRL' progress indicator at 0% full, using an 'asynchronous' mode. The 'Data Volumes' section lists two volumes: 'hr_dv01' and 'hr_dv02', both 4.00 GB in size with a 'CONCAT' layout and 0 mirrors. The 'Alerts' section shows 'No Alerts'.

Primary	State	Disk Group	Checkpoints
seattle	enabled for I/O	hrdg	0

Secondary	Data Status	Replication Status	Logging To	Current Mode	Checkpoints
london	consistent, up-to-date	replicating (connected)	SRL ██████████ 0% full, 0	asynchronous	0

Name	Size	Layout	Mirrors	DCM
hr_dv01	4.00 GB	CONCAT	0	Yes
hr_dv02	4.00 GB	CONCAT	0	Yes
Total Size : 8.00 GB				

The left pane displays the RDSs (and their Primary and Secondary RVGs) that are found on the host list or the selected host. The right pane displays configuration information about the Primary RVG, the replication status of each of the Secondary RVGs in the RDS, and the details of the data volumes of the Primary RVG.

If there are any problems with the configuration of the RDS, the Alerts section at the bottom of the pane shows the misconfiguration messages.

Primary and Secondary RVG Views

- ✓ To display the Primary RVG view, click the Primary host name of the RDS to which the Primary RVG belongs.

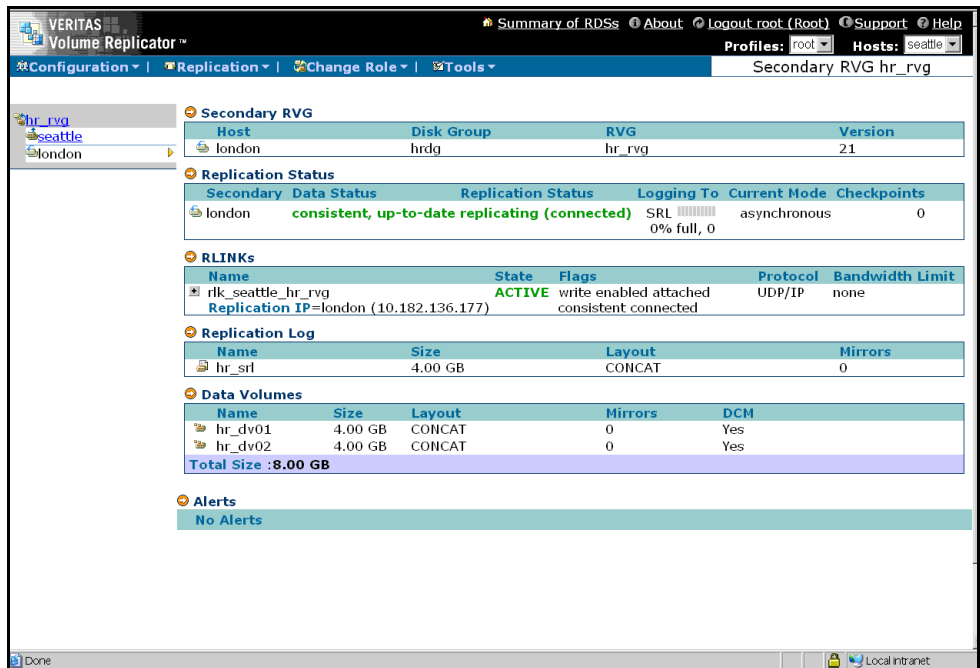
The screenshot shows the VERITAS Volume Replicator web console interface. The main content area displays the configuration for the Primary RVG 'hr_rvg' on host 'seattle'. The interface includes a navigation menu on the left with 'hr_rvg', 'seattle', and 'london'. The main panel is titled 'Primary RVG hr_rvg' and contains several sections:

- Primary RVG:** A table showing the host 'seattle' with disk group 'hrdg' and RVG 'hr_rvg' (Version 21). The state is 'enabled for I/O'.
- Replication Log:** A table showing a replication log entry 'hr_srl' with a size of 4.00 GB and layout 'CONCAT'. There are 0 mirrors and no checkpoints.
- RLINKS:** A table showing an active RLINK 'rlk_london_hr_rvg' with flags 'write enabled attached consistent connected asynchronous', protocol 'UDP/IP', and no bandwidth limit. The replication IP is 'seattle (10.182.136.176)'.
- Data Volumes:** A table showing two data volumes, 'hr_dv01' and 'hr_dv02', both with a size of 4.00 GB and layout 'CONCAT'. There are 0 mirrors and DCM is enabled for both. The total size is 8.00 GB.
- Alerts:** A section indicating 'No Alerts'.

The bottom of the console shows a 'Done' button and a 'Local intranet' icon.



- ✓ To display the Secondary RVG view, click the Secondary host name of the RDS to which the Secondary RVG belongs.



The left pane displays the RDSs (and their Primary and Secondary RVGs) that are found on the host list or the selected host. The right pane displays detailed information about the attributes, and the status of the selected object.

Each view contains a different set of menus specific to the view. For example, the Menu Bar in the RDS view contains the menus Configuration, Replication, Change Role, and Tools, whereas the Menu Bar in the Summary view contains the menus Configuration and Tools.



Navigating VRW

You can navigate VRW using the navigation links. When you log on to VRW, the Summary view displays. You can use the navigation links to navigate to other views.

VRW Conventions

- ◆ Object names or text that appear with underlines indicate that the object name or text is a link to other views related to the selected object name or text.

Using Links in the VRW Views

The following navigation links are available in the top right corner:

- ◆ **About** displays version and copyright information about VVR.
- ◆ **Summary of RDSs** takes you to the Summary view.
- ◆ **Logout** enables you to log out of VRW. The role of the current session's username is displayed in parenthesis.
- ◆ **Support** takes you to the VERITAS Customer Support web page for VVR.
- ◆ **Help** launches the online help for the selected view in a separate browser window.

Using the VRW Menu Bar

If you log on as the root user or an administrative user, the menu bar appears at the top of each VRW view. The menu bar consists of drop-down menus with options that enable you to perform administrative tasks on the selected object. The navigation links **About**, **Summary of RDSs**, **Logout**, **Support** and **Help** are available at the top right corner. The current session's username and role are displayed in parenthesis next to the **Logout** link.

- ◆ To display a menu, move the mouse over a menu name on the Menu bar. Do not click the menu name. Each menu contains a list of options.
- ◆ To choose a menu option, click the required option. After you choose an option, a wizard window or a task dialog box is displayed. For example, if you choose **Pause** from the Replication menu, the Pause Replication dialog box is displayed.

A guest user has view-only privileges, and the menu bar does not contain any menus.



Menus and Menu Options

VRW provides menu options to perform administrative tasks. The VRW administrative tasks are classified into four main categories: Configuration, Replication, Change Role and Tools. Each of these categories is available as a menu on the menu bar. Each menu includes a set of options. Use these menu options to perform an administrative task, using a wizard or dialog box.

Each VVR object view includes the appropriate menus and menu options for that view, so the menus available on the menu bar vary depending on the views.

Menus Available from the Summary View

The Summary view contains the following menus:

- ◆ Configuration
- ◆ Tools

Each menu provides options, each of which represent a corresponding task. For information about the menu options available for the Summary view, see [“Menu Options for Views”](#) on page 12.

Menus Available from the RDS and RVG Views

The RDS and RVG views contain the following menus:

- ◆ Configuration
- ◆ Replication
- ◆ Change Role
- ◆ Tools

Each menu option represents a corresponding task for a specific view as indicated in the following table. Select the appropriate option to initiate tasks for the selected category. If an option is not available in a particular context, the menu option is dimmed to indicate it is not selectable.

Menu Options for Views

The following table shows the menus and menu operations available from each view.

Menu Options for Views

View	Menu	Operations
Summary	Configuration	Create Primary
	Tools	Synchronize Volume, Verify Volume, Command, Edit Profile List, and Edit User List
RDS, Primary RVG and Secondary RVG	Configuration	Create Primary, Add Secondary, Remove Secondary, Remove Primary, Associate Volume, Dissociate Volume, Resize Volume, Resize SRL, Change Replication Settings, and Change IP
	Replication	Start, Stop, Pause, Resume
	Change Role	Take Over, Make Secondary, Migrate, and Replay Failback Log
	Tools	Start Checkpoint, End Checkpoint, Delete Checkpoint, Synchronize Volume, Verify Volume, Synchronize RVG, Verify RVG Offline, Verify RVG Online, Create Snapshot, Display Snapshot, Replay DCM, Command, Edit Profile List, and Edit User List

For a detailed explanation about these tasks, refer to the appropriate section in this document.





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 using the Configuration Worksheet in the *VERITAS Volume Replicator Administrator's Guide*. Follow the best practices or recommendations listed in the *VERITAS Volume Replicator Administrator's Guide* to ensure successful configuration of VVR.

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	<code>hrdg</code>
Primary RVG	<code>hr_rvg</code>
Primary RLINK to Secondary <code>london</code>	<code>rlk_london_hr_rvg</code>
Primary data volume #1	<code>hr_dv01</code>
Primary data volume #2	<code>hr_dv02</code>
Primary SRL volume	<code>hr_srl</code>

Secondary Host name: `london`

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



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](#)

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; that is, the host to which you logged on or which you selected from the drop-down host list.
- ◆ 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.

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

Prerequisites for creating a Primary RVG

- ✓ 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.

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

▼ To create a Primary RVG of an RDS

1. From any view, choose **Configuration > Create 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**.



3. Complete the RVG Name and Disk Group page as follows:

RVG (RDS) Name	Enter a name for the Primary RVG. For example, <code>hr_rvg</code> .
Disk Group	<p>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, <code>hr_dg</code>.</p> <p>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.</p>

4. Click **Next**.

5. Complete the SRL and Data Volume page as follows:

Replication Log	<p>From the drop-down list, select the volume that is to be configured as the SRL. For example, <code>hr_srl</code>.</p> <p>The SRL size must be at least 110MB.</p>
Data Volumes	<p>From the list, select the data volumes to be associated with the RVG. All of the data volumes used by the application MUST be included in the same RVG.</p> <p>To select multiple data volumes from the list, hold down the CTRL key and click the names of the required data volumes.</p> <p>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, <code>hr_dv01</code> and <code>hr_dv02</code>.</p> <p>Note If the checkbox The volumes are VxVM ISP volumes was selected on the first page, the list only displays VxVM ISP data volumes.</p>



6. Click **OK** 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 **Close**. The Primary RVG view displays the newly created Primary RVG.



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. The **Add Secondary** wizard also enables you to set up or change the replication settings between a Primary RVG and the Secondary that is being added. You can change the following replication settings for the Secondary:

- ◆ Mode of replication
- ◆ Latency protection
- ◆ SRL protection

Each setting affects performance and must be set up with care. For details on setting up mode of replication, latency protection and SRL protection, see the *VERITAS Volume Replicator Administrator's Guide*.

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 to 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 second step of the **Add Secondary** wizard.
- ◆ Automatically 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 `r1k_remotehost_rvgname`. The wizard enables you to specify different names for the Primary and Secondary RLINKs.

Prerequisites for adding a Secondary

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. From the Summary view, the RDS view, or the RVG view, click the name of the RDS to which you want to add the Secondary RVG, and then choose **Configuration > Add Secondary**.
2. Complete the Primary and Secondary Host Names page as follows, and then click **Next**:

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, <code>seattle</code> . If you entered the IP address, it must be reachable from the Secondary host.
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, <code>london</code> . If you entered the IP address, it must be reachable from the Primary host.



3. In the RLINKs and Secondary Disk Group Names page, the default values for the Primary RLINK, Secondary RLINK, and the Secondary Disk Group display.

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 as follows, or go to the next step:

Primary RLINK	Enter another Primary RLINK name if the default name is not appropriate.
Secondary RLINK	Enter another Secondary RLINK name if the default name is not appropriate.
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.

4. Click **Next**.
5. In the Replication Settings page, the default values for Replication Mode, SRL Protection, and Latency Protection display.

For information about replication settings, refer to the *VERITAS Volume Replicator Administrator's Guide*.

Note After the Secondary has been successfully added, additional replication settings such as packet size, protocol, and bandwidth limit can be changed as described in [“Changing the Replication Settings”](#) on page 70.

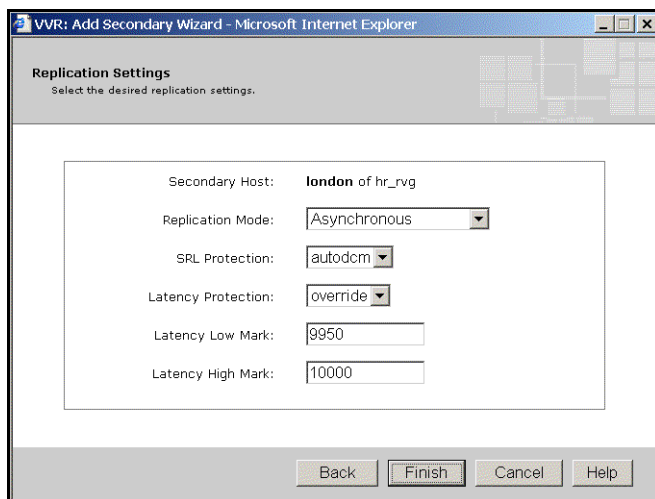


To change the replication settings, complete the Replication Settings page as follows, or go to the next step.

Replication Mode Asynchronous, Synchronous-Override	Asynchronous is selected by default. From the drop-down list, select the required mode of replication. To set the replication to synchronous mode, select Synchronous-Override .
SRL Protection autodcm, dcm, off, override	By default, SRL protection is set to autodcm . 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.
Latency Protection override, off	By default, latency protection is set to override . To enable latency protection and limit the number of updates by which the Secondary is allowed to fall behind, select override . If you select override , the wizard displays the boxes Latency Low Mark and Latency High Mark with default values of 9950 and 10000 respectively. To disable latency protection, select off . To set the latency protection to fail , use the command-line interface (CLI). For details on selecting the appropriate option and using the CLI, refer to the <i>VERITAS Volume Replicator Administrator's Guide</i> .
If Latency Protection is set to override :	
Latency High Mark	Displays only if Latency Protection is set to override . The default is 10000. Enter the required numerical limit.
Latency Low Mark	Displays only if Latency Protection is set to override . The default is 9950. Enter the required numerical limit.



If you select **override** in the Latency Protection box, the page looks like this:



6. Click **Finish**. The Summary page displays the replication settings for the Secondary RVG.
7. Click **OK** to create 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. The following message displays in a window:

The Secondary has been added.

8. Click **Close**. The Secondary RVG view for the newly added Secondary RVG is displayed.

Caution If the data volumes on the Primary contain valid data, synchronize the Secondary data volumes with the Primary data volumes. See [“Synchronizing the Secondary and Starting Replication”](#) on page 25. Otherwise, see [“Setting Up Replication When Data Volumes are Initialized With Zeroes”](#) on page 31.

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 26.

Note For component volumes of a volume set, do not use **Tools > Synchronize RVG**, use **Replication > Start** with the **Using Autosync** option instead.

- ◆ Full synchronization - See [“Setting Up Replication Using Full Synchronization”](#) on page 26.
- ◆ Block-level tape backup - See [“Setting Up Replication Using Block-level Backup and Checkpoint”](#) on page 28.
- ◆ 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, refer to the *VERITAS Volume Replicator Administrator's Guide*.



Setting Up Replication Using Automatic Synchronization

The instructions in this section assume that the RDS has been created using the procedure described in “[Creating a Replicated Data Set](#)” on page 17. You can synchronize the Secondary using automatic synchronization when the application is active or inactive.

Using the **autosync** option enables you to automatically synchronize the Secondary data volumes with the Primary data volumes in an RDS and start replication; **AutoSync** makes the Secondary data volumes up to date with the Primary data volumes. Use this option to synchronize the Secondary when the data volumes contain data and when the application is active or inactive.

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

▼ To set up replication using automatic synchronization

1. From the Secondary RVG view, choose **Replication > Start**.
2. In the Replication Options page, select **Using Autosync** to synchronize the Secondary and start replication using automatic synchronization.
3. Click **OK**. A message displays the status of your request to start replication.
4. Click **Close**. 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

The instructions in this section assume that the RDS has been created using the procedure described in “[Creating a Replicated Data Set](#)” on page 17. You can synchronize the Secondary using full synchronization with checkpoint when the application is active or inactive.

▼ **To set up replication using full synchronization**

1. From the RDS view or Primary RVG view, choose **Tools > Synchronize RVG**.
2. Complete the Synchronize RVG dialog box as follows:

Options: Difference-based synchronization Full synchronization	Select Full synchronization .
Checkpoint Name	Enter a checkpoint name of your choice. For example, ckpt_presync . The Synchronize RVG task uses this name to create a checkpoint.
Secondary Host	From the drop-down list, select one or more Secondary hosts to be synchronized. For example, london .
Display output in reverse order	Select this option to display output in reverse order.

3. Click **OK**. The status of the synchronization displays. Wait until the synchronization completes before you go to the next step.
4. Click **Close** when the synchronization completes.
5. From the Secondary RVG view, choose **Replication > Start**.
6. In the Start Replication dialog box, select **Using checkpoint**. From the drop-down menu, select the checkpoint name which you entered in [step 2](#). This uses the selected checkpoint to start replication.
7. Click **OK**. A message displays the status of your request to start replication.
8. 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 2](#), repeat [step 5](#) through [step 8](#) to start replication for each Secondary RVG.



Setting Up Replication Using Block-level Backup and Checkpoint

The instructions in this section assume that the RDS has been created using the procedure described in “[Creating a Replicated Data Set](#)” on page 17. You can synchronize the Secondary using block-level backup and checkpoint when the application is active or inactive.

▼ To set up replication using block-level backup and checkpoint

1. From the RDS or RVG view, choose **Tools > Start Checkpoint** to start a checkpoint on the Primary.
2. In the Start Checkpoint dialog box, enter a name for the checkpoint and click **OK**. Note down the checkpoint name you entered, for example, **checkpt_presync**.
3. Perform a block-level backup of the data volumes in the Primary RVG.
4. Choose **Tools > End Checkpoint** to end the Primary checkpoint when the backup is complete.
5. To confirm that you want to end the checkpoint, click **OK**.
6. Restore the backup to the Secondary data volumes.
7. When [step 6](#) is completed, check whether the checkpoint you created is still valid by clicking on the checkpoint name in the Primary RVG view. If the checkpoint status shows that it has overflowed, the checkpoint is not valid. Repeat [step 1](#) to [step 6](#).
8. From the Secondary RVG view, choose **Replication > Start**.
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 **OK**. A message displays the result of your request to start replication.
11. Click **Close**. The Secondary RVG 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 Primary RVG view, check whether the `consistent` flag is set on 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, refer to 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, checkpoints that are unusable are displayed in red.



Setting Up Replication Using Difference-Based Synchronization

The instructions in this section assume that the RDS has been created using the procedure described in “[Creating a Replicated Data Set](#)” on page 17. You can synchronize the Secondary using difference-based synchronization with checkpoint when the application is active or inactive.

▼ To set up replication using difference-based synchronization

1. From the RDS view (or RVG view), choose **Tools > Synchronize RVG**.
2. Complete the Synchronize RVG dialog box as follows:

Options: Difference-based synchronization Full synchronization	Select Difference-based synchronization to synchronize the Secondary RVG <code>hr_rvg</code> on <code>london</code> with its Primary RVG on <code>seattle</code> using difference-based synchronization with checkpoint.
Checkpoint Name	Enter a checkpoint name of your choice. For example, checkpt_presync .
Secondary Host	From the drop-down list, select one or more Secondary hosts to be synchronized. For example, london .

3. Click **OK**. The status of the synchronization displays.
4. Start the application only after you have performed this step.
5. From the Secondary RVG view, choose **Replication > Start**.
6. 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.
7. Click **OK**. A message displays the result of your request to start replication.
8. 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 2](#), repeat [step 5](#) through [step 8](#) 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 using the procedure described in “[Creating a Replicated Data Set](#)” on page 17.

Use the **Force Attach** 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 4](#).

▼ To set up replication when data volumes are initialized with zeroes

1. From the Secondary RVG view, choose **Replication > Start**.
2. In the Start Replication dialog box, select **Force Attach**. The Secondary host is selected in the **Secondary Host** box. For example, **london**.
3. Click **OK**. A message displays the result of your request to start replication.
4. Click **Close**. The RDS view for `hr_rvg` on `seattle` is automatically refreshed to display the replication status at a particular point in time.

After completing this step, start the application.





Viewing Configuration and Status Information

3

VERITAS Volume Replicator Web Console (VRW) enables you to get information about the VVR environment using the VRW views. For example, the RDS 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.

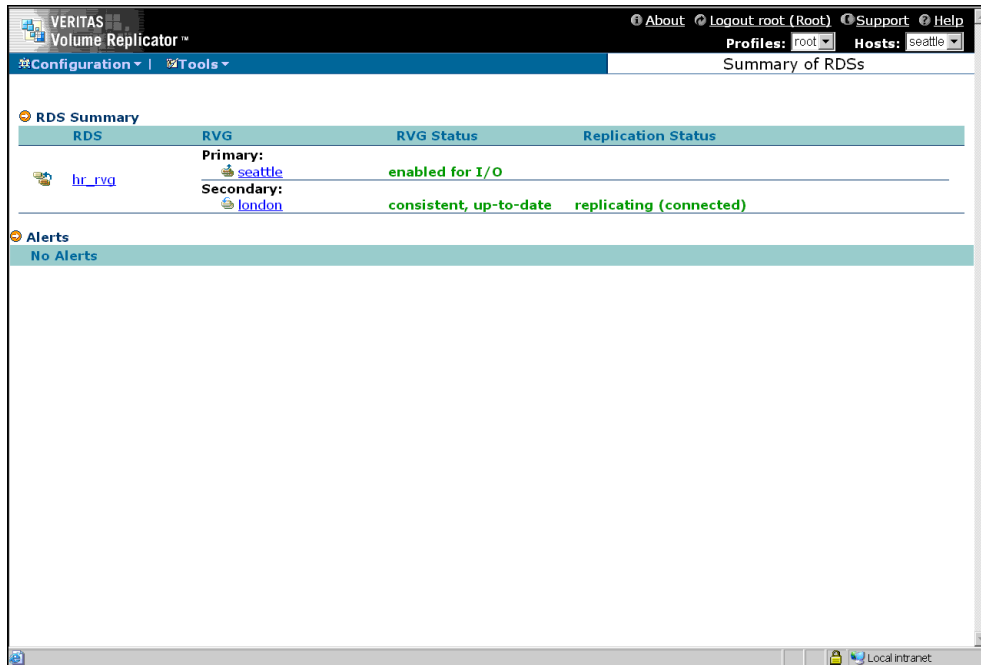
Viewing a Summary of the RDSs

The Summary view appears when you log on to VRW. The Summary view displays information about the RDSs present on the selected hosts, such as the replication status, RVG status, and the names of the Primary and Secondary hosts. The Alerts section in the lower region of the Summary view displays configuration errors in the RDSs, if any. Each line in the Alerts section displays the RDS name along with a brief description of the problem with the configuration.



Using the Summary View

To display the Summary view while you are in another VRW view, click **Summary of RDSs** on the top right corner. The Summary view looks like this:



Selecting Hosts

The upper right corner of each view in VRW displays two drop-down lists: Profiles and Hosts. The Profiles list displays the *profiles* defined for that user. Profiles are used to define a customized list of hosts for which that user can manage and view RDS and RVG data.

When you select a profile, the host list is populated with the hosts defined in that profile and the keyword **All**. When **All** is selected (the default), VRW displays RDS information for all the hosts defined in the current profile. Select a host from the drop-down host list to filter the view, showing only RDSs on that host. Selecting a host also redirects the connection to that host. For more information about profiles and the host list, see [“User Profiles for Central Managing of Hosts”](#) on page 54.

Note If any of the RVGs in an RDS are on a host in the profile, VRW displays information about all RVGs in that RDS. Although some of the RVGs may not reside any host in the host list, they can be viewed and modified as needed.

Information Displayed in the Summary View

The upper region of the view shows the following:

- ◆ The links **About**, **Logout**, **Support**, and **Help**, as described in “[Using Links in the VRW Views](#)” on page 11.
- ◆ Profiles and Hosts drop-down lists, for managing multiple hosts, as described in “[Selecting Hosts](#)” on page 34.
- ◆ The type of the view, Summary of RDSs, displays in the top-right corner, below the navigation links.
- ◆ The menu bar in the Summary view provides options to start the VVR task wizards, which enable you to configure and administer the RDS. For details, see “[Menus Available from the Summary View](#)” on page 12.

The Summary view displays the following information:

RDS	The RDS icon and the name of the RDS. The RDS icon indicates the state of the RDS configuration. For the descriptions of the RDS configuration state indicator, see “ RDS Configuration State Icons ” on page 36. The RDS name is a navigation link to the RDS view.
RVG	
Primary:	Displays Primary RVG icon and 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. For information about the Primary RVG indicators, see “ Conventions for Primary I/O States or Secondary RVG Configuration States ” on page 36.
Secondary:	Displays the Secondary RVG icon and the name of the Secondary host. For more information about the Secondary RVG icons, see “ Conventions for Primary I/O States or Secondary RVG Configuration States ” on page 36.
RVG Status	
Primary:	Shows the state of the RVG. 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. See “ Primary RVG States ” on page 37.
Secondary:	Shows the Data Status of the Secondary RVG. See “ Secondary RVG Data Status ” on page 38.





Replication Status	<p>The replication status to the Secondary in text format. The value of the replication status field. Paused by user; paused due to network disconnection, etc.</p> <p>The replication status n/a means that the RDS is configured incorrectly and the Primary is not known. A text message also displays if a Secondary does not exist.</p> <p>See “Replication Status” on page 39.</p>
---------------------------	--

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.



RDS Configuration State Icons



Icon	Description
	Displays when the RDS does not have any configuration errors.
	Displays when the RDS has some configuration errors. Refer to the <i>VERITAS Volume Replicator Administrator's Guide</i> for more details of the list of 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.

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.

Icon	Description
	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.

RVG Status

The RVG status field shows the current state of the Primary RVG and the data status for the Secondary RVG.

Primary RVG States

The following table shows the different values of the Primary RVG states and their meanings.

Primary RVG States

Value	Meaning
<code>acting_secondary</code>	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.
<code>disabled for I/O</code>	Primary RVG is disabled for I/O, that is, the RVG is stopped.
<code>enabled for I/O</code>	Primary RVG is enabled for I/O, that is, RVG is started.
<code>needs recovery</code>	State after an import or reboot. The <code>vrxvg recover rvg</code> command clears this state.
<code>passthru</code>	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:

Secondary RVG Data Status

Value	Meaning
<code>consistent, behind</code>	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.
<code>consistent, stale</code>	The data on this Secondary is consistent. Replication to this Secondary has been stopped; the Primary RLINK is detached.
<code>consistent, up-to-date</code>	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.
<code>inconsistent</code>	The data on the Secondary volumes is not consistent and the Secondary cannot take over.
<code>needs recovery</code>	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.

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 87. needs failback synchronization—To continue replication, start failback synchronization to this Secondary. See “Failing Back Using Fast-Failback Synchronization” on page 102. For more details, 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.



Value	Meaning
replicating	connected—Replication can take place if there are updates on the Primary data volumes.
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.



Viewing RDS Information

Use the RDS view to get detailed information about a selected Replicated Data Set (RDS).

Using the RDS View

To display the RDS view from any VRW view, click on the RDS name. For example, click the RDS name `hr_rvg`, to display the RDS view for `hr_rvg`. The RDS view looks like this:

The screenshot displays the RDS view for `hr_rvg` in the VERITAS Volume Replicator interface. The window title is "Summary of RDSs". The navigation bar includes "Configuration", "Replication", "Change Role", and "Tools". The main content area shows details for the primary host "seattle" and secondary host "london".

Primary	State	Disk Group	Checkpoints
seattle	enabled for I/O	hrdg	0

Replication Log: Name: hr_srl Size: 4.00 GB Layout: CONCAT Mirrors: 0

Secondary	Data Status	Replication Status	Logging To	Current Mode	Checkpoints
london	consistent, up-to-date replicating (connected)	SRL ██████████	0% full, 0	asynchronous	0

Name	Size	Layout	Mirrors	DCM
hr_dv01	4.00 GB	CONCAT	0	Yes
hr_dv02	4.00 GB	CONCAT	0	Yes

Total Size : 8.00 GB

Alerts
No Alerts

Information Displayed in the RDS View

The upper region of the view shows the following:

- ◆ The links **Summary of RDSs**, **About**, **Logout**, **Support**, and **Help**, as described in "Using Links in the VRW Views" on page 11.
- ◆ Profiles and Hosts drop-down lists, for managing multiple hosts, as described in "Selecting Hosts" on page 34.
- ◆ The type of view and the name of the RDS displays in the top-right corner, below the navigation links.



- ◆ The menu bar in the RDS view provides options to start the VVR tasks, which enable you to configure and administer the RDS. For details, see “[Menus Available from the RDS and RVG Views](#)” on page 12.

The lower region is divided into a left and right pane.

- ◆ The left pane displays a condensed summary view of all RDSs present on the selected hosts. The Primary and Secondary of each RDS is also displayed. Note that these are navigation links.
- ◆ The right pane displays the following details for the selected RDS:

Primary	<p>Displays the following information about the Primary RVG:</p> <p>Primary: the name of the Primary host.</p> <p>State: the RVG state.</p> <p>Disk Group: the name of the Primary disk group.</p> <p>Checkpoints: the total number of checkpoints created on the Primary.</p> <p>Replication Log: the name of the Primary SRL, the size of the SRL, the layout of the SRL, and the number of mirrors.</p> <p>If the Primary is not present in the RDS, the information about the Secondary RVG is displayed.</p>
Replication Status	<p>Displays the status-related information of each Secondary in the selected RDS. This section displays the following information:</p> <p>Secondary: the Secondary host name.</p> <p>Data status: the data status. See “Secondary RVG Data Status” on page 38.</p> <p>Replication status: the replication status. See “Replication Status” on page 39.</p> <p>Logging to: indicates whether the SRL or DCM is currently in use (with respect to the Secondary), the number of writes by which the Secondary is behind, percentage of SRL used.</p> <p>Current mode: shows the replication mode.</p> <p>Checkpoints: the total number of checkpoints created on the RLINK.</p>

Data Volumes	<p>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:</p> <p>Name: Name of the data volume.</p> <p>Size: Size of the data volume.</p> <p>Layout: Layout of the data volume.</p> <p>Mirrors: Number of mirrors of the data volume.</p> <p>DCM: Yes or No depending on whether a DCM is associated to the data volume.</p> <p>The total size of all data volumes in the RVG is displayed in the last row of the table.</p>
Alerts	<p>Displays alert messages for this RDS. For example, an alert message displays when there is some configuration error in the RDS.</p>

Viewing RVG Information

Use the RVG view to get the information about a specific Replicated Volume Group (RVG). The RVG view displays detailed information about the selected RVG. For example, if you click the host name, *seattle*, for the Primary RVG under the RDS *hr_rvg*, the RVG view for the Primary RVG *seattle* is displayed.



Using the RVG Views

To display the RVG view, select the host name for the Primary or Secondary RVG from any VRW view. The Primary or Secondary RVG view displays.

The screenshot displays the VERITAS Volume Replicator Web Console interface. The main content area shows the configuration for the Primary RVG on host 'seattle'. The interface includes a sidebar with a tree view showing 'hr_rvg', 'seattle', and 'london'. The main panel is titled 'Primary RVG hr_rvg' and contains several sections:

- Primary RVG:** A table showing the host 'seattle', disk group 'hrdg', RVG name 'hr_rvg', and version '21'. The state is 'enabled for I/O'.
- Replication Log:** A table showing the replication log entry 'hr_srl' with a size of 4.00 GB and layout 'CONCAT'. There are 0 mirrors and no checkpoints.
- RLINKs:** A table showing the RLINK 'rlk_london_hr_rvg' with state 'ACTIVE', flags 'write enabled attached consistent connected asynchronous', protocol 'UDP/IP', and bandwidth limit 'none'.
- Data Volumes:** A table showing two data volumes: 'hr_dv01' and 'hr_dv02', both with a size of 4.00 GB and layout 'CONCAT'. There are 0 mirrors and DCH is enabled for both. The total size is 8.00 GB.
- Alerts:** A section indicating 'No Alerts'.

The toolbar at the top right of the console displays 'Profiles: root' and 'Hosts: seattle'. The bottom status bar shows 'Done' and 'Local intranet'.

Note that the type of the view, Primary or Secondary, and the name of the RVG is displayed at the far right of the toolbar.

Primary RVG View

To display the Primary RVG view, do one of the following:

- ✓ From the left pane in the RDS view or an RVG view, click the host name for the Primary RVG.
- or
- ✓ From the Summary view, click the host name for the Primary RVG.

The Primary RVG view looks like this:

The screenshot shows the Veritas Volume Replicator interface. The main window title is 'Summary of RDSs'. The left pane shows a tree view with 'hr_rvg' expanded to show 'seattle' and 'london'. The main area displays the 'Primary RVG hr_rvg' configuration. It includes sections for Primary RVG, Replication Log, RLINKs, Data Volumes, and Alerts.

Host	Disk Group	RVG	Version
seattle	hrdg	hr_rvg	21
State: enabled for I/O			

Name	Size	Layout	Mirrors
hr_srl	4.00 GB	CONCAT	0
Checkpoints: No checkpoints			

Name	State	Flags	Protocol	Bandwidth Limit
rlk_london_hr_rvg	ACTIVE	write enabled attached	UDP/IP	none
Replication IP=seattle (10.182.136.176) consistent connected asynchronous				

Name	Size	Layout	Mirrors	DCM
hr_dv01	4.00 GB	CONCAT	0	Yes
hr_dv02	4.00 GB	CONCAT	0	Yes
Total Size: 8.00 GB				

Alerts: No Alerts

Secondary RVG View

To display the Secondary RVG view, do one of the following:

- ✓ From the left pane in the RDS view or an RVG view, click the host name for the Secondary RVG.
- or
- ✓ From the Summary view, click the host name for the Secondary RVG.

The Secondary RVG view is similar to the Primary RVG view. See [“Information Displayed in the RVG Views”](#) on page 46.



Information Displayed in the RVG Views

The layouts for the RVG views are similar to the layout of the RDS view.

- ◆ The upper region of the RVG view is identical to that of the RDS view, except that the name of the RVG being viewed appears in the top right corner of the view under the navigation links. For details, see [“Information Displayed in the RDS View”](#) on page 41.
- ◆ The lower region of RVG view is also divided into a left and a right pane. The left pane contains information that is identical to that of the RDS view, except that the selected RVG is highlighted and is indicated with an arrow to distinguish it from the remaining RVGs. For details, see [“Information Displayed in the RDS View”](#) on page 41.

The lower right pane displays information about the selected RVG under the following categories:

Primary RVG	Displays only in Primary view. Displays the following information about the Primary: Host: Name of the Primary host. Disk Group: Name of the disk group. RVG: Name of the RVG. Version: RVG version. State: The state of the RVG.
Secondary RVG	Displays only in Secondary view. Displays the following information about the Secondary: Host: Name of the Secondary host. Disk Group: Name of the disk group. RVG: Name of the RVG Version: RVG version



Replication Status	<p>Displays only in Secondary view. Displays the status-related information of each Secondary in the selected RDS. This section displays:</p> <p>Secondary: the Secondary host name.</p> <p>Data Status: the data status. See “Secondary RVG Data Status” on page 38</p> <p>Replication Status: the replication status. See “Replication Status” on page 39.</p> <p>Logging To: indicates whether the SRL or DCM is currently in use (with respect to the Secondary), the number of writes by which the Secondary is behind, percentage of SRL used.</p> <p>Current Mode: shows the replication mode.</p> <p>Checkpoints: the total number of checkpoints created on the RLINK.</p>
Replication Log	<p>Displays the following information about the SRL:</p> <p>Name: Name of the SRL.</p> <p>Size: Size of the SRL.</p> <p>Layout: Layout of the SRL.</p> <p>Mirrors: Number of mirrors.</p> <p>Checkpoints: If the Primary RVG is selected, the checkpoints (the RVG or RLINK name followed by the checkpoint name) are listed here. VRW displays the string <code>No Checkpoints</code> if there are no checkpoints.</p>
RLINKs	<p>Displays the following information about the configuration and state of the RLINKs. Local and remote RLINKs are shown in pairs.</p> <p>Name: Name of the RLINK.</p> <p>State: State of the RLINK.</p> <p>Flags: Flags set on the RLINK.</p> <p>Protocol: Transport protocol.</p> <p>Replication IP: Shows the host name used for replication by the RLINK, with the IP address in parenthesis.</p> <p>For details, see “Viewing RLINK Information” on page 49.</p>



Data Volumes	Displays information about the data volumes associated with the RVG. Information includes: Name: Name of the volume. Size: Size of the volume. Layout: Layout of the volume. Mirrors: Number of mirrors. DCM: Shows whether or not DCM is associated with this volume. The total size of all data volumes in the RVG is displayed in the last row of the table.
Alerts	Alert messages display if there is a problem with the configuration. For example, if the RDS to which the RVG belongs is not configured correctly, a message displays.



Viewing RLINK Information

Use the RVG views to view information about the RLINKs. The RLINKs section in Primary and Secondary RVG views displays information about the RLINKs in the Primary and Secondary RVG respectively.

RLINK Information

The RVG view initially displays the following information about the RLINKs in the selected RVG:

Name	Name of the RLINK.
State	State of the connection (RLINK) between the Primary and the Secondary.
Flags	Displays ATTACHED when the Primary and the Secondary are attached; displays DETACHED if the Primary and Secondary are not attached. Other flags such as ENABLED, DISABLED, CONSISTENT, INCONSISTENT, CANT_SYNC, CONNECTED, DISCONNECTED also display in this column. For the meaning of these flags, see the <i>VERITAS Volume Replicator Administrator's Guide</i> .
Protocol	Transport protocol being used for replication.
Bandwidth Limit	Shows the bandwidth limit for VVR if specified, or none .
Replication IP	Shows the host name used for replication by the RLINK, with the IP address in parenthesis.



Viewing Detailed RLINK Information

A (+) icon is available for each displayed RLINK. The view can be expanded to show detailed information for the RLINK.

▼ To expand the view and display information about the attributes of the RLINK

- ❖ Click the (+) icon near the RLINK name. The (+) icon changes to a minus (-) icon.

The expanded view shows the following information:

- ◆ Replication settings of the RLINK
- ◆ Information about the Remote RLINK

▼ To collapse the expanded view

- ❖ Click the (-) icon.

Replication Settings of the RLINK

The expanded view displays the following information about the attributes of the selected RLINK:

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.

Information About the Remote RLINK

The expanded view displays the following information about the remote RLINK:

Replication IP	Name of the remote host and IP in parenthesis.
Name	Name of the remote RLINK.
Protocol	Transport protocol being used for replication.
Bandwidth Limit	Shows the bandwidth limit for VVR if specified, or none .

▼ **To expand the view and display more information for the remote RLINK**

- ❖ Click the (+) icon for the remote RLINK. The (+) icon changes to a minus (-) icon. The expanded view shows the replication settings for the remote RLINK.

▼ **To collapse the expanded view**

- ❖ Click the (-) icon.





Administering VERITAS Volume Replicator

4

VERITAS 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 Users

To use VRW to manage your VVR objects, you must log in as a user with access privileges to VRW. By default, `root` is the only user. VRW enables the `root` user to create additional users to have access to the VRW console and, therefore, to the tasks for administering VVR.

This section describes characteristics of VRW users, including:

- ◆ [User Roles](#)
- ◆ [User Profiles for Central Managing of Hosts](#)

User Roles

You can provide users with VVR administrative access without full `root` privileges. You can also create users with guest privileges, for users who need access to the reporting functionality of VRW, without needing privileges to administrative tasks. VRW provides the following user roles:

Role	Description
<code>root</code>	The root user has full privileges. By default, <code>root</code> is the only user with access to VVR and VRW. The root user has access to all menu operations within VRW, including administering users.



Role	Description
administrator	The administrator role allows users to use VRW to administer VVR. An administrator user has access to all menu operations within VRW, except for administering users.
guest	The guest role allows users to use VRW in view-only mode. A guest user has no access to any menu operations within VRW, except Edit Profile List . The only task guest users can perform is to edit their own user profiles to manage the list of VVR hosts that they can view. Guest users cannot perform any VVR administration tasks.

User Profiles for Central Managing of Hosts

If you install VRW on multiple hosts, VRW provides the ability to view and manage RDSs and RVGs on several hosts in a single console. This functionality is controlled through user profiles and host lists.

Each VRW user has one or more profiles, which define the VVR hosts for which that user can view and manage VVR. When the user logs into VRW, the VRW console provides a centralized view of the RDSs and RVGs on all of the hosts in the user's profile.

To use VRW on multiple hosts with users and user profiles, you need to do the following:

- ◆ [Configuring VRW Host Access](#)
- ◆ [Configuring Users](#)
- ◆ [Configuring User Profiles](#)

Configuring VRW Host Access

Typically, a `root` user manages a group of hosts, and therefore wants to use the same login IDs and passwords for all the hosts in the group. To ensure consistency, the `root` user should choose a particular host, called the *management host*, from which to create and edit VRW users.

Note When editing VRW users, the `root` user should always log into the management host.

For a host to be associated to a user's profile, the host must be running VRW. In addition, the host must allow access from VRW on the management host.

By default, an installation of VRW does not allow other VRW hosts to have access. The `root` user must explicitly configure each host to allow access from VRW on the management host. For example, suppose you install VRW on `seattle` and on `london`.

By default, logging into `seattle` does not allow you to view information on `london`, or vice versa. To manage VVR information on both `seattle` and `london` with a single VRW console, you must explicitly enable remote access between the hosts. (You must also create a user whose profile includes both `seattle` and `london`.)

Note To ensure consistency, each host in the group should allow access from VRW on the other hosts. This allows any changes to user profiles to be propagated to all the hosts so that VRW users see the same user profiles regardless of which host they log into.

Configure hosts as described in “[Configuring Hosts](#)” on page 56.

Configuring Users

The `root` user creates and deletes VRW users. For each VRW user that is created, the `root` user must assign a list of hosts for which that user can view and manage VVR. This list becomes the user’s default profile. The default profile name is the username. Use the Edit User wizard to configure users, as described in “[Editing Users](#)” on page 57.

Configuring User Profiles

Each VRW user can create or edit additional profiles to view a subset of the hosts in that user’s own default profile. Multiple profiles may be useful if a login is used by more than one person or if you want to specify views for particular sets of hosts. For example, the `admin` login could be shared by two administrators who manage RDSs on different sets of hosts. Use the Edit Profile wizard to create or edit profiles, as described in “[Editing Profiles](#)” on page 60.



Configuring Hosts

The `root` user must explicitly configure each host to allow access from VRW on the management host.

Prerequisites:

- ◆ For a host to be associated to a user's profile, the host must be running VRW.
- ◆ Each host name in the list must be resolvable without the domain name.
- ◆ Also configure the host on which the browser is running so that each host name in the host list is resolvable without the domain name.

▼ To enable access between hosts

Configure each host to have a `root` host list which includes itself and the other hosts to be managed by the same `root` user. To create the `root` host list, determine the list of hosts to be managed by the same `root` user. Perform the following steps on each host:

Note The `root` user host list indicates which hosts are available for the `root` user to assign when creating additional VRW users. The `root` user can only assign the hosts from this host list to another VRW user's default host list (default profile).

1. For each host to be added to the host list, log in as `root`.
2. Create a host list entry in the file
`/opt/VRTSweb/VERITAS/vvr/WEB-INF/classes/vvrSystem.properties`.
The host list entry is in one of the following forms:
 - ◆ A plus sign (+) enables any remote VRW hosts to have access to the local host. Use this form to manage all VRW hosts in the enterprise from one console. Include the following host list entry on each host:

```
user.root.root.hostlist=+
```

- ◆ A comma-separated list enables the specified hosts to have access to each other. Use this form to create a list of hosts that are managed by the same `root` user.

```
user.root.root.hostlist=hostname[, hostname...]
```

On each of the hosts in this host list, create an identical host list entry. For example, to enable access between `london` and `seattle`, include the following host list entry on both hosts:

```
user.root.root.hostlist=london, seattle
```

In the above example, the `root` user could log into `london` and create a user with access to `seattle`. Suppose you have a third host, `tokyo`. The `root` user on `london` could not add `toyko` to a user's profile because it is not in the list of managed hosts. To include `tokyo`, include the following host list entry on all three hosts:

```
user.root.root.hostlist=london, seattle, tokyo
```

3. After editing the `/opt/VRTSweb/VERITAS/vvr/WEB-INF/classes/vvrSystem.properties` file, restart VRW.
4. Repeat steps 1-3 on each of the hosts in this host list. Each of the hosts in the host list should contain an identical host list entry.

Editing Users

The Edit User wizard enables the `root` user to add new users, delete users, or change users. Using the Edit User wizard, the `root` user assigns a list of hosts for each user. This list becomes the user's default profile.

Note The user list is updated on all the hosts in the profile list.

Prerequisites for adding, deleting, or changing users:

- ✓ VERITAS Volume Replicator Web Console must be up and running on all the hosts in the profile list.
- ✓ `root` user on all the hosts to be added must have the same password.

Note Each user ID to be added must be a valid login ID on the local host and on each of the hosts in the profile list.

▼ To add users

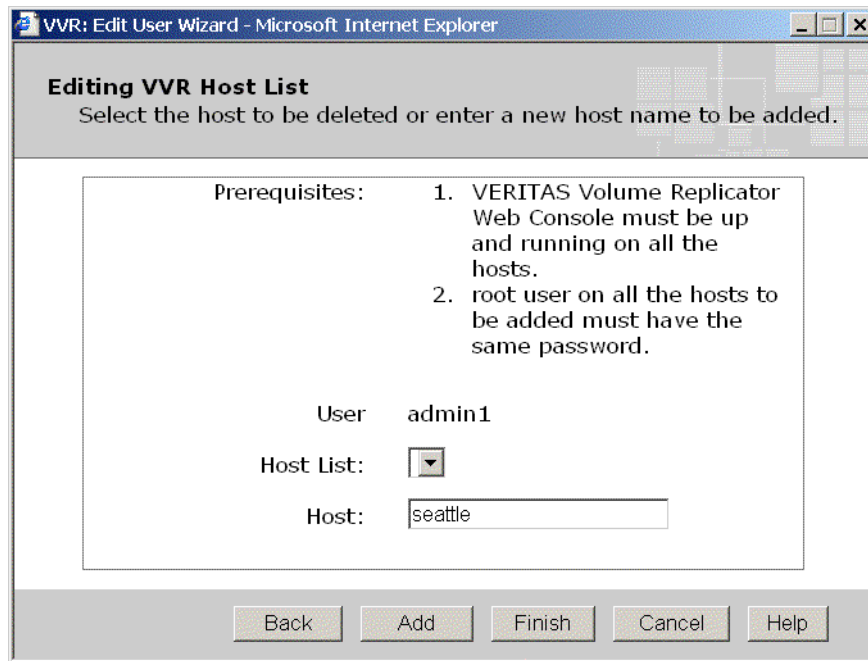
1. Log on as `root`. You must be `root` to edit users.
2. Select **Tools > Edit User List**.



3. Complete the Editing User List page as follows:

User ID:	Enter a valid login ID (username).
User Role:	Select either Guest or Administrator for the User Role

4. Click **Add**.



5. In the Editing VVR Host List page, edit the host list as follows:
 - ◆ To add a host, enter the name of the host in the Host field and click **Add**.
A host list can have up to 8 hosts.
 - ◆ To delete a host, select the name from the drop-down Host List and click **Delete**.
Continue to add or delete hosts, or click **Finish**.
6. In the Summary page, click **OK**. A message displays the status of the request.
7. Click **Close**.



▼ To delete users

1. Log on as `root`. You must be logged in as `root` to edit users.
2. Select **Tools > Edit User List**.
3. Select a user ID from the drop-down User List.
4. Click **Delete**. Continue to add or delete users, or click **Finish**.
5. In the Summary page, click **OK**. A message displays the status of the request.
6. Click **Close**.

▼ To change roles for a user

To change roles for a user, you must delete the user and add the user again with the new role.

1. Log on as `root`. You must be logged in as `root` to edit users.
2. Select **Tools > Edit User List**.
3. Select a user ID from the drop-down User List. Click **Delete**.
4. Reenter the User ID in the Editing User List page, and select the appropriate User Role.
5. Click **Add**.
6. In the Editing VVR Host List page, edit the host list as follows:
 - ◆ To add a host, enter the name of the host in the Host field and click **Add**.
A host list can have up to 8 hosts.
 - ◆ To delete a host, select the name from the drop-down Host List and click **Delete**.
7. To add additional hosts, repeat [step 4](#) through [step 6](#). Otherwise, click **Finish**.
8. In the Summary page, click **OK**. A message displays the status of the request.
9. Click **Close**.



Editing Profiles

VRW users can manage their VVR objects on multiple hosts. Each user can have one or more profiles, which define the list of hosts to access. VRW users other than the `root` user use the Edit Profile wizard to add, delete, and edit profiles. A user can edit only the profiles associated with that user.

Note The profile list is updated on all the hosts in the host list.

Prerequisites for editing a profile:

- ✓ VERITAS Volume Replicator Web Console must be up and running on all the hosts in the host list.
- ✓ *username* user on all the hosts to be added must have the same password.

▼ To add a profile

1. Select **Tools > Edit Profile List**.

VVR: Edit Profile Wizard - Microsoft Internet Explorer

Editing Profile List

Select the profile to be edited, deleted or enter a new profile name to be added.

Prerequisites:

1. VERITAS Volume Replicator Web Console must be up and running on all the hosts in the profiles.
2. Iviste user on all the hosts to be added must have the same password.

Profile List:

Profile:

Add Finish Cancel Help

Note Each user has a default profile that is named with the *username*. Because only the `root` user can edit the default profile, it does not appear in this profile list.

2. In the Editing Profile List page, enter a name for the new profile in the Profile field and select **Add**.
3. In the Editing VVR Host List page, edit the host list as follows:
 - ◆ To add a host, enter the name of the host in the Host field and click **Add**.
A host list can have up to 8 hosts.

Note The default profile list, created by the `root` user for each user, determines the hosts which can be included in additional profiles.

- ◆ To delete a host, select the name from the drop-down Host List and click **Delete**.
4. Click **Finish** when you are done.
 5. Click **OK** to confirm the host list changes. A message indicates the status of the changes on the local host and each of the other hosts in the profile list.
- ▼ **To remove a profile**
1. Select **Tools > Edit Profile List**.
 2. In the Editing Profile List page, select the profile from the drop-down Profile List, and click **Delete**.
 3. Click **Finish** when you are done.
 4. Click **OK** to confirm that you want to delete the profile. A message indicates the status of the changes on the local host and each of the other hosts in the profile list.



▼ **To edit a profile**

1. Select **Tools > Edit Profile List**.
2. In the Editing Profile List page, select the profile from the drop-down Profile List, and select **Edit**.
3. Complete the Editing VVR Host List page as follows:
 - ◆ To add a host, enter the name of the host in the Host field and click **Add**.
A host list can have up to 8 hosts.
 - ◆ To delete a host, select the name from the drop-down Host List and click **Delete**.
4. Click **Finish** when you are done.
5. Click **OK** to confirm the host list changes. A message indicates the status of the changes on the local host and each of the other hosts in the profile list.



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 one or more volumes 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, refer to the information provided in [“Verifying the Data on the Primary and Secondary Volumes”](#) on page 64. To synchronize the Primary and Secondary volumes, see [“Synchronizing Volumes on the Local Host and Remote Hosts”](#) on page 65.

Prerequisites for associating a volume to an RDS

- ✓ 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 using the Synchronize Volume wizard.



▼ **To associate a volume to an RDS**

1. Select any view for the RDS for which the data volume is to be associated.
2. From the RDS or RVG view, choose **Configuration > Associate Volume**.
3. In the Volume Selection page, select the name of one or more data volumes to be associated to the RDS. 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`.
4. Click **Next**.
5. In the confirmation page, check the box to indicate that the prerequisites have been met and click **OK**. 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.
6. In the status page, click **Close**. VRW updates and displays the RDS or RVG view from which you accessed the Associate Volume task.

Verifying the Data on the Primary and Secondary Volumes

The Verify Volume wizard enables you to verify whether there are any differences between remote volumes and the corresponding local volumes. You can use this command when the volumes are not associated with an RVG and the application is inactive (the volumes are not in use). The Verify Volume wizard only reports the amount of data that is different in percentage between remote and local volumes; it does not synchronize remote volumes with local volumes. VVR also allows you to verify the data volumes when replication is in progress. For more information on using the online data verification feature, refer to “[Verifying RVG Online](#)” on page 92.

This command can be useful to make sure that the data volumes are identical before adding them to an RDS.

Synchronizing Volumes on the Local Host and Remote Hosts

The Synchronize Volume wizard enables you to synchronize remote volumes with local volumes when the volumes are not associated with an RVG and the volumes are not in use. The data in the volumes on the local host, where you enter the command, is transferred over the network to the volumes on the remote host.

Use the Synchronize Volume wizard only to synchronize volumes that are not part of an RVG. For example, before adding a volume to an RDS, synchronize the volume using the Synchronize Volume wizard, and then add it to the RDS.

Using the Synchronize Volume wizard, you can synchronize remote volumes with local volumes using one of the following options:

- ◆ Difference-based synchronization
- ◆ Full synchronization

We recommend that the names of the volumes on the local and remote hosts be the same. However, you can synchronize volumes with different names on the local and remote hosts using the Synchronize Volume wizard.

Note The Synchronize Volume wizard cannot be used for an RVG that has volume-set component volumes associated to it.



Resizing a Data Volume in a Replicated Data Set

Use the Resize Volume task to resize a data volume in 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.

Prerequisites for Resizing a Data Volume in an RDS

- ✓ 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 any view for the RDS for which the data volume is to be resized.
2. From the RDS or RVG view, choose **Configuration > Resize Volume**.
3. View the Warnings page and click **Next**.
4. Complete the Resize Volume Options page as follows:

Data Volume	Select the name of the data volume to be resized. For example, <code>hr_dv01</code> .
Increase By	To increase the data volume by a specified size, select Increase By .

Decrease By	To decrease the data volume by a specified size, select Decrease By .
Resize To	To resize the data volume to a specified size, select Resize To . The default size unit is megabytes. To specify a size unit, select either sectors, kilobytes, or gigabytes from the drop-down list. Depending on the selected option, specify the size in the provided field.

5. Click **OK** to process the Resize Volume request.
6. Click **Close** in the status page of the Resize Volume wizard. VRW updates and displays the RDS or RVG view from which you accessed the Resize Volume task.

Dissociating a Data Volume from its Replicated Data Set

Perform the Dissociate Volume task to remove one or more 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 **Replication > Stop**.

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 remove a data volume from its RDS even when the Primary RVG has not been stopped, select the **Remove 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. From the RDS or RVG view, choose **Configuration > Dissociate Volume**.
3. Complete the Volume Selection page as follows:

Data Volume	Select the name of the data volumes that you want to dissociate. 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. We recommend that the applications and the Primary RVG be stopped before dissociating any data volumes.
Option: Dissociate the data volume even if the Primary RVG has not been stopped	Click to dissociate the data volume even if the Primary RVG has not been stopped.

4. Click **OK** to dissociate the volume.
5. Click **Close**. VRW updates and displays the RDS or RVG view from which you accessed the Dissociate Volume task.



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

Use **Configuration > Change Replication Settings** 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. For more information on replication settings, see [“Adding a Secondary to the RDS”](#) on page 20.

▼ To change replication settings to a Secondary in an RDS

1. Select the Secondary RVG view of the RDS for which you want to change the replication settings.
2. Choose **Configuration > Change Replication Settings**.
3. Complete the Change Replication Settings page as follows:

Replication Mode Asynchronous, Synchronous-Override	The current replication mode is displayed by default. From the drop-down list, select the required mode of replication. To set the replication to <i>synchronous</i> mode, select Synchronous-Override . To set the replication to <i>asynchronous</i> 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	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.

4. Click **OK** to change the replication settings for the selected Secondary. A message indicates whether the operation succeeded or failed.
5. In the status page, click **Close**. The Secondary RVG view displays the new replication settings.

Starting Replication to a Secondary

Use **Replication > Start** 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. For more information, refer to [“Synchronizing the Secondary and Starting Replication”](#) on page 25.



Pausing Replication to a Secondary

Use **Replication > Pause** 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. From the Secondary RVG view of the required RDS, choose **Replication > Pause**.
2. In the Pause Replication dialog box, click **OK** to pause replication. A message indicates whether the operation succeeded or failed.
3. In the status page, click **Close**. The Secondary RVG view displays the new replication status.

Resuming Replication to a Secondary

Use **Replication > Resume** 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.

▼ To resume replication to a Secondary

1. From the Secondary RVG view of the required RDS, choose **Replication > Resume**.
2. In the Resume Replication Summary page, click **OK** to resume replication. A message indicates whether the operation succeeded or failed.
3. In the status page, click **Close**. The Secondary RVG view displays the new replication status.

Stopping Replication to a Secondary

Use **Replication > Stop** 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 **Replication > Pause**.

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. From the Secondary RVG view of the required RDS, choose **Replication > Stop**.
2. Complete the Stop Replication dialog box as follows:

Option: 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.
---	--

3. Click **OK** to stop replication. A message indicates whether the stop replication operation succeeded or failed.
4. In the status page, click **Close**. The Secondary RVG view displays the new replication status.



Changing the IP Addresses Used for Replication

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

Prerequisites for changing the IP addresses used for replication:

- ✓ 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. Select the Secondary host for which you want to change the replication settings in the RDS view.
2. From the Secondary RVG view, select **Configuration > Change IP**.
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.
4. Click **OK** to change the replication IP for the selected Secondary. A message indicates whether the operation succeeded or failed.
5. In the status page, click **Close**. The RVG view displays the new replication settings.

Administering the Replicated Data Set

This section describes the following tasks:

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

Removing a Secondary from a Replicated Data Set

Use **Configuration > 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 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.

Prerequisite:

- ✓ Stop replication to the Secondary that is to be removed. Select **Replication > Stop**. For details, see “[Stopping Replication to a Secondary](#)” on page 73.

▼ To remove a Secondary from an RDS

Caution The Remove Secondary operation is irreversible.

1. Select the Secondary RVG to be removed.
2. From the Secondary RVG view, choose **Configuration > Remove Secondary**.
3. In the Remove Secondary dialog box, click **OK** to remove the Secondary. A message indicates whether the Remove Secondary operation succeeded or failed.
4. In the status page, click **Close**. The RDS or RVG view displays the new RDS configuration. If the removed RVG was on the local host, VRW displays the Summary view, which no longer shows the RDS in which the removed RVG belonged. Otherwise, VRW displays the current RDS view.



Removing a Primary RVG

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

Prerequisite for deleting a Primary RVG:

- ✓ All Secondaries in the RDS must be removed. For more information on how to remove a Secondary from an RDS, see [“Removing a Secondary from a Replicated Data Set”](#) on page 75.

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. Select the RDS from which to remove the Primary RVG
2. From the RDS view or Primary RVG view for the RDS, select **Configuration > Remove Primary**.
3. Complete the Remove Primary dialog box by checking the “Force removal of the Primary RVG even if it is not yet stopped” checkbox. Click **OK**. A message indicates whether the Remove Primary operation succeeded or failed.
4. In the status page, click **Close**. The RDS or RVG view displays the new RDS configuration. If the removed RVG was on the local host, VRW displays the Summary view, which no longer shows the RDS in which the removed RVG belonged. Otherwise, VRW displays the current RDS view.

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. VRW enables you to create a maximum of 46 checkpoints.

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

▼ To start a checkpoint

1. Select an RDS or Primary or Secondary RVG view.
2. Select **Tools > Start Checkpoint**.
3. In the Start Checkpoint dialog box, enter a checkpoint name and click **OK**. A message displays the status of the checkpoint.
4. Click **Close**.

Ending Checkpoints

▼ To end a checkpoint

1. Select an RDS or Primary or Secondary RVG view.
2. Select **Tools > End Checkpoint**.
3. In the End Checkpoint dialog box, click **OK** to confirm that you want to end the checkpoint which is currently started. A message displays the status of the checkpoint.
4. Click **Close**.



Viewing Checkpoints

Checkpoints are displayed in the Primary RVG view for the RDS.

- ❖ To display information about a checkpoint, click the checkpoint name, or move the cursor over it.

The information includes usage, % of log used, and the status of the checkpoint (started or complete).

The RDS view indicates the number of checkpoints available.

Deleting Checkpoints

▼ To delete a checkpoint

1. Select an RDS or a Primary or Secondary RVG view.
2. Select **Tools > 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.
4. Click **OK**. A message displays the status of the checkpoint delete operation.
5. Click **Close**.

Creating RVG Snapshots

VRW enables you to create snapshots that are images of the data volumes at a given point in time. The data in the original volume may change; however, the snapshot can still be used as a stable and independent copy for various purposes. VRW provides two methods of creating snapshots:

- ◆ [The Instant Snapshot Feature](#) and
- ◆ [Traditional Snapshots](#)

Note If the secondary RVG is inconsistent, then VVR does not allow you to create snapshots of the volumes under this RVG.

The Instant Snapshot Feature

VRW enables you to create instant volume snapshots using the Snapshot wizard. The wizard takes snapshots of the data volumes in the RVG. However, the snapshot volumes are not part of the RVG. Each data volume in an RVG can have more than one snapshot volume. See the *VERITAS Volume Replicator Administrator's Guide* for more information about instant snapshots.

The instant snapshot feature provides the following methods to create instant snapshots:

- ◆ “Instant Full Snapshots” on page 79
- ◆ “Instant Space-Optimized Snapshots” on page 82
- ◆ “Instant Plex-Breakoff Snapshots” on page 83

Note For VxVM ISP volumes, only instant full and instant space-optimized snapshots can be created; the Cache Size option cannot be used.

Instant Full Snapshots

The Snapshot wizard enables you to create an instant full snapshot of all the volumes in the RVG, at a single point in time. The Snapshot wizard creates the snapshots for all the volumes in the RVG, unless you specify volumes to exclude.

Prerequisites:

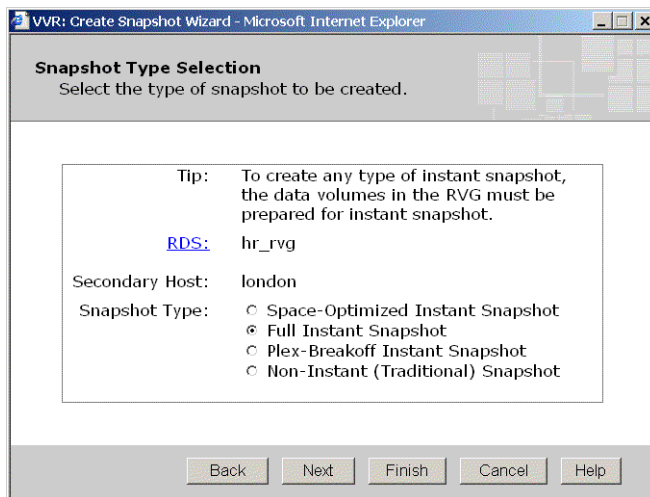
- ✓ All data volumes in the RVG, for which snapshots are to be created, must be prepared for snapshot operation.
- ✓ For each full instant snapshot volume, create the snapshot volume of the same size (or greater) as the original volume and with a prefix. Make sure the snapshot volumes follow a proper naming convention such that whenever required, you can easily relate the snapshot volume names with the original volumes names.
- ✓ When creating plexes, use appropriate prefixes so that specific plexes can also be used for the snapshot operation.

For information on preparing volumes for the snapshot operation, refer to the *VERITAS Volume Replicator Administrator's Guide*.



▼ **To create an instant full snapshot**

1. Pause the replication to the Secondary.
2. Select the RVG for which you want to create the snapshot.
3. From the RVG view, select **Tools > Create Snapshot**.
4. In the Prerequisites Information page, read the information and click **Next**.



5. In the Snapshot Type Selection page, select **Full Instant Snapshot** and click **Next**.
6. Complete the Optional Parameters page as follows:

Snapshot Prefix	Enter a prefix to be used with each of the volume names. This helps to identify the snapshots.
Synchronize the full instant snapshot volumes	Click the checkbox to automatically synchronize the snapshot volumes.

To create an instant full snapshot for all volumes, skip to the summary page by clicking **Finish**. Continue from [step 10](#) on page 81.

To exclude certain volumes from the snapshot, or to specify a different type of snapshot for some volumes, click **Next** to proceed. At any point, click **Finish** to skip to the summary page.

7. To exclude volumes from the snapshot, select the volumes to exclude from the drop-down list in the Exclude Volumes page. No snapshots are created for the selected volumes. Click **Next**.
8. To create space-optimized instant snapshots for some volumes, select the volumes from the drop-down list in the Space-optimized Instant Snapshot Volume page. Enter the cache size or the cache name to be used for the space-optimized instant snapshots. Space-optimized instant snapshots are created for the selected volumes, instead of full snapshots. Click **Next**.
9. To create plex-breakoff instant snapshots for some volumes, select the volumes from the drop-down list in the Plex-Breakoff Instant Snapshot Volume page. Optionally, enter a plex prefix. Plex-breakoff instant snapshots are created for the selected volumes, instead of full snapshots. Click **Finish**.
10. In the Snapshot Summary page, click **OK** to create the specified snapshots, or **Back** to change your selections. A message displays the status of the snapshot creation.
11. Click **Close**.
12. Resume replication on the Secondary.



Instant Space-Optimized Snapshots

The Snapshot wizard can create an instant space-optimized snapshot of all the volumes in the RVG at a single point in time. The Snapshot wizard uses a cache object that functions as a space-optimized persistent store.

The Snapshot wizard enables you to specify cache objects in two ways:

- ◆ by cache name. This method uses a single cache object for all of the space-optimized snapshots. You must create the cache object within the same disk group as the data volumes before creating the snapshot.
- ◆ by cache size. In this case, a separate cache object is created for every space-optimized snapshot.

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

Prerequisites:

- ✓ All data volumes in the RVG, for which snapshots are to be created, must be prepared for snapshot operation.
- ✓ If you are using a single cache object, create the cache object before creating the snapshot.

▼ To create a space-optimized snapshot

1. Pause replication from the Secondary.
2. Select the RVG for which you want to create the snapshot.
3. From the RVG view, select **Tools > Create Snapshot**.
4. In the Prerequisites Information page, read the information and click **Next**.
5. In the Snapshot Type Selection page, select **Space-Optimized Instant Snapshot** and click **Next**.
6. Complete the Optional Parameters page as follows:

Select **Cache Size** and specify a cache size and select the appropriate units, or select **Cache Name** and enter a cache name. Optionally, specify a snapshot prefix.

To create a space-optimized snapshot for all volumes, skip to the summary page by clicking **Finish**. Continue from [step 10](#) on page 83.

To exclude certain volumes from the snapshot, or to specify a different type of snapshot for some volumes, click **Next** to proceed. At any point, click **Finish** to skip to the summary page.

7. To exclude volumes from the snapshot, select the volumes to exclude from the drop-down list in the Exclude Volumes page. No snapshots are created for the selected volumes. Click **Next**.
8. To create full instant snapshots for some volumes, select the volumes from the drop-down list in the Full Instant Snapshot Volume Selection page. Click the checkbox to **Synchronize the full instant snapshot volumes** to use that option. Click **Next**.
9. To create plex-breakoff instant snapshots for some volumes, select the volumes from the drop-down list in the Plex-Breakoff Instant Snapshot Volume Selection page. Optionally, enter a plex prefix. Plex-breakoff instant snapshots are created for the selected volumes, instead of full snapshots. Click **Finish**.
10. In the Snapshot Summary page, click **OK** to create the specified snapshots, or **Back** to change your selections. A message displays the status of the snapshot creation.
11. Click **Close**.
12. Resume replication.

Instant Plex-Breakoff Snapshots

The instant plex-breakoff snapshot feature enables you to create plex-breakoff snapshots just like the traditional snapshot feature.

Prerequisites for creating instant plex-breakoff snapshots

- ✓ All data volumes in the RVG, for which snapshots are to be created, must be prepared for snapshot operation.
- ✓ Make sure the volumes for which you want to create plex-breakoff snapshots already have the appropriate plexes created and are in a SNAPDONE state.
- ✓ When creating plexes, use appropriate prefixes so that specific plexes can also be used for the snapshot operation.

For example, *plexprefix-volume_name*

Note If you do not specify the `plexprefix` attribute when creating the plex-breakoff snapshots, a plex that is in the SNAPDONE state gets selected automatically.



▼ **To create a plex-breakoff snapshot**

1. Pause replication from the Secondary.
2. Select the RVG for which you want to create the snapshot.
3. From the RVG view, select **Tools > Create Snapshot**.
4. In the Prerequisites Information page, read the information and click **Next**.
5. In the Snapshot Type Selection page, select **Plex-Breakoff Instant Snapshot** and click **Next**.

6. Complete the Optional Parameters page as follows:

Optionally, enter a plex prefix.

To create a plex-breakoff snapshot for all volumes, skip to the summary page by clicking **Finish**. Continue from [step 10](#) on page 84.

To exclude certain volumes from the snapshot, or to specify a different type of snapshot for some volumes, click **Next** to proceed. At any point, click **Finish** to skip to the summary page.

7. To exclude volumes from the snapshot, select the volumes to exclude from the drop-down list in the Exclude Volumes page. No snapshots are created for the selected volumes. Click **Next**.
8. To create full instant snapshots for some volumes, select the volumes from the drop-down list in the Full Instant Snapshot Volume Selection page. Click the checkbox to **Synchronize the full instant snapshot volumes** to use that option. Click **Next**.
9. To create space-optimized instant snapshots for some volumes, select the volumes from the drop-down list in the Space-optimized Instant Snapshot Volume Selection page. Select **Cache size** and specify a cache size and select the appropriate units, or select **Cache Name** and enter a cache name. Optionally, specify a snapshot prefix. Space-optimized instant snapshots are created for the selected volumes, instead of full snapshots. Click **Finish**.
10. In the Snapshot Summary page, click **OK** to create the specified snapshots, or **Back** to change your selections. A message displays the status of the snapshot creation.
11. Click **Close**.



Traditional Snapshots

The Snapshot wizard enables you to create snapshots using the traditional snapshot feature, which breaks off mirrors from the data volumes in an RVG, thus providing snapshots of the data volumes in the RVG. The Snapshot wizard takes a snapshot of all the volumes in the RVG at a single point in time. It creates a snapshot volume for each data volume in the RVG. Each data volume in an RVG can have more than one snapshot volume.

The Snapshot wizard enables you to specify a prefix for the names of the snapshot plexes. If you specify the prefix *month*, the name of each snapshot data volume starts with *month*; the resulting snapshot volume is named *month-dv_name*. For example, the data volume `hr_dv01` can have snapshot volumes such as `june-hr_dv01`, `july-hr_dv01`.

See the *VERITAS Volume Replicator Administrator's Guide* for more information about traditional snapshots.

Prerequisite:

- ✓ Create snapshot plexes for each data volume in the RVG

▼ To create a traditional snapshot

1. Pause the replication from the Secondary.
2. Select the RVG for which you want to create the snapshot. From the RVG view, select **Tools > Create Snapshot**.
3. In the Prerequisites Information page, read the information and click **Next**.
4. In the Snapshot Type Selection page, select **Traditional Snapshot** and click **Next**.
5. Complete the Optional Parameters page as follows:
Optionally, enter a prefix.
6. In the Snapshot Summary page, click **OK** to create the specified snapshots, or **Back** to change your selections. A message displays the status of the snapshot creation.
7. Click **Close**.
8. Resume replication from the Secondary.

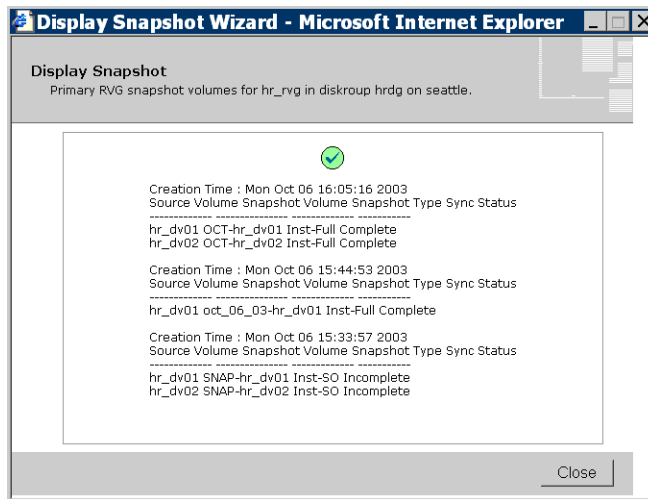


Displaying Snapshots

VVR enables you to display a report on the snapshot volumes that exist for a selected RVG.

▼ To display snapshots

1. Select the RVG for which you want to display a snapshot. From the RVG view, select **Tools > Display Snapshot** to show a report of snapshots for the selected RVG.



2. Click **Close** to return to the RVG view.

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, refer to “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*. To change the replication setting for SRL protection, see “[Changing the Replication Settings](#)” on page 70.

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.

Prerequisite for incrementally resynchronizing the Secondary

- ✓ The RVG must have the `dcm_logging` flag set.

▼ To incrementally resynchronize the Secondary

1. From any view of the RDS for which the resynchronization is to be started, select **Tools > Replay DCM**.
2. Complete the Summary page as follows:

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.
-------------	--

3. Click **OK** 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.

Resizing the Primary SRL

The Primary 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.

Prerequisite

- ✓ 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. Select **Configuration > Resize SRL**.
2. Complete the Resize SRL dialog box as follows:

Increase By	To increase the data volume by a specified size, select Increase By .
Resize To	To resize the data volume to a specified size, select Resize To . The default size unit is megabytes. To specify a size unit, select either sectors, kilobytes, or gigabytes from the drop-down list. Depending on the selected option, specify the size in the provided field.

3. Click **OK**.



Performing Data Verification

VVR provides two methods to validate the data at the Secondary site; online data verification and offline data verification. These methods are used to verify that the data on the Secondary data volumes is identical to the Primary data volumes.

- ◆ Online data verification allows you to validate the data even when the replication is in progress. In this method instead of the actual volumes the point-in-time snapshots are compared. This method is referred to as online data verification.
- ◆ Offline data verification can be performed only when the replication is not active.

Verifying RVG Offline

Offline data verification can be performed only when the replication is not active and can be used in two situations:

- ◆ If you have already created the Primary and Secondary volumes and replication is in progress, you can pause replication when the Secondary is up-to-date. Then perform data verification between the corresponding Primary and Secondary volumes to make sure that the data volumes on the Primary and Secondary are the same. This command performs a comparison of the checksums on the corresponding Primary and Secondary volumes.
- ◆ Before adding new volumes to the RDS, you can validate the data on the new volume. If you find that the Primary data and Secondary data do not match, then you can use some manual means such as backup and restore or some other method to make the data on the new secondary volume the same as the Primary and then add it to the RDS.

VVR enables you to verify whether the data on the Secondary is identical to the data on the Primary data volumes when the application is inactive. The Verify RVG Offline task verifies and reports any differences between the data volumes associated with the Secondary RVG and the corresponding Primary RVG. The Verify RVG Offline task only reports whether the Primary and Secondary volumes are identical or not. It does not make them identical.

Prerequisite:

- ✓ Make sure that any applications writing to the Primary data volumes are stopped before starting the verification of the RVG.



▼ To verify RVG offline

1. Pause replication to the Secondary that is to be verified.
2. Select the Secondary RVG to be verified.
3. From the Secondary RVG view, select **Tools > Verify RVG Offline**.
4. Complete the Verify RVG Offline dialog box as follows:

Secondary Host	Select the Secondary host.
Option: Display output in reverse order.	Select the checkbox only if you want to view the output of Verify RVG in reverse order.

5. Click **OK**. A message displays the status of the operation.
6. Click **Close**.



Verifying RVG Online

The space-optimized snapshots that are created using the **Snapshot** wizard can be used to verify whether the data on the Primary and Secondary RVG volumes is the same. The major advantage of this feature over the Verify RVG Offline task is that you do not need to stop the replication. The verification can be done even while the replication is in progress because the point-in-time snapshots, and not the volumes, are compared. This feature is useful if you want to check the integrity of the data volumes on the Secondary when replication is in progress.

The Verify RVG Online task creates the space-optimized snapshots on the Primary and the Secondary before it proceeds with performing online data verification. The Verify RVG Online task also ensures that the snapshots are taken only after the replication has been paused using the `vxibc freeze` command.

As a result there may be a momentary pause in the replication. It is necessary to freeze the writes so that the snapshots can be taken at an identical point in replication time, on each of the required hosts.

The Verify RVG Online task then verifies the data between the remote and local hosts by comparing the space-optimized snapshots.

Note Verify RVG Online cannot verify an RVG if any of its data volumes is in a volume set.

Prerequisite:

- ✓ Volumes must be prepared for instant snapshots.

To perform online data verification:

1. Select the Secondary RVG for the RDS that you want to verify.
2. From the Secondary RVG view, select **Tools > Verify RVG Online**.
3. Complete the option page as follows:

You must specify either a cache size or a cache name.

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.
-------------------	---



Cache Name	Select Cache Name and specify a name for the pre-created cache object on which the snapshots for the volumes in the specified RVG will be created.
-------------------	--

- Click **Next**. Specify additional options as follows:

Other options: Destroy the cache object and the snapshot volumes after the data verification has completed.	Check this option to destroy the snapshot volume and the cache object after the data verification has proceeded successfully. This is the default.
Keep the cache object but destroy the snapshot volumes.	Check this option to preserve the cache object to be reused when creating future snapshots.
Keep the cache object and the snapshot volumes.	Check this option to preserve the snapshot volumes. The cache object is also preserved along with the snapshot since the snapshot cannot exist without the cache object.

- Click **Finish**.
- The Summary shows. Select the checkbox for the option if you want to display output in reverse order. Click **OK**.
- The Status shows. Click **Close**.

Using the Command Window

Some administrative tasks are not available through the task wizards in VRW. The Command window enables you to enter certain CLI commands without leaving the VRW console. To use the Command window, select **Tools > Command**.





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, refer to 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 **Change Role > Migrate** 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.



Prerequisites for migrating the Primary

- ✓ 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

Perform the following steps to migrate a specified healthy Primary to a Secondary. After the migration, the Secondary becomes the new 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. In the RDS, select the Secondary RVG to which you want to migrate.
4. From the Secondary RVG view, choose **Change Role > Migrate**.
5. In the Migration dialog box, click **OK**. A message indicates whether the migration succeeded or failed.
6. In the status page, click **Close**. The Primary RVG view for the new Primary is displayed. By default, the migration operation starts replication from the new Primary.
7. Restart the application. Because the application was stopped before the migration, an application recovery is not required. To start the application before enabling replication, first, pause replication using **Replication > Pause**, start the application, and then resume replication using **Replication > Resume**.



Example—Migrating From a Healthy Primary

This example explains how to migrate the Primary role from the original Primary `seattle` to the Secondary host `london`. This example gives the RDS views before migration and after migration. Before the migration, the RDS view for `hr_rvg` looks like this:

The screenshot shows the Veritas Volume Replicator interface for RDS hr_rvg. The primary role is on the `seattle` host, and the secondary role is on the `london` host. The replication status is consistent and up-to-date.

Primary	State	Disk Group	Checkpoints
seattle	enabled for I/O	hrdj	0
Replication Log: Name: hr_srl Size: 4.00 GB Layout: CONCAT Mirrors: 0			

Secondary	Data Status	Replication Status	Logging To	Current Mode	Checkpoints
london	consistent, up-to-date	replicating (connected)	SRL ██████████ 0% full, 0	asynchronous	0

Name	Size	Layout	Mirrors	DCH
hr_dv01	4.00 GB	CONCAT	0	Yes
hr_dv02	4.00 GB	CONCAT	0	Yes
Total Size :8.00 GB				

Alerts
No Alerts

- ❖ For instructions on migrating the Primary role from `seattle` to `london`, see “[Migrating the Primary Role](#)” on page 95.



After migration, the RDS view for hr_rvg looks like this:

The screenshot shows the VERITAS Volume Replicator Web Console interface. The main content area displays the configuration for RDS hr_rvg. It is divided into several sections:

- Primary:** A table showing the primary role 'london' with state 'enabled for I/O', disk group 'hrdg', and 0 checkpoints. Below it, the 'Replication Log' shows 'Name: hr_sl', 'Size: 4.00 GB', 'Layout: CONCAT', and 'Mirrors: 0'.
- Replication Status:** A table showing the secondary role 'seattle' with a status of 'consistent, up-to-date replicating (connected)'. It also shows 'Logging To: SRL', 'Current Mode: asynchronous', and 'Checkpoints: 0'.
- Data Volumes:** A table listing two data volumes: 'hr_dv01' and 'hr_dv02', both 4.00 GB in size with a CONCAT layout and 0 mirrors. The total size is 8.00 GB.
- Alerts:** A section indicating 'No Alerts'.

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 selected 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. For more information on failing back, see [“Failing Back to the Original Primary”](#) on page 102.

Failback Options

The Takeover dialog box provides the following options to fail back to the original Primary. Select the required option depending on the method you plan to use for failback.

- ◆ **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 (**Change Role > Replay Failback Log**) on the new Primary.
- ◆ 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 and also automatically synchronizes 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.

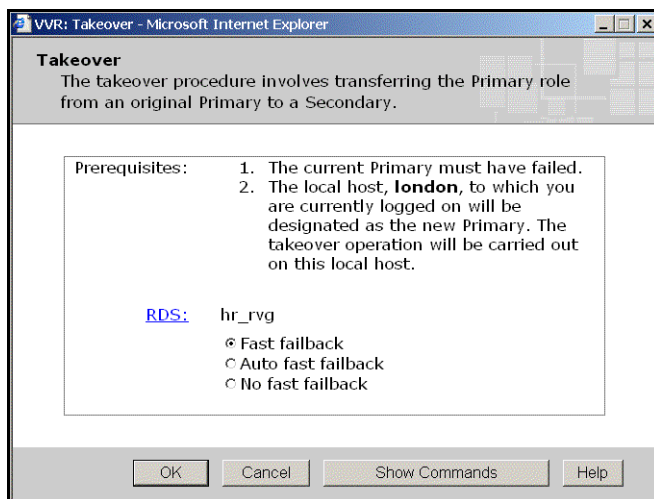


- ◆ If **No fast failback** is chosen, you must synchronize the original Primary after it becomes available either using difference-based synchronization or full synchronization. Use the Synchronize RVG task (**Tools > Synchronize RVG**) to synchronize the original Primary using difference-based synchronization or full synchronization.

Tip We recommend that you use the fast failback synchronization method to synchronize the original Primary with the new Primary. For instructions, see “[Failing Back Using Fast-Failback Synchronization](#)” on page 102.

▼ 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. From the Secondary RVG view on the Secondary host, choose **Change Role > Takeover**. The Takeover dialog box enables you to convert the selected Secondary to be the new Primary.



3. Complete the Takeover Options page as follows:

<p>Options: Fast failback Auto fast failback No fast failback</p>	<p>By default, Fast failback is selected. If this option is not appropriate, select the required option.</p> <ul style="list-style-type: none"> ◆ To enable fast failback logging only, select Fast failback. You can synchronize the data volumes on the original Primary and convert the original Primary to a Secondary using the Replay Failback Log wizard (Change Role > Replay Failback Log). ◆ To convert the original Primary to a Secondary after it becomes available, and also automatically use the DCM to synchronize the data volumes on the original Primary using fast failback, select Auto fast failback. ◆ To change the role from Secondary to Primary without enabling fast failback, select No 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.
---	---

4. Click **OK** to proceed with the Takeover operation. A message displays to indicate whether the takeover succeeded or failed.
5. Click **Close**.
6. 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.
7. 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.



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 It is recommended 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 (**Change Role > Replay Failback Log**). 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. For instructions, see “[Migrating the Primary Role](#)” on page 95.

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. From the RDS or RVG view on any host in the RDS choose **Change Role > Replay Failback Log**.
2. In the Replay Failback Log dialog box, select one of the following options:

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	If not using a cache.

3. Click **OK**. You can check the status of the synchronization using the RDS Summary view.

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.

When the synchronization completes, go to the next step.

4. At a convenient time, stop the application on the new Primary `london`.
5. Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle` using the Migration task (**Change Role > Migrate**). For instructions, see [“Migrating the Primary Role”](#) on page 95.

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



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 (**Change Role > Make Secondary**). For detailed instructions, see [“Converting a Primary to a Secondary”](#) on page 105.
2. Synchronize the data volumes on the original Primary with the data volumes on the new Primary using the difference-based synchronization and checkpoint option of the the Synchronize RVG task (**Tools > Synchronize RVG**). For detailed instructions, see [“Synchronize the Secondary RVG with the Primary RVG”](#) on page 106.
3. Start replication to the new Secondary with the checkpoint by completing the Start Replication task (**Replication > Start**). For detailed instructions, [“Synchronizing the Secondary and Starting Replication”](#) on page 25.
4. Stop the application on the new Primary `london`.
5. 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 95.



Converting a Primary to a Secondary

Use **Change Role > Make Secondary** 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

Perform the following steps to convert an original Primary to a Secondary after one of the existing Secondary hosts has become the new Primary:

1. Select the RDS or RVG view on the Primary host, which is the original Primary that failed and is now restarted.
2. From the RDS or RVG view on the Primary host, choose **Change Role > Make Secondary**.
3. Complete the Make Secondary page as follows:

New Primary	<p>From the drop-down list, select the name of the new Primary.</p> <p>If the original Primary no longer has a list of its original Secondary hosts, then enter the name in the field that is provided.</p> <p>If the name of the new Primary is present in the entry field, the value in the entry field takes precedence over the selected value in the selection box.</p>
-------------	--

4. Click **OK** to convert the original Primary to a Secondary. A message displays to indicate whether the takeover succeeded or failed.
5. Click **Close**.



Synchronize the Secondary RVG with the Primary RVG

Use **Tools > Synchronize RVG** to synchronize the data volumes on the Secondary with the data volumes on the Primary. The Synchronize RVG task is used in the failback procedure to synchronize the data volumes on the original Primary with the data volumes on the new Primary using difference-based synchronization and checkpoint.

▼ To synchronize the Secondary RVG with the Primary RVG

Perform the following steps to synchronize the data volumes on the Secondary with the data volumes on the Primary:

1. From the RDS or RVG view, choose **Tools > Synchronize RVG**.
2. Complete the Synchronize RVG dialog box as follows:

Option: Difference-based synchronization of RVG Full synchronization of RVG	Click the required synchronization option.
Checkpoint Name	Enter a name for the checkpoint.
Secondary Host	From the list, select the name of the Secondary host to synchronize.

3. Click **OK** to synchronize the data volumes on the Secondary with the data volumes on the Primary. VRW displays the status of the synchronization in a message box. A message displays when the synchronization completes.



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 **Change Role > Make Secondary**, and complete the Make Secondary page. Note: In the **New Primary** field, select or enter `london`. For detailed instructions, see [“Converting a Primary to a Secondary”](#) on page 105.

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

2. From the RDS or RVG view, choose **Tools > Synchronize RVG**, and complete the Synchronize RVG dialog box. Note: In the Synchronize RVG dialog box, select **Difference-based synchronization**, type the checkpoint name, and select `seattle` in the Secondary Host field. For detailed instructions, see [“Synchronize the Secondary RVG with the Primary RVG”](#) on page 106.

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` using **Replication > Start**.
4. In the Start Replication page, select the **checkpoint** option, and select the checkpoint `checkpt_presync` from the **checkpoint** field 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 **Change Role > Migrate**. For instructions, see [“Migrating the Primary Role”](#) on page 95. Replication from the original Primary `seattle` to the original Secondary `london` is started by default after the migration is completed.





Additional Settings for VRW

6

This chapter describes specific configuration and management tasks that can be performed on VRW after the Volume Replicator Web Console has been installed successfully. These procedures are generally not required, if the default VRW installation settings are adequate for your requirements. These tasks are available for making specific settings to VRW, if required.

You can also make changes to the VERITAS Web Server. Refer to “[Administering VERITAS Web Server](#)” on page 113 for more information.

Note This chapter assumes that `VRTSweb` and `VRTSvrw` are installed in the default directories `/opt/VRTSweb` and `/opt/VRTSvrw` respectively, that is, the installation base directory is root (`/`). If the `VRTSweb` or `VRTSvrw` is installed in a different base directory the paths would be `/<BaseDirPath>/opt/VRTSweb` and `/<BaseDirPath>/opt/VRTSvrw` respectively.

Configuring the VRW Application

The following parameters that can change the behavior of VRW are configurable:

- ◆ User session timeout
- ◆ The path of VRW log file `vrw.log`
- ◆ The maximum size of VRW log file `vrw.log`



Configuring VRW Using the Configuration File

The configuration properties are stored in VRW's configuration file. The VRW application uses the following configuration file:

File Name	Default Location
web.xml	/opt/VRTSweb/VERITAS/vvr/WEB-INF

In general, to change the default properties of VRW, edit the `web.xml` file:

1. Edit the appropriate fields in the configuration file as described in the following sections.
2. Save the changes.
3. Close the configuration file.
4. In order to see the changes in the VRW application, you must stop and then restart VRW. For more information on how to restart, see [“Starting and Stopping the VRW Application”](#) on page 111.

Editing the Configuration File

This section gives some examples of editing the configuration file.

User Session Timeout

```
<session-timeout>30</session-timeout>
```

The default idle timeout value for a user session is set to 30 minutes. This requires you to log in to VRW again after a session has been idle for 30 minutes. To disable the timeout, set the timeout value to -1.

Path of VRW Log File `vwr.log`

```
<context-param>  
  <param-name>LogFilePath</param-name>  
  <param-value>/var/vx/vrw/log/vrw.log</param-value>  
</context-param>
```

The default directory where the log file is stored is `/var/vx/vrw/log`, and the default name of the log file is `vrw.log`. For timely support, we recommend you not change the path or name of the log file.

Maximum Size of VRW Log File `vrw.log`

```
<context-param>
  <param-name>LogFileRolloverSize</param-name>
  <param-value>10485760</param-value><!-- 10MB -->
</context-param>
```

The maximum log file size is specified in bytes. The default size for the log file is set to 10 MB.

Starting and Stopping the VRW Application

In order to apply the changes to the configuration parameters of VRW, you need to stop and start VRW. The VRW application can be started and stopped without affecting the VERITAS Web Server and other running applications on that server.

1. Stop VRW using the following command:

```
# /opt/VRTSvrw/bin/vrw stop
```

2. Start using the following command:

```
# /opt/VRTSvrw/bin/vrw start
```





Administering VERITAS Web Server



VERITAS Web Server (VRTSweb) is a Web Server component shared by various VERITAS Web consoles, including VERITAS Cluster Server, VERITAS Volume Replicator, and VERITAS Traffic Director.

This document describes how to administer VRTSweb and provides instructions for common configuration tasks. Note that changes to the VRTSweb configuration apply to all Web consoles sharing the Web server.

Note The Web server is installed at the path `/opt/VRTSweb/` on UNIX systems. On Windows systems, the default installation path is `C:\Program Files\VERITAS\VRTSweb`.




Getting Started

1. Access the Web server using an existing port number, for example, `http://hostname:8181/`.
2. Accept the self-signed certificate (issued by VERITAS) to proceed.

You can prevent this certificate from appearing every time you connect to the console by installing a CA-signed certificate. See “[Configuring a CA-Signed SSL Certificate](#)” on page 122 for instructions.

3. The browser displays the **Home** and **Configuration** tabs.




To view and select the available VERITAS Web consoles, click the **Home** tab in the top left corner of the page.

To view and configure ports, SMTP recipients, SMTP servers, and logging, click **Configuration** in the top left corner of the page.

Reviewing the Web Server Configuration

1. Access the Web server using an existing port number; for example, `http://hostname:8181/`.
2. Click the **Configuration** tab.



The **Configured Ports** table lists information about the configured ports.

The **SMTP Recipients** table displays information about configured SMTP recipients and the SMTP server.

The **Logging** table lists the log levels for various Web server components.



Configuring Ports for VRTSweb

By default, VRTSweb is configured to serve HTML content on two ports: 8181 (HTTP) and 8443 (HTTPS). Additionally, VRTSweb uses port 14300 as an administrative port.

If you use any of these ports for another application on the system, you must configure VRTSweb to use different ports.

Port 8181 is the non-secure port, used for backward compatibility; 8443 is the secure SSL port. Users accessing the Web server on the non-secure port are redirected to the secure port.

When accessing content over the secure port, VRTSweb presents a self-signed SSL certificate (issued by VERITAS) to the browser. You must accept the certificate before accessing the secure Web consoles. The SSL protocol prevents malicious users from sniffing Web console data from the network.

Retrieving the List of Ports

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui listports
```

The output displays the list of configured ports and their protocols.

▼ From the Web Console

1. Access the Web server using an existing port number; for example, `http://hostname:8181/`.
2. Click the **Configuration** tab.

The **Configured Ports** table on the right side of the Configuration page lists the ports.

Adding Ports

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui addport portno protocol bind_ip_address
```

The variable *portno* represents the port number to be added. The variable *protocol* represents the protocol for the port. HTTP specifies a normal HTTP port, HTTPS specifies a secure SSL port.

Web servers using the HTTP port can be accessed at `http://hostname:portno/`.

Web servers using the HTTPS port can be accessed at `https://hostname:portno/`.

The optional variable *bind_ip_address* specifies that the new port be bound to a particular IP address instead of each IP address on the system. Use this option to restrict Web server access to specific administrative subnets. If specified, the IP address must be available on the system before the Web server is started. Otherwise, the Web server fails to start.

For example:

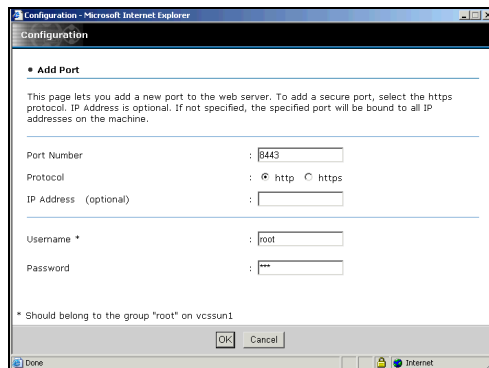
```
# /opt/VRTSweb/bin/webgui addport 443 HTTPS 101.1.1.2  
# /opt/VRTSweb/bin/webgui addport 80 HTTP.
```

▼ From the Web console

1. Access the Web server using an existing port number; for example, `http://hostname:8181/`.
2. Click the **Configuration** tab.
3. Click **Add Port** on the left side of the Configuration page.



4. In the Add Port dialog box:



- a. Enter the port number to be added.
- b. Choose the HTTP option to add a normal port; choose the HTTPS option to add a secure SSL port.

Web servers using the HTTP port can be accessed at `http://hostname:portno/`.

Web servers using the HTTPS port can be accessed at `https://hostname:portno/`.
- c. Enter an IP address to bind the new port to a specific IP address instead of each IP address on the system. Ensure the IP address is available on the system before starting the Web server. Use this attribute to restrict Web server access to specific administrative subnets.
- d. Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
- e. Click **OK**.

Deleting Ports

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui delport <portno> [bind_ip_address]
```

The variable *portno* represents the port number to be deleted. If the port was bound to a particular IP address, use the *bind_ip_address* option.

You must ensure that at least one port remains configured for the Web server.

For example:

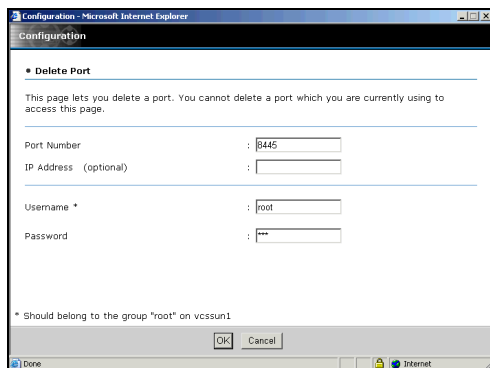
```
# /opt/VRTSweb/bin/webgui delport 443 101.1.1.2
# /opt/VRTSweb/bin/webgui delport 80
```

▼ From the Web console

1. Access the Web server using an existing port number; for example, `http://hostname:8181/`.
2. Click the **Configuration** tab.
3. Click **Delete Port** on the left side of the Configuration page.



4. In the Delete Port dialog box:



- a. Enter the port number to be deleted. You cannot delete the port being used to access the Web page.
- b. If the port was bound to a particular IP address, enter the IP address.
- c. Enter the name and password for a user having superuser (administrative) privileges on Web server system.
- d. Click OK.

Changing the Administrative Port

You can change the administrative port for VRTSweb only from the command line.

1. Stop the Web server:

```
# $VRTSWEB_HOME/bin/webgui stop force
```

2. Set the administrative port to a new value:

```
# $VRTSWEB_HOME/bin/webgui adminport new_port_no
```

3. Restart the Web server:

```
# $VRTSWEB_HOME/bin/webgui restart
```

Managing VRTSweb SSL Certificates

When serving content over the secure port, VRTSweb presents a self-signed SSL certificate (issued by VERITAS) to the browser. This section describes how you can manage the certificate.

Note Certificate management commands are available only via the command line interface. Commands that modify the certificate require a server restart. You can use the `webgui restart` command to restart the Web server.

Viewing SSL Certificate Information

To view information about the configured SSL certificate, run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui cert display
```

Creating a Self-Signed SSL Certificate

To create a custom self-signed SSL certificate for VRTSweb, run the following interactive command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui cert create
```

The command guides you through the process of creating a new certificate.

```
Please answer the following questions to create a self-signed SSL
certificate. This is required to enable the HTTPS protocol for the
web server.
+++++
With what hostname/IP will you access this web server?
[thor106]:thor106
What is the name of your organizational unit? [Unknown]:Engineering
What is the name of your organization? [Unknown]:Your Company
What is the name of your City or Locality? [Unknown]: Mountain View
What is the name of your State or Province? [Unknown]:California
What is the two-letter country code for this unit? [Unknown]:US
Is CN=thor106, OU=Engineering, O=Your Company, L=Mountain View,
ST=California, C=US correct? [no]:yes
Certificate created successfully
```

Note You must restart the server for the new certificate to take effect.



Exporting SSL Certificate to a File

You can export the public key associated with an SSL certificate to a file. This key can then be imported into other applications that will trust the VRTSweb instance.

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui cert export cert_file [rfc]
```

If the VRTSweb SSL certificate does not exist, the command prompts you to create one. If you specify the RFC option, the key output is encoded in a printable format, defined by the Internet RFC 1421 standard.

For example:

```
# /opt/VRTSweb/bin/webgui cert export /myapp/vrtsweb.cer rfc
```

Configuring a CA-Signed SSL Certificate

By default, VRTSweb presents a self-signed SSL certificate every time you access VRTSweb over the SSL port. You can install a certificate signed by a Certificate Authority (CA) like Verisign.com or Thawte.com.

1. If you do not have a self-signed certificate with information that can be verified by the CA, create one.

```
# $VRTSWEB_HOME/bin/webgui cert create
```

See [“Creating a Self-Signed SSL Certificate”](#) on page 121 for more information.

2. Generate a Certificate Signing Request (CSR) for the certificate. Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui cert certreq certreq_file
```

The variable *certreq_file* specifies the file to which the CSR will be written. The file is written using the Public-Key Cryptography Standard PKCS#10.

For example:

```
# /opt/VRTSweb/bin/webgui cert certreq /myapp/vrtsweb.csr
```

3. Submit the CSR to a certification authority, who will issue a CA-signed certificate.

4. Import the CA-issued certificate to VRTSweb. Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui import ca_cert_file
```

The variable `cert_file` represents the certificate issued to you by the certification authority.

For example:

```
# /opt/VRTSweb/bin/webgui cert import /myapp/vrtsweb.cer
```

Note that the import command fails if the CA root certificate is not a part of the trust store associated with VRTSweb. If the command fails, add the CA root certificate to the VRTSweb trust store:

```
# $VRTSWEB_HOME/bin/webgui cert trust ca_root_cert_file
```

For example:

```
# /opt/VRTSweb/bin/webgui cert trust /myapp/caroot.cer
```

Once the certificate used to sign the CSR is added to VRTSweb trust store, you can import the CA-assigned certificate into VRTSweb.

5. Restart VRTSweb:

```
# $VRTSWEB_HOME/bin/webgui restart
```

Cloning the VRTSweb SSL Certificate

You can clone the VRTSweb SSL keypair into a keystore and use the cloned VRTSweb certificate for another application or Web server. Visit <http://java.sun.com> for more information about keystores.

```
# $VRTSWEB_HOME/bin/webgui cert clone keystore storepass alias  
keypass
```

If a clone keystore exists, the command renames it to `keystore.old`. If the VRTSweb SSL certificate does not exist, the command prompts you to create one.

For example:

```
# /opt/VRTSweb/bin/webgui webgui cert clone  
/myapp/myserv.keystore mystorepass myalias mykeypass
```



Configuring SMTP Notification for VRTSweb

You can configure VRTSweb to send out email notifications about events associated with the Web server. For example:

- ◆ The Web server is starting/stopping [severity: INFORMATION]
- ◆ The Web console is starting/stopping [severity: INFORMATION]
- ◆ The Web server's allocated heap size very close to the maximum allowed [severity: SEVERE]

To send an email notification, VRTSweb needs to know the IP address or hostname of a configured SMTP server. The SMTP server address is also made available to all the Web consoles running on the Web server, thereby avoiding the need to configure the SMTP server at multiple places.

Retrieving the Name of the Configured SMTP Server

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui smtp getserver
```

The command displays the SMTP server address or hostname, if it is configured.

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. The **SMTP Recipients** table on the right side of the page displays the configured SMTP server.

Setting the SMTP Server

▼ From the command line

Run any of the following commands on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui smtp setserver server_ip/hostname
# $VRTSWEB_HOME/bin/webgui smtp delserver
```

The `setserver` command sets the SMTP server to the specified hostname/IP address. The `delserver` command deletes the SMTP server setting and disables SMTP notification.

For example:

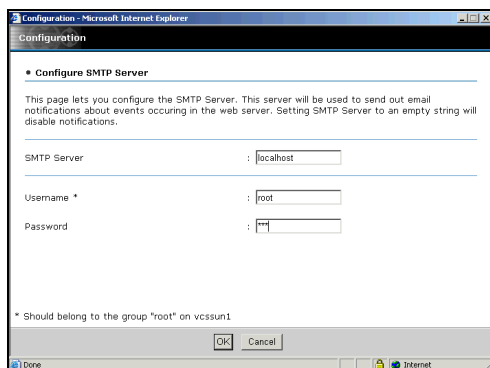
```
# /opt/VRTSweb/bin/webgui smtp setserver smtphost.company.com
# /opt/VRTSweb/bin/webgui smtp setserver 101.1.2.3
# /opt/VRTSweb/bin/webgui smtp delserver
```

▼ From the Web Console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. Click **Configure SMTP Server** on the left side of the Configuration page.



4. In the Configure SMTP Server dialog box:



- a. Enter the IP address or hostname of the SMTP server to be used for notification. An empty string will disable notification.
- b. Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
- c. Click OK.

Retrieving Configured SMTP Recipients

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui smtp listrcpt
```

This command retrieves the email addresses of the configured recipients, the notification severity level, and the notification locale.

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. The **SMTP Recipients** table on the right side of the Configuration page lists the configured SMTP recipients.

Adding an SMTP Recipient

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui smtp addrcpt email\  
[severity=<INFO|WARN|ERROR|SEVERE>] \  
[locale=<en|any_other_installed_locale>]
```

The variable *email* represents the email address of the new recipient.

The optional attribute *severity* represents the threshold for receiving Web server events. It can assume one of the following values: INFO|WARN|ERROR|SEVERE. If no value is specified for this attribute, it takes the default ERROR level.

The optional attribute *locale* specifies the locale in which the notification is to be sent. If no value is specified for this attribute, it takes the default locale of the system.

To retrieve the list of installed locales, run the following command:

```
# $VRTSWEB_HOME/bin/webgui smtp listlocales
```

For example:

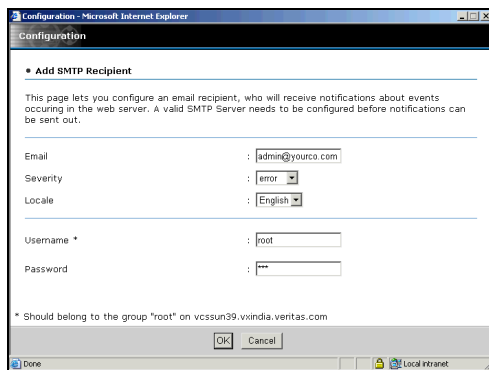
```
# /opt/VRTSweb/bin/webgui smtp addrcpt admin@company.com  
severity=INFO locale=ja_JP  
# /opt/VRTSweb/bin/webgui smtp addrcpt admin@company.com  
severity=ERROR  
# /opt/VRTSweb/bin/webgui smtp addrcpt admin@company.com
```

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. Click **Add SMTP Recipient** on the left side of the Configuration page.



4. In the Add SMTP Recipient dialog box:



- a. Enter the email address of the new recipient.
- b. From the **Severity** list, select the threshold for receiving Web server events. You can select one of the following values: INFO | WARN | ERROR | SEVERE.
- c. From the **Locale** list, select the locale in which notification is to be sent.
- d. Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
- e. Click **OK**.



Deleting an SMTP Recipient

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui smtp delrcpt email
```

The variable *email* represents the email address of the recipient to be deleted.

For example:

```
# /opt/VRTSweb/bin/webgui smtp delrcpt admin@company.com
```

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. Click **Delete SMTP Recipient** on the left side of the Configuration page.
4. In the Delete SMTP Recipient dialog box:

- a. Enter the email address of the recipient to be deleted.
- b. Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
- c. Click **OK**.



Configuring VRTSweb Logging

You can configure the amount of logs generated by individual VRTSweb components. VRTSweb comprises the following components:

- ◆ Web server
- ◆ Web applications
- ◆ Other components

You can set the logging threshold for each component separately. The lower the threshold, the more are the logs generated. VERITAS recommends setting log levels to lower values only for debugging.

Most of the logs are located at:

- ◆ `/var/VRTSweb/log` (for UNIX)
- ◆ `%VRTSWEB_HOME%\log` (for Windows),

Individual VERITAS Web consoles choose their own locations for their logs. See the documentation of the specific Web console for more information.

Retrieving Log Levels

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui log
```

This returns the logging thresholds for various components and the limit and rollover count of various log files for VRTSweb.

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. The **Logging** table on the right side of the Configuration page lists the log levels for various components of the Web server. Note that the table does not display the limit and rollover count of various log files; you must use the command line to retrieve this information.

Modifying Log Levels

▼ From the command line

Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui log [server=level] [webapps=level]  
[other=level]
```

You can specify any of the following values for the variable *level* for each Web server component: FINE | FINER | FINEST | CONFIG | INFO | WARNING | SEVERE.

Set the level to a lower value to generate more logs. FINEST is the lowest level while SEVERE is the highest level.

For example:

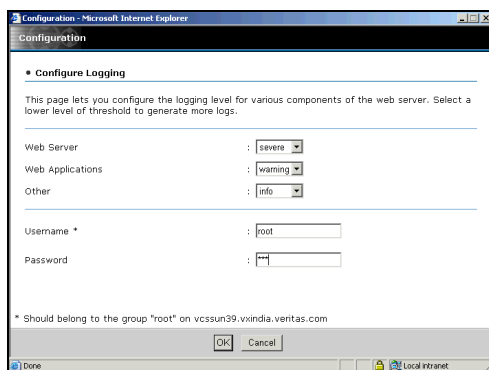
```
# /opt/VRTSweb/bin/webgui log server=FINEST webapps=INFO  
other=ERROR  
# /opt/VRTSweb/bin/webgui log server=INFO
```

▼ From the Web console

1. Access the Web server using an existing port number. For example, `http://hostname:8181/`
2. Click the **Configuration** tab.
3. Click **Configure Logging** on the left side of the Configuration page.



4. In the Configure Logging dialog box:



- a. Select the logging levels for the Web server, Web applications, and for other components.
- b. Enter the name and password for a user having superuser privileges on the Web server system.
- c. Click OK.

Modifying Maximum Size Limit and Rollover Count for Logs

You can modify the maximum size limit and rollover count for logs maintained by VRTSweb only from the command line. Run the following command on the system where VRTSweb is installed:

```
# $VRTSWEB_HOME/bin/webgui log
[vrtsweb_size=size]           [vrtsweb_count=count]
[command_size=size]          [command_count=count]
[binary_size=size]           [binary_count=count]
[jvm_size=size]               [jvm_count=count]
[protocol_client_size=size]   [protocol_client_count=count]
[protocol_server_size=size]   [protocol_server_count=count]
[out_size=size]               [out_count=count]
[err_size=size]               [err_count=count]
[webapps_size=size]          [webapps_count=count]
```

For example:

```
# /opt/VRTSweb/bin/webgui log vrtsweb_size=100000 vrtsweb_count=4
# /opt/VRTSweb/bin/webgui log err_size=200000
# /opt/VRTSweb/bin/webgui log webapps_count=4
```



The following table describes the command parameters:

Parameter	Description
vrtsweb_size	The size of the file <code>_vrtsweb.log</code> , which contains the Web server logs and the tomcat container related logs.
vrtsweb_count	The count for the file <code>_vrtsweb.log</code> .
command_size	The size of the file <code>_command.log</code> , which contains the logs related to administrative commands.
command_count	The count for the file <code>_command.log</code> .
binary_size	The size of the file <code>_binary.log</code> , which contains the binary representation of other log files.
binary_count	The count for the file <code>_binary.log</code> .
jvm_size	The size of the file <code>_jvm.log</code> , which contains JVM-related measurements. The file records memory consumed by the JVM at various times.
jvm_count	The count for the file <code>_jvm.log</code> .
protocol_client_size	The size of the file <code>_protocol_client.log</code> , which contains the communication sent (and received) by various utilities to the server.
protocol_client_count	The count for the file <code>_protocol_client.log</code> .
protocol_server_size	The size of the file <code>_protocol_server.log</code> , which contains the communication sent (and received) by the running server to various utilities.
protocol_server_count	The count for the file <code>_protocol_server.log</code> .
out_size	The size of the file <code>_out.log</code> , which contains messages logged to the standard output stream of the JVM.
out_count	The count for the file <code>_out.log</code> .
err_size	The size of the file <code>_err.log</code> , which contains messages logged to the standard error stream of the JVM, including any stack traces.
err_count	The count for the file <code>_err.log</code> .
webapps_size	The default size for log files of all Web applications running VRTSweb. Individual Web applications can override this default value.
webapps_count	The count for log files of all Web applications running VRTSweb. Individual Web applications can override this default value.



Modifying the Maximum Heap Size for VRTSweb

The default maximum allowed heap size for the VRTSWeb Java Virtual Machine (JVM) is 256MB. This prevents the Web server from increasing its memory footprint over the specified limit. However, for environments with a large number of VERITAS Web consoles sharing the same VRTSweb instance or with Web consoles managing large configurations, it may be necessary to modify this maximum limit.

You can modify the maximum heap size only from the command line.

```
# $VRTSWEB_HOME/bin/webgui maxheap new_size_in_MB
```

For example:

```
# /opt/VRTSweb/bin/webgui maxheap 512
```

You must restart the Web server after specifying a new limit.

```
# $VRTSWEB_HOME/bin/webgui restart
```

To display the current limit, run the command without specifying a new limit.

```
# $VRTSWEB_HOME/bin/webgui maxheap
```

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

Problem After Successfully Logging on to VVR

Message

Cannot communicate with vradmind on *hostname*; vradmind might be down or is not yet installed.

Description

The Summary page may display this error message even after you have successfully logged on to VRW.

Solution

This error indicates one of the following possibilities:

- ✓ You may not have a proper license for VVR.
Make sure that you have installed a valid VVR license.
- ✓ The `vradmind` command has been restarted and is running on a different port.
Restart the VRW application to use the new `vradmind` port. See [“Starting and Stopping the VRW Application”](#) on page 111.
- ✓ The `vradmind` daemon is not running on the system.
Start the `vradmind` daemon.



Application Page is Not Displayed in the Browser

Message

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

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/webgui listapps
Web server is offline
```

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

```
# /opt/VRTSvrw/bin/vrw 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 Close Button on Certain Wizards

Problem

The Close button may be not be available on some wizards and dialog windows.

Solution

This error can happen when you are using a version of Web browser that is not supported by VRW. Certain versions of Netscape 6.x have problems supporting JavaScript.

Correct RDS Name is not Displayed

Problem

The wizard may not display the correct name of the currently selected RDS. As a result, the wizard does not perform as expected.

Solution

This problem occurs if Internet Explorer (IE) is running on a PCI card in a Solaris environment. To resolve this problem do the following:

1. Select **Tools->Internet Options** from the IE menu bar.
2. Click the **Settings...** button.
3. In the Settings dialog box, click **Every visit to the page**, located under the **Check for newer versions of stored pages** heading.
4. Click **OK** to close the Settings window.
5. Click **OK** to close the Internet Options window.



Log in redirects the user back to the login page

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.

Accessibility and VRW



VERITAS products meet federal accessibility requirements for software as defined in Section 508 of the Rehabilitation Act:

- ◆ <http://www.access-board.gov/508.htm>

Keyboard shortcuts are available for all major graphical user interface (GUI) operations and menu items. VERITAS products are compatible with operating system accessibility settings as well as a variety of assistive technologies. All manuals are also provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

Navigation and Keyboard Shortcuts

VRW uses standard operating system navigation keys and keyboard shortcuts. For its unique functions, VRW uses its own navigation keys and keyboard shortcuts which are documented below.

Navigation in the Web Console

VRW supports standard browser-based navigation and shortcut keys for the following browsers:

- ◆ Internet Explorer 5.5 and 6.0
- ◆ Netscape Navigator 6.2 and 7.0

All VERITAS GUIs use the following keyboard navigation standards:

- ◆ Tab moves the focus to the next active area, field, or control, following a preset sequence. Shift+Tab moves the focus in the reverse direction through the sequence.
- ◆ Ctrl+Tab exits any Console area that you internally navigate with Tab.
- ◆ Up and Down arrow keys move focus up and down the items of a list.
- ◆ Alt in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.



- ◆ Either Enter or the Spacebar activates your selection. For example, after pressing Tab to select Next in a wizard panel, press the Spacebar to display the next page.

Support for Accessibility Settings

VERITAS software responds to operating system accessibility settings. On UNIX systems, you can change the accessibility settings using desktop preferences or desktop controls

Support for Assistive Technologies

VRW is compatible with IBM Home Page Reader version 3.0. It has been tested with Internet Explorer version 5.5 and IBM HPR version 3.0.

Index

A

- access modes, 54
- Add Secondary wizard, 20
- adding user roles, 53
- adding users, 57
- admin role, 54
- administrative tasks, 11
- Associate Volume wizard, 63
- automatic synchronization, 26
- AutoSync, 26

B

- bandwidth limit, changing, 70

C

- changing roles, 59
- checkpoint attach, using, 28
- checkpoint option, 27, 30
- checkpoints
 - creating, 28, 77
 - deleting, 78
 - ending, 28
 - viewing, 78
- configuration files, 110
- conventions, VRW, 11
- Create Primary wizard, 17
- creating
 - instant plex breakoff snapshot, 83
 - instant snapshot, 79
 - snapshots, overview, 78
 - space-optimized snapshot, 82
- creating a checkpoint, 28
- creating RVG snapshots, 78

D

- data volumes, dissociating, 68
- deleting users, 59
- detailed views, 7
- dissociate data volumes from an RDS, 68
- Dissociate Volume wizard, 67

E

- ending a checkpoint, 28
- errors, page display problem, 136
- example, scenario, 16
- examples
 - failing back using difference-based synchronization, 107
 - failing back using fast failback, 102

F

- failback
 - fast failback, 100
 - recommended method, 100
 - synchronizing data volumes, 106
 - using Make Secondary wizard, 105
- failback options, 99
- failing back
 - difference-based synchronization, 107
 - fast-failback synchronization, 102
 - overview, 102
- fast failback, 99
- force attach, 31
- full synchronization, 106

G

- guest role, 54

H

- heap size for VRTSweb, 134
- host list
 - see hosts, profiles*
- hosts
 - adding or removing in default profile, 59
 - adding or removing in user profile, 61
 - configuring VRW access, 54

I

- IP addresses, changing, 74

L

- latency high mark



- changing, 71
 - viewing, 50
- latency low mark
 - changing, 71
 - viewing, 50
- latency protection, changing, 70
- links, in views, 11
- logging
 - VRTSweb, 130
- logs
 - for VRTSweb, 130
- M**
- Make Secondary wizard, 105
- managing hosts with user profiles
 - , 54
- menu bar, 11
- menu options, 12
- Migrate wizard, 95
- multiple hosts, managing with user profiles, 54
- N**
- navigating VRW, 6, 11
- No fast-failback option, 100
- O**
- options, menu, 12
- P**
- packet size, changing, 70
- page display problems, 136
- Pause Replication wizard, 72
- pausing replication, 72
- Primary RVG
 - creating, 17
 - prerequisites for creating, 17
 - transferring role, 95
- Primary RVG view, 9
- Primary, deleting, 76
- profiles
 - configuring additional profiles, 55
 - default user profile, 55
- R**
- RDS
 - administering volumes, 63
 - associating volumes, 63
 - creating, 17
 - creating a Primary RVG, 17
 - pausing replication, 72
 - RDS view, using, 41
 - removing a Secondary, 75
 - removing data volumes, 67, 68
 - resizing data volumes, 66
 - status information, 35
 - stopping replication, 73
 - Synchronize RVG wizard, 106
 - transferring RVG roles, 95
- RDS states, Primary RVGs, 36
- RDS view, 8
- RDS, icons, 36
- Remove Primary wizard, 76
- Remove Secondary wizard, 75
- removing a Secondary RVG, 75
- replication
 - pausing, 72
 - resuming, 72
 - stopping, 73
- replication mode, changing, 70
- replication protocol, changing, 70
- replication settings
 - bandwidth limit, 70
 - packet size, 70
 - replication mode, 70
 - replication protocol, 70
 - SRL protection, 70
- replication settings, changing, 22
- replication, overview of setting up, 15
- Resize Volume wizard, 66
- resizing data volumes, 66
- Resume Replication wizard, 72
- RLINK, viewing information, 49
- roles, changing, 59
- root role, 53
- RVG snapshots, creating, 78
- RVG states, Primary RVGs, 36
- RVG, icons, 36
- S**
- Secondary RVG
 - pausing replication, 72
 - removing a Secondary, 75
 - stopping replication, 73
- Secondary RVG view, 10
- Secondary, adding to RDS, 20
- Secondary, removing, 75
- server port, 4
- SMTP notification, configuring for VRTSweb, 124

- SMTP server, retrieving name of, 124
- snapshots
 - overview, 78
- snapshots of RVG
 - creating
 - instant plex breakoff, 83
 - instant snapshot, 79
 - space-optimized, 82
 - displaying information, 86
- SRL protection, changing, 70
- Stop Replication wizard, 73
- Summary view, 6
- synchronization
 - difference-based, 104
 - full, 106
- Synchronize RVG wizard, 106
- synchronize Secondary, methods, 25
- synchronizing Secondary
 - automatic synchronization, 26

T

- Takeover wizard, 99
- Taking over from a failed Primary, 100
- tasks, VRW, 11
- Technical assistance, xiv

U

- user accounts, session timeout, 110
- user profiles, editing, 60
- user roles, 53
- user roles, adding, 53
- users, adding, 57
- users, deleting, 59
- using checkpoint option, 27, 30

V

- view-only mode, 54
- views
 - links, using, 11
 - Primary RVG view, 9
 - RDS view, 8
 - Secondary RVG view, 10
 - Summary view, 6
- volumes

- adding to Secondary RVG, 21
- administering, 63
- associating to an RDS, 63
- removing from an RDS, 67
- resizing, 66

VRTSweb

- adding ports, 117
- adding SMTP recipients, 127
- deleting ports, 119
- deleting SMTP recipients, 129
- logging, 130
- modifying log levels, 131
- notification for, 124
- ports for, 116
- retrieving log levels, 130
- retrieving ports, 116
- setting heap size, 134
- setting SMTP server, 125

VRW

- conventions, 11
- navigating, 6, 11
- server port, 4

VRW application

- restarting, 111
- stopping, 111

W

Web wizards

- Add Secondary, 20
- Associate Volume, 63
- Create Primary, 17
- Dissociate Volume, 67
- list of, 12
- Make Secondary, 105
- Migrate, 95
- Pause Replication, 72
- Remove Secondary, 75
- Resize Volume, 66
- Resume Replication, 72
- Stop Replication, 73
- Synchronize RVG, 106
- Takeover, 99



